

Please note that Cypress is an Infineon Technologies Company.

The document following this cover page is marked as “Cypress” document as this is the company that originally developed the product. Please note that Infineon will continue to offer the product to new and existing customers as part of the Infineon product portfolio.

Continuity of document content

The fact that Infineon offers the following product as part of the Infineon product portfolio does not lead to any changes to this document. Future revisions will occur when appropriate, and any changes will be set out on the document history page.

Continuity of ordering part numbers

Infineon continues to support existing part numbers. Please continue to use the ordering part numbers listed in the datasheet for ordering.



Traveo™ II Automotive Body Controller High Family Architecture Technical Reference Manual (TRM)

Document No. 002-24401 Rev. *D

September 21, 2020

Cypress Semiconductor
An Infineon Technologies Company
198 Champion Court
San Jose, CA 95134-1709
www.cypress.com
www.infineon.com

Copyrights

© Cypress Semiconductor Corporation, 2018-2020. This document is the property of Cypress Semiconductor Corporation and its subsidiaries ("Cypress"). This document, including any software or firmware included or referenced in this document ("Software"), is owned by Cypress under the intellectual property laws and treaties of the United States and other countries worldwide. Cypress reserves all rights under such laws and treaties and does not, except as specifically stated in this paragraph, grant any license under its patents, copyrights, trademarks, or other intellectual property rights. If the Software is not accompanied by a license agreement and you do not otherwise have a written agreement with Cypress governing the use of the Software, then Cypress hereby grants you a personal, non-exclusive, nontransferable license (without the right to sublicense) (1) under its copyright rights in the Software (a) for Software provided in source code form, to modify and reproduce the Software solely for use with Cypress hardware products, only internally within your organization, and (b) to distribute the Software in binary code form externally to end users (either directly or indirectly through resellers and distributors), solely for use on Cypress hardware product units, and (2) under those claims of Cypress's patents that are infringed by the Software (as provided by Cypress, unmodified) to make, use, distribute, and import the Software solely for use with Cypress hardware products. Any other use, reproduction, modification, translation, or compilation of the Software is prohibited.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CYPRESS MAKES NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS DOCUMENT OR ANY SOFTWARE OR ACCOMPANYING HARDWARE, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. No computing device can be absolutely secure. Therefore, despite security measures implemented in Cypress hardware or software products, Cypress shall have no liability arising out of any security breach, such as unauthorized access to or use of a Cypress product. CYPRESS DOES NOT REPRESENT, WARRANT, OR GUARANTEE THAT CYPRESS PRODUCTS, OR SYSTEMS CREATED USING CYPRESS PRODUCTS, WILL BE FREE FROM CORRUPTION, ATTACK, VIRUSES, INTERFERENCE, HACKING, DATA LOSS OR THEFT, OR OTHER SECURITY INTRUSION (collectively, "Security Breach"). Cypress disclaims any liability relating to any Security Breach, and you shall and hereby do release Cypress from any claim, damage, or other liability arising from any Security Breach. In addition, the products described in these materials may contain design defects or errors known as errata which may cause the product to deviate from published specifications. To the extent permitted by applicable law, Cypress reserves the right to make changes to this document without further notice. Cypress does not assume any liability arising out of the application or use of any product or circuit described in this document. Any information provided in this document, including any sample design information or programming code, is provided only for reference purposes. It is the responsibility of the user of this document to properly design, program, and test the functionality and safety of any application made of this information and any resulting product. "High-Risk Device" means any device or system whose failure could cause personal injury, death, or property damage. Examples of High-Risk Devices are weapons, nuclear installations, surgical implants, and other medical devices. "Critical Component" means any component of a High-Risk Device whose failure to perform can be reasonably expected to cause, directly or indirectly, the failure of the High-Risk Device, or to affect its safety or effectiveness. Cypress is not liable, in whole or in part, and you shall and hereby do release Cypress from any claim, damage, or other liability arising from any use of a Cypress product as a Critical Component in a High-Risk Device. You shall indemnify and hold Cypress, its directors, officers, employees, agents, affiliates, distributors, and assigns harmless from and against all claims, costs, damages, and expenses, arising out of any claim, including claims for product liability, personal injury or death, or property damage arising from any use of a Cypress product as a Critical Component in a High-Risk Device. Cypress products are not intended or authorized for use as a Critical Component in any High-Risk Device except to the limited extent that (i) Cypress's published data sheet for the product explicitly states Cypress has qualified the product for use in a specific High-Risk Device, or (ii) Cypress has given you advance written authorization to use the product as a Critical Component in the specific High-Risk Device and you have signed a separate indemnification agreement.

Cypress, the Cypress logo, Spansion, the Spansion logo, and combinations thereof, WICED, PSoC, CapSense, EZ-USB, F-RAM, and Traveo are trademarks or registered trademarks of Cypress in the United States and other countries. For a more complete list of Cypress trademarks, visit cypress.com. Other names and brands may be claimed as property of their respective owners.

Content Overview



Section A: Overview	21
1. Introduction	22
2. Getting Started	30
3. Document Construction	33
Section B: CPU Subsystem	36
4. CPU Subsystem (CPUSS)	37
5. Inter-Processor Communication	48
6. Protection Unit	53
7. Direct Memory Access	71
8. Code Flash	99
9. Work Flash	116
10. SRAM Interface	126
11. BootROM	133
12. Interrupts	144
13. Device Security	158
14. Chip Operational Modes	160
15. Fault Subsystem	162
Section C: System Resources Subsystem (SRSS)	166
16. Power Supply and Monitoring	167
17. Device Power Modes	186
18. Clocking System	197
19. Reset System	217
20. Watchdog Timer	222
21. Real-Time Clock	239
Section D: Input/Output Subsystem Overview	245
22. I/O System	246
Section E: Digital Subsystem	273
23. Serial Communications Block (SCB)	274
24. CAN FD Controller	328
25. Timer, Counter, and PWM	389
26. Local Interconnect Network (LIN)	455
27. Cryptography Block	480

28.	Event Generator (EVTGEN)	483
29.	Trigger Multiplexer	493
30.	FlexRay Controller	499
31.	Ethernet MAC	545
32.	Serial Memory Interface	581
33.	SDHC Host Controller	626
34.	Audio Subsystem	639
Section F: Analog Subsystem		653
35.	SAR ADC	654
Section G: Program and Debug Overview		683
36.	Program and Debug Interface	684
37.	Nonvolatile Memory Programming	694
38.	Flash Boot	740

Contents



Section A: Overview	21
1. Introduction	22
1.1 Device Characteristics	22
1.1.1 CPU Subsystem	22
1.1.2 Communication	22
1.1.3 Serial Memory Interface	23
1.1.4 Secure Digital High Capacity	23
1.1.5 Audio Subsystem	23
1.1.6 Miscellaneous	23
1.2 Top Level Architecture	24
1.2.1 CPU Subsystem	24
1.2.2 System Resources	25
1.2.3 Peripherals	27
1.2.4 I/Os	29
2. Getting Started	30
2.1 Support	30
2.2 Product Upgrades	30
2.3 Development Kits	30
2.3.1 Evaluation Board	30
2.3.2 Base Board	31
2.3.3 Sample Driver Library (SDL)	31
2.3.4 Development Tools	32
2.3.5 Cypress Programmer (CYP)	32
2.4 Application Notes	32
3. Document Construction	33
3.1 Major Sections	33
3.2 Documentation Conventions	33
3.2.1 Register Conventions	33
3.2.2 Numeric Naming	33
3.2.3 Units of Measure	34
3.2.4 Acronyms and Abbreviations	34
Section B: CPU Subsystem	36
4. CPU Subsystem (CPUSS)	37
4.1 Features	37
4.2 How It Works	37
4.3 Address Map	38
4.4 Operating Modes and Privilege Levels	38
4.5 Instruction Set	39
4.6 TCM Interface	39

4.6.1	TCM Read Wait State	39
4.6.2	TCM ECC	39
4.7	Cache ECC Fault Reporting	44
4.8	Registers	44
4.8.1	Arm Specification Registers	44
4.8.2	MCU Control Registers	47
5.	Inter-Processor Communication	48
5.1	Features	48
5.1.1	IPC Channel	48
5.1.2	IPC Interrupt	49
5.1.3	IPC Channels and Interrupts	49
5.2	Implementing Locks	50
5.3	Message Passing	50
5.4	Registers	52
6.	Protection Unit	53
6.1	Features	53
6.2	Configuration	54
6.2.1	Block Diagram	54
6.2.2	Protection Unit Structure	54
6.2.3	Master with Missing Access Attributes	55
6.3	Protection Context	55
6.3.1	Protection Context Configuration	55
6.3.2	Protection Context 0 and 1	56
6.4	Protection Structure	56
6.4.1	Address Region	56
6.4.2	Access Control Attributes	57
6.4.3	Protection Violation	58
6.4.4	Protection of Protection Structures	58
6.4.5	MPU	59
6.4.6	SMPU	59
6.4.7	PPU	60
6.4.8	Protection Structure Types	62
6.5	SWPU	65
6.5.1	SWPU Layout	65
6.5.2	SWPU Configuration	67
6.6	Registers	69
7.	Direct Memory Access	71
7.1	Peripheral DMA (P-DMA)	71
7.1.1	Overview	71
7.1.2	Channels	72
7.1.3	Descriptors	73
7.1.4	Interrupts	75
7.1.5	P-DMA Controller Status Registers	76
7.1.6	P-DMA Controller Design	76
7.1.7	Functionality	79
7.1.8	P-DMA Descriptor Structure	81
7.2	Memory DMA (M-DMA)	85
7.2.1	Overview	85
7.2.2	Channels	85
7.2.3	Descriptors	86

7.2.4	Interrupts.....	89
7.2.5	Control and Active Registers.....	89
7.2.6	M-DMA Controller Design	89
7.2.7	Examples	90
7.2.8	M-DMA Descriptor Structure	91
7.3	Registers.....	96
8.	Code Flash	99
8.1	Features.....	99
8.2	Configuration	99
8.2.1	Block Diagram.....	99
8.2.2	Flash Controller.....	100
8.2.3	Flash Geometry	108
8.2.4	Over-The-Air (OTA) Support	110
8.3	Operation	114
8.3.1	SROM APIs.....	114
8.4	Registers.....	114
9.	Work Flash	116
9.1	Features.....	116
9.2	Configuration	116
9.2.1	Block Diagram.....	116
9.2.2	Flash Controller.....	117
9.2.3	Flash Geometry	119
9.2.4	Over-the-Air (OTA) Support	121
9.2.5	Address Map of Work Flash.....	122
9.3	Operation	124
9.3.1	Read	124
9.3.2	SROM APIs.....	124
9.4	Registers.....	125
10.	SRAM Interface	126
10.1	Features.....	126
10.2	Configuration	127
10.2.1	Block Diagram.....	127
10.2.2	Wait States.....	127
10.2.3	Operation	128
10.2.4	Write Buffer	128
10.3	ECC Details	129
10.3.1	ECC Parity Generation for SRAM Write Accesses	129
10.3.2	ECC Syndrome Generation for SRAM Read Accesses.....	129
10.3.3	ECC Error Injection	130
10.3.4	ECC Parity Generation by Software	130
10.4	RAM Retention Configuration	131
10.5	Registers.....	132
11.	BootROM	133
11.1	Features.....	133
11.2	ROM Controller.....	133
11.2.1	Wait States.....	133
11.3	ROM Boot Process	134
11.3.1	Life-Cycle Stages and Protection States	134
11.3.2	Multicore Boot.....	134
11.3.3	Secure Boot	134

11.3.4	Protection Setting.....	134
11.3.5	Debug and Test Access Restrictions.....	138
11.3.6	SWD/JTAG Initialization.....	139
11.3.7	Waking up from Hibernate.....	139
11.3.8	Disable Watchdog Timer.....	139
11.3.9	ROM Boot Flow Chart.....	140
11.4	MMIO Registers and eFuse Used by ROM Boot.....	141
11.4.1	MMIO Registers.....	141
11.4.2	eFuse Bits.....	142
12.	Interrupts	144
12.1	Features.....	144
12.2	How It Works.....	145
12.3	Interrupts and Exceptions – Operation.....	146
12.3.1	Interrupt/Exception Handling.....	146
12.3.2	Level Interrupts.....	146
12.3.3	Exception Vector Table.....	147
12.4	Exception Sources.....	149
12.4.1	Reset Exception.....	149
12.4.2	Non-Maskable Interrupt Exception.....	149
12.4.3	HardFault Exception.....	149
12.4.4	Memory Management Fault Exception.....	149
12.4.5	Bus Fault Exception.....	149
12.4.6	Usage Fault Exception.....	149
12.4.7	Supervisor Call (SVCall) Exception.....	150
12.4.8	PendSV Exception.....	150
12.4.9	SysTick Exception.....	150
12.5	Interrupt Sources.....	150
12.6	Exception Priority.....	152
12.7	Enabling and Disabling Interrupts.....	152
12.8	Exception States.....	153
12.8.1	Pending Exceptions.....	153
12.9	Stack Usage for Exceptions.....	153
12.10	Interrupts and Low-Power Modes.....	154
12.11	Exception – Initialization and Configuration.....	154
12.12	Registers.....	155
13.	Device Security	158
13.1	Features.....	158
13.2	How It Works.....	158
13.2.1	Life-Cycle Stages.....	158
13.2.2	Memory and Peripheral Protection.....	159
13.2.3	Flash Write and eFuse Read/Write Protection.....	159
13.2.4	Hardware-based Cryptography.....	159
14.	Chip Operational Modes	160
14.1	Boot.....	160
14.2	User.....	160
14.3	Trusted.....	160
14.4	Debug.....	161
15.	Fault Subsystem	162
15.1	Fault Report Structure.....	163
15.2	Fault and Reset.....	164

15.3	Fault and Power Modes.....	164
15.4	Register List.....	165
Section C: System Resources Subsystem (SRSS)		166
16. Power Supply and Monitoring		167
16.1	Features.....	167
16.2	Power Supply.....	168
16.2.1	Core Regulators.....	169
16.2.2	Power Pins and Rails.....	169
16.2.3	Power Sequencing Requirements	169
16.2.4	Power Supply Sources.....	169
16.2.5	Usage of High-Current Regulator Controller.....	170
16.3	Voltage Monitoring.....	181
16.3.1	Power-On-Reset (POR).....	181
16.3.2	Brownout-Detection (BOD).....	181
16.3.3	Over-Voltage Detection (OVD).....	182
16.3.4	Low-Voltage-Detection (LVD).....	182
16.3.5	Over-Current Detection.....	183
16.3.6	Voltage Monitoring by ADC.....	183
16.4	Register List.....	185
17. Device Power Modes		186
17.1	Features.....	186
17.2	Device Power Modes.....	187
17.2.1	Active and Sleep Modes	187
17.2.2	DeepSleep Mode	188
17.2.3	Hibernate Mode	189
17.2.4	Other Operational States	189
17.3	Power Mode Transitions	190
17.3.1	Power-up Transitions	191
17.3.2	Low-Power Mode Transition	191
17.3.3	Wakeup.....	194
17.3.4	Internal Reset Transitions	194
17.3.5	Powering Down/Brownout/Overvoltage	194
17.3.6	Debugger Effect on Device Power Modes	194
17.4	Summary	195
17.5	Register List.....	196
18. Clocking System		197
18.1	Block Diagram	197
18.2	Clock Sources.....	199
18.2.1	Internal Main Oscillator (IMO).....	199
18.2.2	External Crystal Oscillator (ECO)	199
18.2.3	External Clock (EXT_CLK)	200
18.2.4	Internal Low-speed Oscillator (ILO)	200
18.2.5	Watch Crystal Oscillator (WCO).....	201
18.2.6	ECO Prescaler	201
18.2.7	LPECO.....	201
18.2.8	LPECO Prescaler.....	201
18.3	Clock Generation	201
18.3.1	PLL without SSCG and Fractional Operation	201
18.3.2	PLL with SSCG and Fractional Operation (400-MHz PLL)	202
18.3.3	Frequency Locked Loop (FLL).....	203

18.4	Clock Trees.....	205
18.4.1	Path Clocks.....	205
18.4.2	High-Frequency Root Clocks.....	206
18.4.3	Low-Frequency Root Clocks.....	206
18.4.4	Timer Clock.....	206
18.4.5	Clock Output Function.....	206
18.5	CLK_HF Distribution.....	207
18.5.1	CLK_MEM.....	207
18.5.2	CLK_SLOW.....	207
18.5.3	CLK_FAST_0 and CLK_FAST_1.....	207
18.5.4	CLK_PERI.....	207
18.5.5	PCLK.....	207
18.5.6	CLK_GR.....	207
18.5.7	CLK_TRC_DBG.....	207
18.6	Peripheral Clock Dividers.....	207
18.6.1	Fractional Clock Dividers.....	208
18.6.2	Peripheral Clock Divider Configuration.....	208
18.7	Clock Calibration Counters.....	210
18.8	Clock Supervision (CSV).....	210
18.8.1	Overview.....	210
18.8.2	CSV Operation.....	210
18.9	Registers.....	214
19.	Reset System	217
19.1	Reset Sources.....	218
19.1.1	Power-on Reset.....	218
19.1.2	Brownout Detection Reset.....	218
19.1.3	Over-Voltage Detection Reset.....	218
19.1.4	Over-Current Reset.....	219
19.1.5	External Reset.....	219
19.1.6	Programmable Reset.....	219
19.1.7	Watchdog Timer Reset.....	219
19.1.8	Internal System Reset.....	219
19.1.9	Fault Detection Reset.....	219
19.1.10	Clock-Supervision Reset.....	219
19.1.11	Hibernate Wakeup Reset.....	219
19.2	Identifying Reset Sources.....	219
19.3	Register List.....	221
20.	Watchdog Timer	222
20.1	Features.....	222
20.2	Block Diagram.....	223
20.3	Basic Watchdog Timer.....	223
20.3.1	Overview.....	223
20.3.2	Watchdog Reset.....	226
20.3.3	Watchdog Interrupt.....	227
20.4	Multi-Counter Watchdog Timer.....	228
20.4.1	Overview.....	228
20.4.2	How It Works.....	229
20.4.3	Enabling and Disabling MCWDT.....	234
20.4.4	Watchdog Reset.....	235
20.4.5	Watchdog Interrupt.....	235
20.5	Reset Cause Detection.....	236

20.6	Debug Mode	236
20.7	CPU Select	237
20.8	Register List	237
21.	Real-Time Clock	239
21.1	Features	239
21.2	Block Diagram	240
21.3	Power Supply	240
21.4	Clocking	240
21.5	Reset	241
21.6	Real-Time Clock	241
21.6.1	Reading RTC User Registers	242
21.6.2	Writing to RTC User Registers	242
21.7	WCO/LPECO Calibration	242
21.7.1	Absolute Accuracy Calibration	242
21.8	Alarm Feature	243
21.9	Backup Registers	244
21.10	Real Time Clock Registers	244
Section D:	Input/Output Subsystem Overview	245
22.	I/O System	246
22.1	Features	246
22.2	GPIO Interface Overview	246
22.3	I/O Cell Architecture	248
22.4	High Speed I/O (HSIO)	249
22.5	Digital Input Buffer	249
22.6	Digital Output Driver	250
22.6.1	Drive Modes	250
22.6.2	Slew Rate Control	252
22.7	High-Speed I/O Matrix	254
22.8	I/O State on Power Up	255
22.9	Behavior in Low-Power Modes	255
22.10	Interrupt	256
22.11	Peripheral Connections	257
22.11.1	Firmware-Controlled GPIO	257
22.11.2	Analog I/O	258
22.11.3	Serial Communication Block (SCB)	258
22.12	Smart I/O	258
22.12.1	Overview	258
22.12.2	Block Components	259
22.12.3	Routing	266
22.12.4	Operation	266
22.12.5	Example Application	267
22.13	Registers	271
Section E:	Digital Subsystem	273
23.	Serial Communications Block (SCB)	274
23.1	Features	274
23.2	Block Diagram	275
23.2.1	AHB-Lite Bus Interface	275
23.2.2	Trigger Interface	275
23.2.3	Serial Protocol Interfaces	276

23.2.4	Clock and Reset Interface.....	276
23.2.5	Block Enable.....	276
23.2.6	Interrupt Interface.....	276
23.3	Operation Modes.....	278
23.3.1	Buffer Modes.....	278
23.3.2	Clocking Modes.....	279
23.4	Serial Peripheral Interface (SPI).....	281
23.4.1	Features.....	281
23.4.2	General Description.....	281
23.4.3	SPI Modes of Operation.....	282
23.4.4	SPI Buffer Modes.....	286
23.4.5	Clocking and Oversampling.....	290
23.4.6	SPI Master SELECT Output Timing Control.....	293
23.4.7	SPI Parity Functionality.....	294
23.4.8	Loop-back.....	294
23.4.9	Enabling and Initializing SPI.....	294
23.4.10	I/O Pad Connection.....	295
23.4.11	SPI Registers.....	296
23.5	UART.....	297
23.5.1	Features.....	297
23.5.2	General Description.....	297
23.5.3	UART Modes of Operation.....	297
23.5.4	Clocking and Oversampling.....	307
23.5.5	Loop-back.....	307
23.5.6	Enabling and Initializing UART.....	307
23.5.7	I/O Pad Connection.....	308
23.5.8	UART Registers.....	310
23.6	Inter Integrated Circuit (I2C).....	311
23.6.1	Features.....	311
23.6.2	General Description.....	311
23.6.3	Terms and Definitions.....	311
23.6.4	I2C Modes of Operation.....	312
23.6.5	I2C Buffer Modes.....	314
23.6.6	Clocking and Oversampling.....	317
23.6.7	Loop-back.....	320
23.6.8	Enabling and Initializing the I2C.....	320
23.6.9	I/O Pad Connections.....	321
23.6.10	I2C Registers.....	321
23.7	SCB Interrupts.....	322
23.7.1	SPI Interrupts.....	322
23.7.2	UART Interrupts.....	323
23.7.3	I2C Interrupts.....	323
23.8	Registers.....	325
23.8.1	SPI Registers.....	325
23.8.2	UART Registers.....	326
23.8.3	I2C Registers.....	326
24.	CAN FD Controller	328
24.1	Overview.....	328
24.1.1	Features.....	328
24.1.2	Features Not Supported.....	329
24.2	Configuration.....	329
24.2.1	Block Diagram.....	329

24.2.2	Dual Clock Sources	329
24.2.3	Interrupt Lines	329
24.3	Functional Description	330
24.3.1	Operation Modes	330
24.3.2	Timestamp Generation	335
24.3.3	Timeout Counter	336
24.3.4	RX Handling	337
24.3.5	TX Handling	343
24.3.6	FIFO Acknowledge Handling	346
24.3.7	Configuring the CAN Bit Timing	346
24.4	Message RAM	350
24.4.1	Message RAM Configuration	350
24.4.2	RX Buffer and FIFO Element	350
24.4.3	TX Buffer Element	352
24.4.4	TX Event FIFO Element	353
24.4.5	Standard Message ID Filter Element	354
24.4.6	Extended Message ID Filter Element	356
24.4.7	Trigger Memory Element	357
24.4.8	ECC for Message RAM	359
24.4.9	Message RAM OFF	360
24.4.10	RAM Watchdog (RWD)	361
24.5	TTCAN Operation	361
24.5.1	Reference Message	361
24.5.2	TTCAN Configuration	362
24.5.3	TTCAN Gap Control	367
24.5.4	Stop Watch	368
24.5.5	Local Time, Cycle Time, Global Time, and External Clock Synchronization	368
24.5.6	Synchronization Triggers	370
24.5.7	TTCAN Error Level	371
24.5.8	TTCAN Message Handling	372
24.5.9	TTCAN Interrupt and Error Handling	373
24.5.10	Level 0	374
24.5.11	Synchronization to External Time Schedule	375
24.6	Setup Procedures	376
24.6.1	General Program Flow	376
24.6.2	Clock Stop Request	376
24.6.3	Message RAM OFF Operation	376
24.6.4	Message RAM ON Operation	376
24.6.5	Consolidated Interrupt Handling	377
24.6.6	Procedures Specific to M_TTCAN Channel	377
24.7	Registers	386
25.	Timer, Counter, and PWM	389
25.1	Features	389
25.2	Block Diagram	390
25.2.1	Enabling and Disabling Counters in TCPWM Block	390
25.2.2	Clocking	390
25.2.3	Trigger Inputs	391
25.2.4	Synchronization of Multiple Counters	394
25.2.5	Trigger Outputs	396
25.2.6	Internal Events	397
25.2.7	Interrupts	400
25.2.8	Debug Mode	400

25.2.9	PWM Outputs.....	400
25.2.10	Power Modes	403
25.3	Operation Modes	403
25.3.1	Timer Mode.....	404
25.3.2	Capture Mode	411
25.3.3	Quadrature Decoder Mode	414
25.3.4	Pulse Width Modulation (PWM) Mode	425
25.3.5	Pulse Width Modulation with Dead Time Mode	443
25.3.6	Pulse Width Modulation Pseudo-Random Mode (PWM PR)	445
25.3.7	Shift Register (SR).....	450
25.4	Design Configuration Parameters.....	453
25.5	Recovery.....	453
25.6	Initialize.....	453
25.7	Pin Status	453
25.8	TCPWM Registers	454
26.	Local Interconnect Network (LIN)	455
26.1	Features.....	455
26.1.1	LIN	455
26.1.2	UART	455
26.2	Block Diagram	456
26.2.1	Internal Bus Interface.....	456
26.2.2	Test Registers	456
26.2.3	LIN Channel	456
26.3	Clocking.....	456
26.3.1	Baud Rate and Sample Point.....	456
26.4	LIN Message Frame Format.....	458
26.4.1	Break and Synchronization Fields	458
26.4.2	PID Field	459
26.4.3	Data Fields.....	459
26.4.4	Checksum Field	460
26.5	Timeout Operation	460
26.6	Wakeup.....	460
26.6.1	Wakeup Signal Transmission.....	460
26.6.2	Wakeup Signal Reception.....	460
26.6.3	Wake up in Low Power Mode	461
26.7	External Transceiver Control	461
26.8	Test Modes	461
26.8.1	Interrupt Test.....	461
26.8.2	Loop-back Mode	461
26.8.3	Error Injection Mode.....	462
26.9	Operation	463
26.9.1	LIN Operation.....	463
26.9.2	UART Operation	467
26.10	Noise Filter.....	467
26.10.1	Example	467
26.11	Interrupts.....	470
26.11.1	Overview	470
26.11.2	Transmission Interrupts	473
26.11.3	Reception Interrupts.....	474
26.11.4	Error and Status Interrupts.....	476
26.12	Registers	479

27. Cryptography Block	480
27.1 Features Overview	480
27.2 System Diagram	481
27.3 Block Diagram	481
27.4 Function Description	482
27.4.1 Operating Mode	482
27.4.2 Memory Map and Register Definitions	482
27.4.3 Instruction Set	482
28. Event Generator (EVTGEN)	483
28.1 Features	483
28.2 Block Diagram	483
28.2.1 Enabling and Disabling EVTGEN Block	484
28.2.2 Counters	484
28.2.3 Counter Status	485
28.2.4 Comparator Structures	486
28.2.5 Interrupts	488
28.2.6 Use Case	489
28.2.7 Register List	490
29. Trigger Multiplexer	493
29.1 Features	493
29.2 Description	493
29.3 Trigger Multiplexing	494
29.4 Trigger Functionality	496
29.5 Registers	498
30. FlexRay Controller	499
30.1 Overview	499
30.2 Configuration	500
30.2.1 Block Diagram	500
30.3 Operations	501
30.3.1 Communication Cycle	501
30.3.2 Clock Synchronization	503
30.3.3 Error handling	505
30.4 Communication Controller States	506
30.4.1 Communication Controller State	506
30.4.2 DEFAULT_CONFIG State	508
30.4.3 CONFIG state	508
30.4.4 MONITOR_MODE	508
30.4.5 READY State	508
30.4.6 WAKEUP State	508
30.4.7 STARTUP State	511
30.4.8 NORMAL_ACTIVE State	515
30.4.9 NORMAL_PASSIVE State	516
30.4.10 HALT State	516
30.5 Network Management	516
30.6 Filtering and Masking	517
30.6.1 Slot Counter Filtering	517
30.6.2 Cycle Counter Filtering	517
30.6.3 Channel ID Filtering	518
30.6.4 FIFO Filtering	518
30.7 Transmission Procedure	518

30.7.1	Static Segment.....	518
30.7.2	Dynamic Segment.....	518
30.7.3	Transmission Buffer	518
30.7.4	Frame Transmission	519
30.7.5	Null Frame Transmission	519
30.8	Reception Procedure	519
30.8.1	Reception Buffer	519
30.8.2	Frame Reception	520
30.8.3	Null Frame Reception	520
30.9	FIFO Function.....	520
30.9.1	Details	520
30.9.2	FIFO Settings.....	521
30.9.3	Access to FIFO	521
30.10	Message Handling	521
30.10.1	Message Buffer Reconfiguration.....	521
30.10.2	Host Access to Message RAM.....	522
30.10.3	FlexRay Protocol Controller Access to Message RAM	526
30.11	Message RAM	527
30.11.1	Header Partition	528
30.11.2	Data Partition	529
30.11.3	Message Buffer Assignment	530
30.11.4	Parity Check.....	530
30.11.5	Parity Error Handling.....	532
30.12	Interrupts.....	533
30.12.1	Error and Status Interrupts.....	533
30.13	Timers and Stop Watch	533
30.13.1	Timer 0	533
30.13.2	Timer 1	533
30.13.3	Stop Watch	533
30.14	Test Modes	533
30.14.1	Asynchronous Transmit Mode (ATM).....	533
30.14.2	Loop Back Mode	534
30.14.3	RAM Test Mode	534
30.14.4	I/O Test Mode.....	534
30.14.5	Additional Status Information	534
30.15	Traveo II Specific Functions	535
30.15.1	Enable/Disable FlexRay Controller	535
30.15.2	Timer 0 Trigger Output.....	535
30.15.3	Stop Watch Event Trigger Input	535
30.15.4	DMA Trigger Interface for Input/Output Buffer Access	535
30.16	FlexRay Registers	542
31.	Ethernet MAC	545
31.1	Overview	545
31.1.1	Supported Features	545
31.2	Block Diagram	546
31.3	Ethernet MAC Operation	546
31.3.1	DMA Interface	546
31.3.2	Transmit Scheduling Algorithm	556
31.3.3	MAC Transmitter	558
31.3.4	MAC Receiver.....	558
31.3.5	Checksum Offload for IP, TCP, and UDP	558
31.3.6	Jumbo Frame Support	559

31.3.7	MAC Filtering Block	560
31.3.8	IEEE 1588 and IEEE 802.1AS Support	562
31.3.9	MAC 802.3 Pause Frame Support	565
31.3.10	MAC PFC Priority-based Pause Frame Support	566
31.3.11	Energy Efficient Ethernet Support	567
31.3.12	MDIO Interface	568
31.3.13	Interrupts	568
31.3.14	Media Independent Interfaces	568
31.3.15	Clocks to EMAC	573
31.3.16	Power Modes	573
31.4	Register List	574
32.	Serial Memory Interface	581
32.1	Features	581
32.2	System Diagram	582
32.3	Block Diagram	582
32.3.1	Clocks	583
32.4	Functional Description	584
32.4.1	Operating Modes	584
32.4.2	Off-chip Interfaces	587
32.4.3	AXI Interface	606
32.4.4	Triggers	610
32.4.5	Interrupts	610
32.4.6	Monitor Signals	610
32.5	Supply Rails and Power Domains	611
32.5.1	Power Modes	611
32.6	Sub Block Descriptions	611
32.6.1	Address Space	611
32.6.2	TX and RX FIFOs	612
32.6.3	Interrupts and Triggers	614
32.6.4	Cache	615
32.6.5	Arbitration	616
32.6.6	Cryptography	616
32.6.7	Serial Memory Interface Logic	619
32.7	SMIF Registers	624
33.	SDHC Host Controller	626
33.1	Features	626
33.1.1	Features Not Supported	626
33.2	Block Diagram	627
33.3	Clocking	628
33.3.1	Clock Gating	628
33.3.2	Base Clock (CLK_HF _x) Configuration	628
33.3.3	Card Clock (SDCLK) Configuration	628
33.3.4	Timeout (TOUT) Configuration	628
33.4	Bus Speed Modes	629
33.5	Power Modes	629
33.5.1	Standby Mode	629
33.6	Interrupts to CPU	629
33.6.1	SDIO Interrupt	630
33.7	I/O Interface	630
33.8	Packet Buffer SRAM	631
33.8.1	Packet Buffer Full/Empty	631

33.9	DMA Engine.....	631
33.10	Initialization Sequence.....	632
33.10.1	Enabling SDHC.....	632
33.10.2	Card Detection.....	633
33.10.3	SDHC Initialization.....	635
33.10.4	Clock Setup.....	635
33.11	Error Detection.....	635
33.12	Register List.....	636
34.	Audio Subsystem	639
34.1	Features.....	639
34.2	Architecture.....	640
34.3	Digital Audio Interface Formats	641
34.3.1	Standard I2S Format.....	641
34.3.2	Left Justified (LJ) Format	643
34.3.3	Time Division Multiplexed (TDM) Format.....	643
34.4	Clocking Polarity and Delay Options	644
34.5	Interfacing with Audio Codecs	646
34.6	Clocking Features.....	646
34.7	FIFO Buffer and DMA Support	648
34.8	Interrupt Support.....	650
34.9	Watchdog Timer	651
34.10	MCLK Output Function	651
34.11	Register List.....	652
Section F:	Analog Subsystem	653
35.	SAR ADC	654
35.1	Features.....	654
35.2	Block Diagram	656
35.3	Operation	656
35.3.1	SAR ADC Conversion Flow	657
35.3.2	Result Data Format.....	658
35.3.3	Acquisition/Sample Time.....	659
35.4	SARMUX	660
35.4.1	Preconditioning	660
35.4.2	Overlap Diagnostic.....	661
35.4.3	SARMUX Diagnostics	661
35.5	SAR Sequencer	662
35.5.1	Analog Input Selection	662
35.5.2	External Analog Multiplexer	662
35.5.3	Port Selection.....	663
35.5.4	Averaging.....	663
35.5.5	Range Detect.....	663
35.5.6	Pulse Detect.....	664
35.5.7	Double Buffer	666
35.5.8	Group Coherency.....	667
35.5.9	Status.....	667
35.6	Triggering and Scheduling.....	667
35.6.1	Channel Grouping.....	667
35.6.2	Triggers	668
35.6.3	Arbitration, Preemption, and Acquisition Scheduling	669
35.6.4	Debug Freeze	671
35.6.5	Auto Idle Power Down	671

35.6.6	Channel Disable/Software Abort	671
35.7	Output Triggers and Interrupts	671
35.7.1	Trigger Outputs	671
35.7.2	Channel Interrupts	672
35.7.3	Group Interrupts	672
35.8	Calibration	672
35.8.1	Analog Calibration	672
35.8.2	Alternate Calibration	673
35.8.3	Coherent Calibration Update	673
35.9	Temperature Measurement	675
35.9.1	Example Measurement Flow	675
35.9.2	Temperature Sensor Calibration and SFlash Address	675
35.9.3	Temperature Calculation	677
35.10	Diagnostic Reference Generator	679
35.10.1	Diagnostic Reference Configuration	679
35.11	Reference Buffer	679
35.12	Registers	681
Section G: Program and Debug Overview		683
36. Program and Debug Interface		684
36.1	Features	684
36.2	Functional Description	684
36.2.1	Debug Access Port (DAP)	685
36.2.2	ROM Tables	685
36.2.3	Trace	686
36.2.4	Embedded Cross-Triggering	686
36.3	Serial Wire Debug (SWD) Interface	687
36.3.1	SWD Timing Details	688
36.3.2	ACK Details	688
36.3.3	Turnaround (Trn) Period Details	689
36.4	JTAG Interface	689
36.5	Programming the TVII-B-H Device	692
36.5.1	SWD Port Acquisition	692
36.5.2	SWD Programming Mode Entry	692
36.5.3	SWD Programming Routines Executions	692
36.6	Registers	693
37. Nonvolatile Memory Programming		694
37.1	Functional Description	694
37.2	System Call Implementation	696
37.2.1	System Call via CM0+ or CM7_0 or CM7_1	696
37.2.2	System Call via DAP	696
37.2.3	Exiting from a System Call	696
37.3	SROM API Library	697
37.4	System Calls	699
37.4.1	BlankCheck	699
37.4.2	BlowFuseBit	700
37.4.3	Calibrate	701
37.4.4	CheckFactoryHash	702
37.4.5	CheckFMStatus	703
37.4.6	Checksum	704
37.4.7	ComputeBasicHash	706
37.4.8	ConfigureFMInterrupt	707

37.4.9	EnterFlashMarginMode	708
37.4.10	EraseAll	708
37.4.11	EraseResume	709
37.4.12	EraseSector	710
37.4.13	EraseSuspend	711
37.4.14	ExitFlashMarginMode	712
37.4.15	GenerateHash	713
37.4.16	SwitchOverRegulators	714
37.4.17	ConfigureRegulator	715
37.4.18	ProgramRow	717
37.4.19	ProgramWorkFlash	719
37.4.20	ReadFuseByte	720
37.4.21	ReadFuseByteMargin	721
37.4.22	ReadSWPU	722
37.4.23	ReadUniqueID	722
37.4.24	SetEnforcedApproval	723
37.4.25	SiliconID	724
37.4.26	SoftReset	726
37.4.27	TransitiontoRMA	727
37.4.28	TransitiontoSecure	727
37.4.29	DirectExecute	728
37.4.30	WriteRow	730
37.4.31	WriteSWPU	732
37.4.32	DebugPowerUpDown	734
37.4.33	LoadRegulatorTrims	735
37.4.34	OpenRMA	736
37.5	System Call Status	737
37.6	eFuse Memory	739
37.6.1	Features	739
37.6.2	Customer eFuses	739
38.	Flash Boot	740
38.1	Features	740
38.2	Using Flash Boot	741
38.2.1	Flash Boot Shared Functions	741
38.2.2	Using a Bootloader	742
38.3	Flash Boot Internals	747
38.3.1	Definitions	747
38.3.2	SFlash Address Mapping	748
38.3.3	Flash Boot Flow	749
38.3.4	Data Structures	754
38.3.5	Internal Bootloader	759

Section A: Overview



This section encompasses the following chapters:

- [Introduction chapter on page 22](#)
- [Getting Started chapter on page 30](#)
- [Document Construction chapter on page 33](#)

1. Introduction



The Traveo™ II Body High (TVII-B-H) device is a Traveo II microcontroller targeted at automotive systems such as high-end body control units. These devices have two Arm® Cortex®-M7 CPUs for primary processing, and a Cortex-M0+ CPU for peripheral and security processing. These devices contain embedded peripherals supporting Controller Area Network with Flexible Data rate (CAN FD), Local Interconnect Network (LIN), Gigabit Ethernet, and FlexRay. Traveo II is manufactured on an advanced 40-nm process. TVII-B-H incorporates Cypress' low-power flash memory, multiple high-performance analog and digital peripherals, and enables the creation of a secure computing platform.

1.1 Device Characteristics

1.1.1 CPU Subsystem

- Two 350-MHz 32-bit Arm Cortex-M7 CPUs, each with
 - Single-cycle multiply
 - Single/double-precision floating point unit (FPU)
 - 16-KB data cache, 16-KB instruction cache
 - Memory protection unit (MPU)
 - 16-KB instruction and 16-KB data Tightly-Coupled Memory (TCM)
- 100-MHz 32-bit Arm Cortex M0+ CPU with single-cycle multiply and MPU
- Interprocessor communication in hardware
- Three DMA controllers – two controllers to support peripheral-to-memory (and vice versa) and one for memory-to-memory data transfers
- Up to 8384 KB of code-flash with an additional, up to 256 KB of work-flash and an internal SRAM of up to 1024 KB
 - Flash programming on JTAG/SWD interface
 - Read-While-Write (RWW) allows updating the code-flash and work-flash while executing from it
 - Single- and dual-bank modes (specifically for Firmware Over-The-Air (FOTA) update)
- Crypto engine to support enhanced Secure Hardware Extension (eSHE) and Hardware Secure Module (HSM). The crypto engine and software support the following functions:
 - RSA-3072, ECC-256, ECC-384, SHA-2, SHA-3, AES-128/256, and 3DES
 - True random number generator (TRNG) and pseudo random number generator (PRNG)
 - Hash function
 - Galois/Counter Mode (GCM)
- Hardware error correction (SECDED ECC) on all safety-critical memories (SRAM and flash)

1.1.2 Communication

- High-speed CAN FD communication supporting up to 8 Mbps data rate
- Serial interface to support various serial communication (UART/SPI/I²C)
- LIN master/slave support by hardware compliant with ISO 17987
- 10/100/1000 Mbps Ethernet MAC interfaces with Audio Video Bridging (AVB) and Precision Time Protocol (PTP) support conforming to IEEE-802.3az. These interfaces support the following PHY interfaces:

- ❑ Media-independent interface (MII)
- ❑ Reduced media-independent interface (RMII)
- ❑ Gigabit media-independent interface (GMII)
- ❑ Reduced gigabit media-independent interface (RGMII)
- FlexRay interface (V2.1) configurable for single or dual data-channels for fault tolerance, supporting data rates up to 10 Mbps

1.1.3 Serial Memory Interface

- Supports SPI (single, dual, quad, or octal) or HyperBus interface, with on-the-fly encryption and decryption along with execute-in-place (XIP)

1.1.4 Secure Digital High Capacity

- Secure digital high capacity (SDHC) interface supporting embedded MultiMediaCard (eMMC), or secure digital (SD), or secure digital input output (SDIO) compliant to eMMC 5.1, SD 6.0, and SDIO 4.10 specifications
- Supports data rates up to SD SDR50 (single data rate) or eMMC 52 MHz DDR

1.1.5 Audio Subsystem

- Inter-IC sound (I²S) interfaces based on NXP I²S bus specifications are supported to connect digital audio devices
- Supports I²S, left justified, or time division multiplexed (TDM) audio formats
- Supports independent transmitter or receiver operation, each in master or slave mode


1.1.6 Miscellaneous

- Low-power 2.7-V to 5.5-V operation, with two robust brownout detect (BOD) and over-voltage detect (OVD) options
- Programmable GPIOs, and Smart I/O to perform Boolean operations on signals going to and from I/O pins
- High-performance 12-bit analog-to-digital converter (ADC)
 - ❑ Supports 12-bit resolution and sampling rates up to 1 Msps
- High-current linear regulator
 - ❑ Supports generation of 1.1-V nominal core supply from a 2.7-V to 5.5-V input supply
 - ❑ Supports control of external low-dropout (LDO) or power management IC (PMIC) regulators
- Hardware watchdog function
- Real-time clock with auto-calibration
- Timing and pulse-width modulation with support for timer, capture, quadrature, pulse-width modulation (PWM outputs), PWM with dead time (PWM_DT), pseudo-random PWM (PWM_PR), and shift-register (SR) modes; some PWM channels also support stepper motor control
- DeepSleep and Hibernate power modes for low-power solution
- Event generator to support cyclic wakeup from DeepSleep mode and peripheral trigger in active power mode
- AEC-Q100 qualification for all temperature range
- ASIL-B level functional safety
- Debugging via JTAG controller (interface compliant IEEE-1149.1-2001) and Arm SWD port
- Supports Arm Embedded Trace Macrocell (ETM) Trace
 - ❑ Data trace using SWD
 - ❑ Instruction and data trace using JTAG

1.2 Top Level Architecture

Figure 1-1 shows the TVII-B-H architecture block diagram, giving a simplified view of the interconnection between subsystems and blocks. The device has four major subsystems: CPU, system resources, peripherals, and I/O.

Figure 1-1. TVII-B-H Architecture Diagram



The device provides extensive support for programming, testing, debugging, and tracing of both hardware and firmware. Debug-on-chip functionality enables in-system debugging using the production device. It does not require special interfaces, debugging pods, simulators, or emulators. The JTAG interface is fully compatible with industry-standard third-party probes such as I-jet, J-Link, and GHS. The debug circuits are enabled by default. The microcontroller provides a high level of security with robust flash protection and the ability to disable features such as debug. Additionally, each device interface can be permanently disabled for applications concerned with phishing attacks due to a maliciously reprogrammed device or attempts to defeat security by starting and interrupting flash programming sequences. All programming, debug, and test interfaces are disabled when maximum device security is enabled.

1.2.1 CPU Subsystem

1.2.1.1 CPU

The TVII-B-H CPU subsystem contains a 32-bit Arm Cortex-M0+ CPU with MPU, two 32-bit Arm Cortex-M7 CPUs each with MPU, single/double-precision FPU, and 16-KB data and instruction caches. This subsystem also includes P-/M-DMA controllers, a cryptographic accelerator, up to 8384 KB

of code-flash, up to 256 KB of work-flash, up to 1024 KB of SRAM, and 64 KB of ROM.

The Cortex-M0+ CPU provides a secure, un-interruptible boot function. This guarantees that after the boot function is complete, system integrity is valid and privileges are enforced. Shared resources such as flash, SRAM, and peripherals can be accessed through bus arbitration, and exclusive accesses are supported by an inter-processor communication (IPC) mechanism using hardware semaphores.

Each Cortex-M7 CPU has 16 KB of instruction and 16 KB of data TCM with an option of programmable read wait states. Each TCM is clocked by the associated Cortex-M7 CPU clock.

1.2.1.2 DMA Controllers

TVII-B-H has two types of DMA controllers: P-DMA and M-DMA. P-DMA is used for peripheral-to-memory and memory-to-peripheral data transfers and provides low latency for a large number of channels. Each P-DMA controller uses a single data-transfer engine that is shared by the associated channels. General-purpose channels have a rich interconnect matrix including P-DMA cross triggering, which enables demanding data-transfer scenarios. Dedicated channels have a single triggering input (such as an ADC channel) to handle common transfer needs. M-DMA is used for memory-to-memory data

transfers and provides high memory bandwidth for a small number of channels. M-DMA uses a dedicated data-transfer engine for each channel. They support independent accesses to peripherals using the AHB multi-layer bus.

1.2.1.3 Flash

TVII-B-H has up to 8384 KB (8128 KB with a 32-KB sector size, and 256 KB with an 8-KB sector size) of code-flash with an additional work-flash of up to 256 KB (192 KB with a 2-KB sector size, and 64 KB with a 128-B sector size). Work-flash is optimized for reprogramming many more times than code-flash. Code-flash supports Read-While-Write (RWW) operation allowing flash to be updated while the CPU is active. Both the code-flash and work-flash areas support dual-bank operation for over-the-air (OTA) programming.

1.2.1.4 SRAM with 32-KB Retention Granularity

TVII-B-H has up to 1024 KB of SRAM with three independent controllers. The first controller SRAM0 provides deep-sleep retention in 32 KB increments while SRAM1/2 is selectable between fully retained and not retained.

1.2.1.5 ROM

TVII-B-H has 64 KB of ROM that contains boot and configuration routines. This ROM enables secure boot and authentication of user flash to guarantee a secure system.

1.2.1.6 Cryptography Accelerator for Security

The cryptography accelerator implements (3)DES block cipher, AES block cipher, SHA hash, cyclic redundancy check, pseudo random number generation, true random number generation, and a vector unit to support asymmetric key cryptography such as RSA and ECC.

Depending on the part number, this block is either completely or partially available or not available at all.

1.2.2 System Resources

1.2.2.1 Power System

The power system ensures that the supply voltage levels meet the requirements of each power mode, and provides a full-system reset when these levels are not valid. Internal power-on reset (POR) guarantees full-chip reset during the initial power ramp.

Three brownout detection (BOD) circuits monitor the external supply voltages (V_{DD} , V_{DDA} , V_{CCD}). The BOD on V_{DD} and V_{CCD} are initially enabled and cannot be disabled. The BOD on V_{DDA} is initially disabled and can be enabled

by the user. For the external supplies V_{DD} and V_{DDA} , BOD circuits are software configurable with two settings; a 2.7-V minimum voltage, which is robust for all internal signaling and a 3.0-V minimum voltage, which is also robust for all I/O specifications (that are guaranteed at 2.7 V). The BOD on V_{CCD} is provided as a safety measure and is not a robust detector.

Three over-voltage detection (OVD) circuits are provided for monitoring external supplies (V_{DD} , V_{DDA} , V_{CCD}), and over-current detection circuits (OCD) for monitoring internal and external regulators. OVD thresholds on V_{DD} and V_{DDA} are configurable with two settings; a 5.0-V and 5.5-V maximum voltage.

Two voltage detection circuits are provided to monitor the external supply voltage (V_{DD}) for falling (low-voltage detector, LVD) and rising levels (high-voltage detector, HVD), each configurable for one of the 26 selectable levels.

All BOD, OVD, and OCD circuits on V_{DD} and V_{CCD} generate a reset, because these protect the CPUs and fault logic. The BOD and OVD circuits in V_{DDA} can be configured to generate either a reset, or a fault.

1.2.2.2 Regulators

TVII-B-H contains three regulators that provide power to the low-voltage core transistors: DeepSleep, core internal, and core external. These regulators accept a 2.7 – 5.5-V V_{DD} supply and provide a low-noise 1.1-V supply to various parts of the device. These regulators are automatically enabled and disabled by hardware and firmware when switching between power modes. The core internal and core external regulators operate in Active mode, and provide power to the CPU subsystem and associated peripherals.

DeepSleep. The deep-sleep regulator is used to maintain power to a small number of blocks when in DeepSleep mode. These blocks include the ILO and WDT timers, BOD detector, SCB0, SRAM memories, Smart I/O, and other configuration memories. The deep-sleep regulator is enabled when in DeepSleep mode, and when the core internal regulator is disabled. It is disabled when XRES_L is asserted (LOW) and when the core internal regulator is disabled.

Core Internal. The core internal regulator supports load currents up to 300-mA, and is operational during device start up (boot process), and in Active/Sleep modes.

Core External. To support worst case loading, with both CM7 CPUs and the CM0+ CPU at their maximum clock frequency and all integrated peripherals operating, the core external regulator must handle load current up to 600 mA. While the control and monitor circuits for the core external regulator are internal to TVII-B-H, the power regulating element (NPN pass transistor, PMIC, or LDO) is external.

This reduces the overall power dissipation within the TVII-B-H package, while maintaining a well-regulated core supply. The core external regulator may be implemented with either an external NPN pass transistor, linear regulator (LDO), or PMIC. Each implementation requires different external components on the PCB, and different connections to TVII-B-H for both regulation and control.

1.2.2.3 Clock System

The TVII-B-H clock system provides clocks to all subsystems that require them, and glitch-free switching between different clock sources. In addition, the clock system ensures that no metastable conditions occur.

The clock system for TVII-B-H consists of the 8-MHz IMO, ILOs, watchdog timers, PLLs, an FLL, clock supervisors (CSV), an ECO, and a WCO.

The clock system supports three main clock domains: CLK_HF, CLK_SLOW, and CLK_LF.

- CLK_HF_x are the active domain clocks. Each can use any of the high frequency clock sources including IMO, EXT_CLK, ECO, FLL, or PLL.
- CLK_SLOW provides reference clock for the CM0+ CPU, Crypto, P-/M-DMA, and other slow infrastructure blocks of CPU subsystem.
- CLK_LF is a DeepSleep domain clock and provides source for MCWDT or RTC modules. The reference clock for the CLK_LF domain is selectable from ILO0, ILO1, WCO, or disabled.

1.2.2.4 IMO Clock Source

The IMO is the frequency reference in TVII-B-H when no external reference is available or enabled. The IMO operates at a frequency of 8 MHz $\pm$ 1%. The internal trim settings for the IMO can be dynamically updated to provide a tolerance of less than 1%.

1.2.2.5 ILO Clock Source

An ILO is a low-power oscillator, which generates clocks for a watchdog timer when in DeepSleep mode. There are two ILOs to ensure CSV capability in the DeepSleep mode. ILO-driven counters can be calibrated to the IMO, WCO, or ECO to improve their accuracy. ILO1 is also used for clock supervision.

1.2.2.6 PLL and FLL

A PLL or FLL may be used to generate high-speed clocks from the IMO, ECO, or EXT_CLK. The FLL provides a much faster lock than the PLL in exchange for a small amount of frequency error and a lower maximum output frequency. 400-MHz PLLs supports spread spectrum clock generation (SSCG) and down spreading (down spreading is where the maximum frequency of the spread spectrum clock is the same as that of the non-modulated clock).

1.2.2.7 Clock Supervisor

Each clock supervisor (CSV) allows one clock (reference) to supervise the behavior of another clock (monitored). Each CSV has counters for both monitored and reference clocks. Parameters for each counter determine the frequency of the reference clock as well as the upper and lower frequency limits of the monitored clock. If the frequency range comparator detects a stopped clock or a clock outside the specified frequency range, an abnormal state is signaled and either a reset or an interrupt is generated.

1.2.2.8 EXT_CLK

Dedicated GPIOs can be used to provide an external clock. This clock can be used as the source clock for either the PLL or FLL, or can be used directly by the CLK_HF domain.

1.2.2.9 ECO

The ECO provides high-frequency clocking using an external crystal connected to the ECO_IN and ECO_OUT pins. It supports fundamental mode (non-overtone) quartz crystals. When used in conjunction with the PLL, it generates CPU and peripheral clocks up to device's maximum frequency. ECO accuracy depends on the selected crystal. If the ECO is disabled, the associated pins can be used for any of the available I/O functions.

1.2.2.10 WCO

The WCO is a low-power, watch-crystal oscillator intended for real-time-clock applications. It requires an external crystal connected to the WCO_IN and WCO_OUT pins. The WCO can also be configured as a clock reference for CLK_LF, which is the clock source for the MCWDT and RTC.

1.2.2.11 Reset

TVII-B-H can be reset from a variety of sources, including software. Reset events are asynchronous and guarantee reversion to a known state. The reset cause (POR, BOD, OVD, overcurrent, XRES_L, WDT, MCWDT, software reset, fault, CSV, Hibernate wakeup, or debug) is recorded in a register, which is sticky through reset and allows software to determine the cause of the reset. An XRES_L pin is available for external reset.

1.2.2.12 Power Modes

TVII-B-H supports six power modes.

- Active - All peripherals are available
- Low-Power Active (LPACTIVE) - Low-power profile of Active mode where all peripherals and the CPUs are available, but with limited capability
- Sleep - All peripherals except the CPUs are available

- Low-Power Sleep (LPSLEEP) - Low-power profile of Sleep mode where all peripherals except the CPUs are available, but with limited capability
- DeepSleep - Only peripherals that work with CLK_LF are available
- Hibernate - The device and I/Os are in High-Z state and the device resets on wakeup

1.2.3 Peripherals

1.2.3.1 Peripheral Clock Dividers

Integer and fractional clock dividers are provided for peripheral and timing purposes.

1.2.3.2 Peripheral Protection Unit

The peripheral protection unit (PPU) controls and monitors unauthorized access from all masters (CPU, P-/M-DMA, Crypto, and the debug interface) to the peripherals. It allows or restricts data transfers on the bus infrastructure. The access rules are enforced based on specific properties of a transfer, such as an address range for the transfer and access attributes (such as read/write, user/privilege, and secure/non-secure).

1.2.3.3 12-bit SAR ADC

TVII-B-H contains three 1-Msps SAR ADCs, which can be clocked at up to 26.67 MHz and provide a 12-bit result in 26 clock cycles.

The references for all three SAR ADCs come from a dedicated pair of inputs: VREFH and VREFL.

Each ADC has a sequencer, which autonomously cycles through the configured channels (sequencer scan) with zero-switching overhead (that is, the aggregate sampling bandwidth, when clocked at 26.67 MHz, is equal to 1 Msps whether it is for a single channel or distributed over several channels). The sequencer switching is controlled through a state machine or firmware. The sequencer prioritizes trigger requests, enables the appropriate analog channel, controls ADC sampling, initiates ADC data conversion, manages results, and initiates subsequent conversions for repetitive or group conversions without CPU intervention.

Each SAR ADC has an analog multiplexer used to connect the signals to be measured to the ADC. It has 32 GPIO_STD inputs, one special GPIO_STD input for motor-sense, and six additional inputs to measure internal signals such as a band-gap reference, a temperature sensor, and power supplies. The device supports synchronous sampling of one motor-sense channel on each of the three ADCs.

TVII-B-H has one temperature sensor that is shared by all three ADCs. The temperature sensor must only be sampled by one ADC at a time. Software post processing is required to convert the temperature sensor reading into kelvin or Celsius values.

To accommodate signals with varying source impedances and frequencies, it is possible to have different sample times programmed for each channel. Each ADC also supports range comparison, which allows fast detection of out-of-range values without having to wait for a sequencer scan to complete and for the CPU firmware to evaluate the measurement for out-of-range values.

The ADCs are not usable in the DeepSleep and Hibernate modes as they require a high-speed clock. The ADC input reference voltage (VREFH) range is 2.7 V to V_{DDA} and VREFL is V_{SSA} .

1.2.3.4 Timer/Counter/PWM Block (TCPWM)

The TCPWM block consists of 16-bit and 32-bit counters with user-programmable period. Some of the 16-bit counters are optimized for motor-control operations. Each TCPWM counter contains a capture register to record the count value at the time of an event, a period register (used to either stop or auto-reload the counter when its count is equal to the period register), and compare registers to generate signals that are used as PWM duty-cycle outputs.

Each counter within the TCPWM block supports several functional modes such as timer, capture, quadrature, PWM with dead-time insertion (PWM_DT, 8-bit), pseudo-random PWM (PWM_PR), and shift-register.

In motor control applications, the counter within the TCPWM block supports Enhanced Quadrature mode with features such as asymmetric PWM generation, dead-time insertion (16-bit), and association of different dead times for PWM output signals.

The TCPWM block also provides true and complement outputs, with programmable offset between them, to allow their use as deadband complementary PWM outputs. The TCPWM block also has a kill input (only for the PWM mode) to force outputs to a predetermined state; for example, this may be used in motor-drive systems when an overcurrent state is detected and the PWMs driving the FETs need to be shut off immediately (no time for software intervention).

1.2.3.5 Serial Communication Blocks (SCB)

TVII-B-H contains up to 11 serial communication blocks, each configurable to support I²C, UART, or SPI.

I²C Interface. An SCB can be configured to implement a full I²C master (capable of multi-master arbitration) or slave interface. Each SCB configured for I²C can operate at speeds of up to 1 Mbps (Fast-mode Plus) and has flexible buffering options to reduce the interrupt overhead and latency of the CPU. In addition, each SCB supports FIFO buffering for receive and transmit data, which, by increasing the time for the CPU to read the data, reduces the need for clock stretching. The I²C interface is compatible with Standard, Fast-mode, and Fast-mode Plus devices as

specified in the NXP I²C-bus specification and user manual (UM10204). The I²C-bus I/O is implemented with GPIO in open-drain modes.

UART Interface. When configured as a UART, each SCB provides a full-featured UART with maximum signaling rate determined by the configured peripheral-clock frequency and over-sampling rate. It supports infrared interface (IrDA) and SmartCard (ISO 7816) protocols, which are minor variants of the UART protocol. It also supports the 9-bit multiprocessor mode that allows addressing of peripherals connected over common RX and TX lines. Common UART functions such as parity, number of stop bits, break detect, and frame error are supported. FIFO buffering of transmit and receive data allows greater CPU service latencies to be tolerated.

The LIN protocol is supported by the UART. LIN is based on a single-master multi-slave topology. The LIN bus has one master node and multiple slave nodes. The SCB UART supports only LIN slave functionality. Compared to the dedicated LIN blocks, an SCB/UART used for LIN requires a higher level of software interaction and increased CPU load.

SPI Interface. The SPI configuration supports full Motorola SPI, TI Synchronous Serial Protocol (SSP, essentially adds a start pulse that is used to synchronize SPI-based codecs), and National Microwire (a half-duplex form of SPI). The SPI interface can use the FIFO. SCB also supports EZSPI[7] mode.

SCB0 supports the following additional features:

- Operable as a slave in DeepSleep mode
- I²C slave EZ (EZI2C) mode with up to 256-byte data buffer for multi-byte communication without CPU intervention
- I²C slave externally-clocked operations
- Command/response mode with a 512-byte data buffer for multi-byte communication without CPU intervention

1.2.3.6 CAN FD

TVII-B-H contains CAN FD controller blocks, each supporting one or more CAN FD channel. All CAN FD controllers are compliant with the ISO 11898-1:2015 standard; an ISO 16845:2015 certificate is available. It also implements the time-triggered CAN (TTCAN) protocol specified in ISO 11898-4 (TTCAN protocol levels 1 and 2) completely in hardware. All functions concerning the handling of messages are implemented by the RX and TX handlers. The RX handler manages message acceptance filtering, transfer of received messages from the CAN core to a message RAM, and provides receive-message status. The TX handler is responsible for the transfer of transmit messages from the message RAM to the CAN core, and provides transmit-message status.

1.2.3.7 Local Interconnect Network (LIN)

TVII-B-H contains LIN channels that support transmission/reception of data following the LIN protocol according to ISO standard 17987. Each LIN channel connects to an external transceiver through a three-pin interface (including an enable function) and supports master and slave functionality. It also supports classic and enhanced checksum, along with break detection during message reception and wake-up signaling. Break detection, sync field, checksum calculations, and error interrupts are handled in hardware.

1.2.3.8 FlexRay Interface

TVII-B-H supports FlexRay interface with channel A and an optional channel B, conforming to FlexRay protocol specification 2.1, and supports up to a 10-Mbps data rate. Message buffers are configurable as TX, RX, or RXFIFO, and are filtered based on FrameID, cycle count, and message ID.

1.2.3.9 Ethernet MAC

TVII-B-H supports Ethernet channels with transfer rates of 10, 100, or 1000 Mbps. The input/output frames and flow control are compliant to the Ethernet/IEEE 802.3az standards, and also IEEE-1588 precision-time protocol (PTP). The device supports full-duplex data transport using external PHY devices. The MAC supports glue-free connection to PHYs through IEEE standard MII, RMII, GMII, and RGMII interfaces. The device also supports audio-video bridging (AVB).

1.2.3.10 External Memory Interface

In addition to the 8MB of internal flash memory, TVII-B-H supports direct connection to an external flash or RAM memory. This connection is made through either a HyperBus or SPI. HyperBus allows connection to HyperFlash and HyperRAM devices, while SPI (single, dual, quad, or octal SPI) can connect with serial flash memory. Code stored in memory connected through this interface allows execute-in-place (XIP) operation, which does not require the instructions to be first copied to internal memory, and on-the-fly encryption and decryption for environments requiring secure external data and code.

1.2.3.11 SDHC Interface

TVII-B-H supports secure digital host capacity (SDHC) interface, which conforms to SD 6.0, Secure digital input output (SDIO) 4.10, and embedded multimedia card (eMMC) 5.1 specifications, along with host control interface (HCI) 4.2 specification. The interface supports system DMA (SDMA), advance DMA (ADMA2, ADMA3), and command queuing (CQ) features. Supports data rates of SD DS (default speed, 4-bits at 25 MHz), SD HS (high speed, 4-bits at 50 MHz, and eMMC 52-MHz DDR (8-bits at 52-MHz card clock).

1.2.3.12 One-Time-Programmable (OTP) eFuse

TVII-B-H devices contain a 1024-bit OTP eFuse memory that can be used to store and access a unique and unalterable identifier or serial number for each device. eFuses are also used to control the device lifecycle (manufacturing, programming, normal operation, end-of-life, and so on) and the security state.

1.2.3.13 Event Generator

The event generator supports generation of interrupts and triggers in the Active mode and interrupts in the DeepSleep mode. The event generators are used to trigger a specific device function (execution of an interrupt handler, a SAR ADC conversion, and so on) and provide a cyclic wakeup mechanism from the DeepSleep mode. They provide CPU-free triggers for device functions, and reduce CPU involvement in triggering device functions, thus reducing overall power consumption and processing overhead.

1.2.3.14 Trigger Multiplexer

TVII-B-H supports connecting various peripherals using trigger signals. Triggers are used to inform a peripheral of the occurrence of an event or change of state. These triggers are used to effect or initiate some action in other peripherals. The trigger multiplexer is used to route triggers from a source peripheral to a destination. Triggers provide active logic functionality and are supported in the Active mode.

1.2.4 I/Os

TVII-B-H has three types of programmable I/Os: GPIO Standard, GPIO Enhanced, and HSIO Standard.

The I/Os are organized as logical entities called ports, which are a maximum of eight bits wide. During power-on, Hibernate, and reset, the I/Os are forced to the High-Z state.

Every I/O can generate an interrupt (if enabled) and each port has an interrupt request (IRQ) and interrupt service routine (ISR) associated with it.

1.2.4.1 GPIO

GPIO Standard (GPIO_STD). Supports standard automotive signaling across the 2.7-V to 5.5-V V_{DDIO} range. GPIO Standard I/Os have multiple configurable drive levels, drive modes, and selectable input levels.

GPIO Enhanced (GPIO_ENH). Supports extended functionality automotive signaling across the 2.7-V to 5.5-V V_{DDIO} range with higher currents at lower voltages (full I²C timing support and slew-rate control).

Both GPIO_STD and GPIO_ENH implement the following:

- Configurable input threshold (CMOS, TTL, or Automotive)
- Hold mode for latching previous state (used to retain the I/O state in DeepSleep mode)
- Analog input mode (input and output buffers disabled)

1.2.4.2 HSIO

HSIO Standard (HSIO_STD). These I/Os are optimized exclusively for high-speed signaling (supports automotive signaling across the 2.7-V to 3.6-V V_{DDIO} range) and do not support slew-rate control, DeepSleep operation, POR mode control, analog connections, or non-CMOS signaling levels. HSIO_STD supports high-speed peripherals such as QSPI, HyperBus, Ethernet, and SDHC controller. It also supports programmable drive strength. These I/Os are available only in Active mode and retain state in DeepSleep modes.

1.2.4.3 Drive Modes

All I/Os support the following programmable drive modes:

- High impedance
- Resistive pull-up
- Resistive pull-down
- Open drain with strong pull-down
- Open drain with strong pull-up
- Strong pull-up or pull-down
- Weak pull-up or pull-down

1.2.4.4 Port Nomenclature

Px.y describes a particular bit “y” available within an I/O port “x.”

For example, P4.2 reads “port 4, bit 2”.

1.2.4.5 Smart I/O

Smart I/O allows Boolean operations on signals going to the GPIO from the device subsystems or on signals coming into the chip. Operation can be synchronous or asynchronous and the blocks operate in all device power modes except for the Hibernate mode.

2. Getting Started



2.1 Support

Free support for Traveo II products is available online at www.cypress.com. Resources include training seminars, application notes, CRM technical support email, knowledge base, and application support engineers.

For application assistance, visit www.cypress.com/support/ or call 1-800-541-4736.

2.2 Product Upgrades

Cypress provides scheduled upgrades and version enhancements for the sample driver library (SDL) (which can be used only to evaluate the microcontroller and is not for production usage) and Cypress programmer (CYP) free of charge. Upgrades are available at www.cypress.com. Critical updates to system documentation are also provided in the Documentation section.

2.3 Development Kits

The CYTVII-B-H kits enable customers to evaluate and design with the TVII-B-H series of devices. The CYTVII-B-H evaluation kit is a combination of the following:

- CYTVII-B-E-BB: Traveo II base board that includes peripherals such as CANFD, LIN, and CXPI that evaluate some of the key features of the TVII-B-H series of devices.
- CYTVII-B-H-8M-XX-CPU: This CPU board supports Ethernet, SMIF (SPI/Hyper Memory), UART, I2S, and so on for device evaluation purposes.

Quick start guides for the respective kits can be downloaded from www.cypress.com. Refer to the respective device datasheet to understand the supported peripherals.

2.3.1 Evaluation Board

The CYTVII-B-H evaluation board comprises of two boards. The CPU board (CYTVII-B-H-8M-XX-CPU) should be plugged into the base board (CYTVII-B-E-BB). Together with the base board, the CPU board provides access to all peripherals that are available on the base board, resulting in a fully-featured evaluation platform. When plugged together, the CPU board and Traveo II base board is referred to as the evaluation kit. The CYTVII-B-H-8M-XX-CPU has the Cortex debug connector and Cortex "Debug + ETM" connector for the debug interface. The evaluation boards will be used to evaluate device performance and development of software.

Figure 2-1. Evaluation Board Hook Up with 320-BGA CPU Board




Table 2-1. CPU Board Part Number

Board	TVII-B-H Series Device
CYTVII-B-H-8M-320-CPU	320-BGA
CYTVII-B-H-8M-176-CPU	176-TEQFP
CYTVII-B-H-8M-272-CPU	272-BGA

2.3.2 Base Board

Table 2-2. Resource List of CYTVII-B-E-BB

Components	Description	Availability
CAN FD	Connector (D-sub 9-pin) and Transceiver (TJA1057GT, 6 channels and TJA1145T/FDJ, 4 channels)	10 channels
LIN	Connector (D-sub 9-pin) and Transceiver (TJA1021T/20)	6 channels
I ² C	Pin header	8 channels
CXPI	Connector (D-sub 9-pin) and Transceiver (S6BT112A)	1 channel
ADC	Potentiometer to control an ADC channel	1 channel
FlexRay	Connector (D-sub 9-pin) and Transceiver (AS8221)	2 channels
Potentiometer	To control ADC (Alps RK09K1130A8G)	1 number
User Button	User controlled	5 numbers
User LEDs	Green LED	10 numbers

2.3.3 Sample Driver Library (SDL)

Cypress provides a sample driver library (SDL) including startup as sample software. The SDL provides a simple interface to access various peripherals and can be used for evaluation, hardware bring-up, benchmarks, feasibility studies, and solution demos. It also serves as a reference to customers for drivers that are not covered by the official AUTOSAR products. The SDL cannot be used for production purposes because it does not qualify to any automotive standards. The SDL integrates

device header files, startup code, and peripheral drivers. The SDL contains a set of firmware drivers that provide APIs to access the device-specific resources.

2.3.4 Development Tools

Table 2-3. Supported Development Tools

Vendor	Emulators/Probes	Compiler
GHS	GHS probe (5.6.5)	MULTI V7 (version 7.1.4)
IAR	I-JET	EWARM (8.42)
iSYSTEM	iC5000	–
Lauterbach	TRACE 32-ICE	–

2.3.5 Cypress Programmer (CYP)

CYP is a freely available programmer targeted to program code flash, work flash, and supervisory flash supported by the Traveo II MCU family of devices. It uses either the Cypress-specific MiniProg4 or Segger J-Link to perform the required activities and is a command line based utility.

2.4 Application Notes

See application note *AN220118 – Getting Started with Traveo II* for additional information on Traveo II device capabilities such as:

- Hardware connection information
- SDL folder structure, driver support, and its usage with third-party IDEs
- Startup sequence related to the device and individual cores

3. Document Construction



This document includes the following sections:

- [Section B: CPU Subsystem on page 36](#)
- [Section C: System Resources Subsystem \(SRSS\) on page 166](#)
- [Section D: Input/Output Subsystem Overview on page 245](#)
- [Section E: Digital Subsystem on page 273](#)
- [Section F: Analog Subsystem on page 653](#)
- [Section G: Program and Debug Overview on page 683](#)

3.1 Major Sections

For ease of use, information is organized into sections and chapters that are divided according to device functionality.

- **Section** – Presents the top-level architecture, how to get started, and conventions and overview information of the product.
- **Chapter** – Presents the chapters specific to an individual aspect of the section topic. These are the detailed implementation and information for some aspect of the integrated circuit.
- **Glossary** – Defines the specialized terminology used in this technical reference manual (TRM).
- **Registers Technical Reference Manual** – Supplies all device register details summarized in the technical reference manual. This is an additional document.

3.2 Documentation Conventions

This document uses only four distinguishing font types, besides those found in the headings.

- The first is the use of *italics* when referencing a document title or file name.
- The third is the use of Times New Roman font, distinguishing equation examples.
- The fourth is the use of `Courier New` font, distinguishing code examples.

3.2.1 Register Conventions

Register conventions are detailed in the *Traveo II Body Controller High Registers TRM*.

3.2.2 Numeric Naming

Hexadecimal numbers are represented with all letters in uppercase with an appended lowercase 'h' (for example, 14h or 3Ah) and hexadecimal numbers may also be represented by a '0x' prefix, the C coding convention. Binary numbers have an appended lowercase 'b' (for example, 01010100b or 01000011b). Numbers not indicated by an 'h' or 'b' are decimal.

3.2.3 Units of Measure

This table lists the units of measure used in this document.

Table 3-1. Units of Measure

Abbreviation	Unit of Measure
bps	bits per second
°C	degrees Celsius
dB	decibels
dBm	decibels-milliwatts
fF	femtofarads
G	Giga
Hz	Hertz
k	kilo, 1000
K	kilo, 2 ¹⁰
KB	1024 bytes, or approximately one thousand bytes
Kbit	1024 bits
kHz	kilohertz (1000)
kΩ	kilohms
MHz	megahertz
MΩ	megaohms
μA	microamperes
μF	microfarads
μs	microseconds
μV	microvolts
μVrms	microvolts root-mean-square
mA	milliamperes
ms	milliseconds
mV	millivolts
nA	nanoamperes
ns	nanoseconds
nV	nanovolts
Ω	ohms
pF	picofarads
pp	peak-to-peak
ppm	parts per million
SPS	samples per second
s	sigma: one standard deviation
V	volts

3.2.4 Acronyms and Abbreviations

This table lists the acronyms used in this document.

Table 3-2. Acronyms and Abbreviations

Acronym	Description
A/D	analog to digital
ABS	absolute
ADC	analog-to-digital converter
AES	Advanced Encryption Standard

Table 3-2. Acronyms and Abbreviations

Acronym	Description
AHB	AMBA (advanced microcontroller bus architecture) high-performance bus, Arm data transfer bus
Arm	Advanced RISC Machine, a CPU architecture
ASIL	automotive safety integrity level
AUTOSAR	Automotive Open System Architecture
BD	buffer descriptor
BLE	Bluetooth Low Energy
BOD	brown-out detection
CAN FD	controller area network with flexible data rate
CBS	credit-based shaping
CFI	canonical format indicator
CMOS	complementary metal-oxide-semiconductor
CPHA	(SPI) clock phase
CPOL	(SPI) clock polarity
CPU	central processing unit
CPUSS	CPU subsystem
CRC	cyclic redundancy check, an error-checking protocol
CSV	clock supervisor
DDR	double data rate (also see SDR)
DES	data encryption standard
DMA	direct memory access
DW	data wire, same as P-DMA
ECC	error correcting code
ECO	external crystal oscillator
EEE	Energy Efficient Ethernet (IEEE Std 802.3az)
EOF	end of frame
ETM	embedded trace macrocell
FCS	frame check sequence
FIFO	first in first out
FLL	frequency locked loop
FPU	floating point unit
GHS	Green Hills tool chain with IDE
GPIO	general-purpose input/output
HSIOM	high-speed I/O matrix
HSM	hardware security module
IF	interface
I/O	input/output
IP	Internet protocol
IPG	inter-packet gap
I ² C	Inter-Integrated Circuit, a communications protocol
ILO	internal low-speed oscillator
IMO	internal main oscillator
IPC	inter-processor communication
IrDA	infrared interface
IRQ	interrupt request
JTAG	Joint Test Action Group
LAN	local area network (IEEE Std 802)

Table 3-2. Acronyms and Abbreviations

Acronym	Description
LIN	local interconnect network, a communications protocol
LLDP	link layer discovery protocol (IEEE Std 802.1AB)
LPI	low-power idle (IEEE Std 802.3az)
LVD	low-voltage detection
MAC	media access control (IEEE Std 802)
MCU	microcontroller Unit
MCWDT	multi-counter watchdog timer
M-DMA	memory-direct memory access
MDC	management data clock
MII	media independent interface
MISO	master-in slave-out
MMIO	memory mapped I/O
MOSI	master-out slave-in
MPU	memory protection unit
NSP	non-standard preamble
NVIC	nested vectored interrupt controller
OTA	over-the-air programming
OTP	one-time programmable
OVD	over voltage detection
P-DMA	peripheral-direct memory access same as DW
PCS	physical coding sublayer
PFC	priority-based flow control (IEEE Std 802.1Qbb)
PLL	phase-locked loop
PHY	physical sublayer
POR	power-on reset
PPB	private peripheral bus
PPPoE	point-to-point protocol over Ethernet
PPU	peripheral protection unit
PRNG	pseudo-random number generator
PTP	precision time protocol (IEEE Std 1588)
PWM	pulse-width modulation
RAM	random access memory
RISC	reduced-instruction-set computing
ROM	read-only memory
RTC	real-time clock
RX	reception
SAR	successive approximation register
SCB	serial communication block
SCL	I ² C serial clock
SDA	I ² C serial data
SDR	single data rate (also see DDR)
SECDDED	single error correction double error detection
SerDes	serializer/deserializer
SHA	secure hash algorithm
SHE	secure hardware extension
SFD	start of frame delimiter

Table 3-2. Acronyms and Abbreviations

Acronym	Description
SGMII	serial Gigabit media independent interface
SMIF	serial memory interface
SMPU	shared memory protection unit
SNAP	subnetwork access protocol
SOF	start of frame
SPI	serial peripheral interface, a communications protocol
SRAM	static random-access memory
SWD	single wire debug
SYNC	LIN synchronization field
Tbit	bit period
TCP	transfer control protocol
TCPWM	timer/counter pulse-width modulator
TTL	transistor-transistor logic
TRNG	true random number generator
TS	timestamp
TSU	timestamp unit
TX	transmission
UART	universal asynchronous transmitter receiver, a communications protocol
UDP	user datagram protocol
VLAN	virtual LAN (IEEE Std 802.1Q)
WCO	watch crystal oscillator
WDT	watchdog timer reset
XIP	execute-in-place
XRES_L	external reset I/O pin (Active Low)
XTAL	crystal

Section B: CPU Subsystem




This section encompasses the following chapters:

- CPU Subsystem (CPUSS) chapter on page 37
- Inter-Processor Communication chapter on page 48
- Protection Unit chapter on page 53
- Direct Memory Access chapter on page 71
- Code Flash chapter on page 99
- Work Flash chapter on page 116
- SRAM Interface chapter on page 126
- BootROM chapter on page 133
- Interrupts chapter on page 144
- Device Security chapter on page 158
- Chip Operational Modes chapter on page 160
- Fault Subsystem chapter on page 162

Top Level Architecture

CPU System Block Diagram



4. CPU Subsystem (CPUSS)



The CPU subsystem is based on dual 32-bit Arm® Cortex® CPUs. The two Cortex-M7 are the main CPUs. The Cortex-M7 is designed for short interrupt response time, high code density, and high 32-bit throughput while maintaining a strict cost and power consumption budget. An additional Cortex-M0+ CPU can implement security, safety, and protection features.

This section provides only an overview of the Arm Cortex CPUs in Traveo II. For details, see the Arm documentation sets for [Cortex-M7](#) and [Cortex-M0+](#).

4.1 Features

The Traveo II Arm Cortex CPUs have the following features:

- Little-endian byte ordering for data, memory, and CPU registers
- Each Cortex-M7 has 16-KB instruction cache and 16-KB data cache with ECC support.
- Flash Controller has 8-KB cache for Cortex-M0+, described in [Code Flash chapter on page 99](#).
- Besides the shared system SRAM each Cortex-M7 has 16-KB of ITCM and 16-KB of DTCM (instruction/data tightly coupled memory) with ECC support.
See the device datasheet for address allocation of ITCM and DTCM.
- Cortex-M7 has a [floating-point unit](#) (FPU) with single and double precision that supports single-cycle digital signal processing (DSP) instructions, and a [memory protection unit](#) (MPU). Cortex-M0+ has an MPU.
- The Cortex-M7 supports a subset of the Thumb instruction set (defined in the [Arm®v7-M Architecture Reference Manual](#)). The Cortex-M0+ supports the Armv6-M Thumb instruction set (defined in the [Arm®v6-M Architecture Reference Manual](#)). See [4.5 Instruction Set](#).
- Both CPUs have [nested vectored interrupt controllers](#) (NVIC) for rapid and deterministic interrupt response.
- Both CPUs have extensive debug support. For details, see the [Program and Debug Interface chapter on page 684](#).
 - SWJ: combined serial wire debug (SWD) and Joint Test Action Group (JTAG) ports
 - Breakpoints
 - Watchpoints
 - Trace: Cortex-M7: instrumentation trace macrocell (ITM) and embedded trace macrocell (ETM) with embedded trace buffer (ETB) and trace port interface unit (TPIU). Cortex-M0+: micro trace buffer (MTB)
- Inter-processor communication (IPC) hardware - see the [Inter-Processor Communication chapter on page 48](#).

4.2 How It Works

Both Cortex CPUs are 32-bit processors with a 32-bit data path, 32-bit registers, and a 32-bit memory interface. They support a wide variety of instructions in the Thumb instruction set. The CPUs support two operating modes (see [4.4 Operating Modes and Privilege Levels](#)).

The Cortex-M7 instruction set includes:

- Signed and unsigned, 32×32 → 32-bit and 32×32 → 64-bit, multiply and multiply-accumulate, all single-cycle
- Signed and unsigned 32-bit divides that take two to 12 cycles
- DSP instructions
- Complex memory-load and store access
- Complex bit manipulation

The Cortex-M7 FPU has its own set of registers and instructions. It is compliant with the ANSI/IEEE Std 754-2008, IEEE Standard for Binary Floating-Point Arithmetic.

The Cortex-M0+ has a single cycle 32x32 -> 32-bit signed multiplication instruction.

4.3 Address Map

Both CPUs have a fixed address map, with shared access to memory and peripherals except the PPB area, which is a private address space for each core. The 32-bit (4 GB) address space is divided into the regions shown in [Table 4-1](#). Note that code can be executed from the code and SRAM regions.

Table 4-1. Address Map for Cortex-M7 and Cortex-M0+

Address Range	Name	Use
0x00000000 - 0x1FFFFFFF	Code	Program code region. You can also put data here. It includes the exception vector table, which starts at address 0.
0x20000000 - 0x3FFFFFFF	SRAM	Data region. You can also execute code from this region.
0x40000000 - 0x5FFFFFFF	Peripheral	All peripheral registers. You cannot execute code from this region.
0x60000000 - 0x9FFFFFFF	External RAM	SMIF. You can execute code from this region.
0xA0000000 - 0xDFFFFFFF	External device	External device memory
0xE0000000 - 0xE00FFFFF	PPB	Peripheral registers within the CPU core.
0xE0100000 - 0xFFFFFFFF	Device	Device-specific system registers.

Notes:

- Gaps in the address space are reserved. Do not access these gaps; if accessed, it can result in hard faults or BUS ERROR depending on which bus segment or peripheral an address is allocated to.

4.4 Operating Modes and Privilege Levels

Both CPUs support two operating modes and two privilege levels:

■ Operating Modes:

- Thread Mode – used to execute application software. The processor enters Thread mode when it comes out of reset.
- Handler Mode – used to handle exceptions. The processor returns to Thread mode when it has finished all exception processing.

■ Privilege Levels:

- Unprivileged – the software has limited access to the MSR and MRS instructions, and cannot use the CPS instruction. It cannot access the system timer, NVIC, or system control block. It may have restricted access to memory or peripherals.
- Privileged – the software can use all the instructions and has access to all resources.

In Thread mode, the CONTROL register controls whether software execution is privileged or unprivileged. In Handler mode, software execution is always privileged.

Only privileged software can write to the CONTROL register to change the privilege level. Unprivileged software can use the SVC instruction to transfer control to privileged software.

In Handler mode, the MSP is always used. The exception entry and return mechanisms automatically update the CONTROL register, which may change whether MSP/PSP is used.

In Thread mode, use the MSR instruction to set the stack pointer bit in the CONTROL register. When changing the stack pointer, use an ISB instruction immediately after the MSR instruction. This ensures that instructions after the ISB execute using the new stack pointer.

4.5 Instruction Set

Cortex-M0+ is based on the Armv6-M architecture and supports all the 16-bit Thumb instructions defined by the Armv7-M architecture excluding CBZ, CBNZ, and IT. In addition, it supports the following 32-bit Thumb instructions: BL, DMB, DSB, ISB, MRS, and MSR.

Cortex-M7 is based on the Armv7E-M architecture (Armv7-M with DSP extension) and supports all Thumb instructions defined by that architecture, including the optional floating point instructions.

For details, see one of the Arm Cortex Generic User Guides, Technical Reference Manuals or Architecture Reference Manuals.

4.6 TCM Interface

Cortex-M7 has three TCM interfaces:

- ITCM, 64-bit data
- D0TCM, 32-bit data
- D1TCM, 32-bit data

The TCM interface implements the following functionality:

- Programmable read wait state (0 or 1).
- ECC functionally
 - Single-bit error automatic correction
 - Single- and multi-bit error detection and fault reporting
 - ECC error injection

4.6.1 TCM Read Wait State

The ITCM interface has a programmable read wait states. The ITCM and DTCM wait states are independently programmable through the CPUSS_CM7_x_CTL.ITCM_READ_WS and CPUSS_CM7_x_CTL.DTCM_READ_WS fields. When the CPUSS_CM7_x_CTL.ITCM_READ_WS or DTCM_READ_WS is '1', read access of TCM that set to '1' will get two cycles.

Note that TCM writes are always zero wait states.

Table 4-2. TCM Read Wait State Configuration

Register	Bit Field and Bit Name	Description
CPUSS_CM7_x_CTL[31:0]	ITCM_READ_WS	ITCM read wait states (writes have no wait states). 0: 0 wait state. 1: 1 wait state
	DTCM_READ_WS	DTCM read wait states (writes have no wait states). 0: 0 wait state 1: 1 wait state.

Note: The 'x' in the register name denotes the Cortex-M7 core number.

4.6.2 TCM ECC

Both ITCM and DTCMs have ECC with SECDED scheme (Single Error Correction, Dual Error Detection) for functional safety. ITCM and DTCMs interfaces have different data widths the ECC is slightly different:

- ITCM 64-bit data/8-bit ECC
- DTCMs 32-bit data/7-bit ECC

4.6.2.1 TCM Read Modify Write Configuration

To calculate the correct ECC parity, all data bits in the stored word must be included. Therefore, when partial write is done (for example, a byte write to the DTCM), it works as follows.

1. The stored data word is read from memory.
2. Read data word is merged with the partial data as new data word.
3. ECC parity of the new data word is calculated.
4. The new data word and calculated ECC parity are stored to memory.

The Cortex-M7 core already does all this provided the CPU-specific registers CM7_ITCMCR.RMW and CM7_DTCMCR.RMW bits are set. When ECC is used, they must be set to '1' for correct operation.

CM7_ITCMCR.RMW and CM7_DTCMCR.RMW bit initialization out of reset are controlled by the CPUSS_CM7_x_CTL.INIT_RMW_EN[1:0] register.

Note that this Cortex-M7 feature also works for partial write from the AHB slave interface (AHBS) to the TCMs.

Table 4-3. TCM Read-Modify-Write Configuration

Register	Bit Field and Bit Name	Description
CPUSS_CM7_x_CTL[31:0]	INIT_RMW_EN[1:0]	TCM read-modify-write enable initialization after reset: Bit 0: ITCM. '0': disabled; '1': enabled. Bit 1: DTCM. '0': disabled; '1': enabled. Note: When TCM ECC is enabled, the read-modify-write functionality should be enabled. This prevents partial (sub-word) writes to the TCM.

Note: The 'x' in the register name denotes the Cortex-M7 core number.

If they are disabled at reset then the following code example can be used to enable them in software:

```

CM7_ITCMCR EQU 0xE000EF90
CM7_DTCMCR EQU 0xE000EF94
LDR r11, =CM7_ITCMCR
LDR r0, [r11]
ORR r0, r0, #0x1:SHL:1 ; Set CM7_ITCMCR.RMW field
ORR r0, r0, #0x1:SHL:2 ; Set CM7_ITCMCR.RETEN field
STR r0, [r11]
LDR r11, =CM7_DTCMCR
LDR r0, [r11]
ORR r0, r0, #0x1:SHL:1 ; Set CM7_DTCMCR.RMW field
ORR r0, r0, #0x1:SHL:2 ; Set CM7_DTCMCR.RETEN field
STR r0, [r11]
DSB
ISB
  
```

See the Arm documentation sets for Cortex-M7 for more details.

4.6.2.2 ECC Failure Detection

If the ECC logic detects a single-bit error, calculate the corrected data and then write the corrected data back to the SRAM. Therefore, subsequent read (retry) for the same data will get the corrected data. However, if the ECC logic detects that the data has more than one-bit error or hard single-bit errors, then data cannot be corrected.

Thus, Cortex-M7 may retry the same read a few cycles later and cause an ECC failure detection again. Note that the retry read is not guaranteed because the Cortex-M7 core sometimes issues speculative reads.

In this case, if the retry read does happen, then this creates a live-lock situation; that is, the Cortex-M7 repeats the same read over and over and the TCM logic keeps responding with Retry. Because of the Retry, it is guaranteed that the firmware will not use the erroneous data.

This live-lock will need to be resolved by a high-level interrupt, for example, the Cortex-M7 watchdog timer (WDT) interrupt or an interrupt from the fault structure. Note that this may include the Cortex-M0+ processor resetting and rebooting the Cortex-M7.

4.6.2.3 ECC Fault Reporting

Both the correctable and the non-correctable ECC errors are reported to the central fault structure. For both correctable and non-correctable errors, a separate set of registers and a separate fault report channel is used.

This allows a more important/urgent non-correctable error to be reported immediately even if there was a preceding correctable error.

Note however that the three TCM interfaces share the registers and fault report channels for correctable and the non-correctable errors. This means that when two or more TCM interfaces detect an error at (almost) the same time, only one error will be reported and the other is lost

4.6.2.4 TCM Initialization

The TCM memories need to be initialized before reading when ECC is enabled to avoid unwanted ECC faults.

ITCM writes for initialization may result in 64-bit read-modify-writes from CM7. The reads in this process may also cause ECC faults. To avoid this, use CPUSS_CM7_x_CTL.ITCM_ECC_CHECK_DIS bit. This bit need to be set (only) during initialization of ITCM to avoid ECC checking during initialization.

Table 4-4. TCM Initialization

Register	Bit Field and Bit Name	Description
CPUSS_CM7_x_CTL[31:0]	ITCM_ECC_CHECK_DIS	Disable ECC checking and thus fault reports. This also disables ECC correction (required to enable initialization). Intended use is for initialization. This bit is ignored when TCM_ECC_EN = 0.

Note: The 'x' in the register name denotes the Cortex-M7 core number.

4.6.2.5 ECC Error Injection

The TCM interface ECC logic supports ECC error injection through the following registers:

- CPUSS_CM7_x_CTL.ITCM_ECC_INJ_EN
- CPUSS_CM7_x_CTL.DTCM_ECC_INJ_EN
- CPUSS_ECC_CTL.WORD_ADDR
- CPUSS_ECC_CTL.PARITY

If CPUSS_CM7_x_CTL.ITCM_ECC_INJ_EN or CPUSS_CM7_x_CTL.DTCM_ECC_INJ_EN set to '1', a write transfer to word address specified in the CPUSS_ECC_CTL.WORD_ADDR field is performed, ECC parity generation uses the parity specified in the CPUSS_ECC_CTL.PARITY field rather than the calculated parity. Reading the same location will cause ECC fault (**Note:** Set the RETRY bit to '0' in CM7 DTCMCR/ITCMCR before the read to avoid live-lock during ECC injection test.

The RETRY can be enabled back to '1' after the ECC injection test is completed.).

Table 4-5. TCM ECC Error Injection Register

Register	Bit Field and Bit Name	Description
CPUSS_CM7_x_CTL[31:0]	ITCM_ECC_INJ_EN	ITCM ECC error injection enable: 0: Disabled. 1: Enabled.
	DTCM_ECC_INJ_EN	DTCM ECC error injection enable: 0: Disabled. 1: Enabled.
CPUSS_ECC_CTL[31:0]	WORD_ADDR[23:0]	Specifies the word address where an error will be injected. - On a write transfer to this SRAM address and when the corresponding RAM0/RAM1/RAM2/ITCM/DTCM ECC_INJ_EN bit is '1', the parity (PARITY) is injected.
	PARITY[31:24]	ECC parity to use for ECC error injection at address WORD_ADDR. For the DTCM that has only seven parity bits, PARITY[30:24] is used as ECC parity.

Note: The 'x' in the register name denotes the Cortex-M7 core number.

4.6.2.6 ECC Parity Generation by Software

To inject the ECC error for fault generation, ECC parity must be generated by software. Follow this procedure to generate an 8-bit ECC parity for ITCM. Parity generation calculates an 8-bit Parity[7:0] over a 64-bit data word W[63:0]. First, a 128-bit ECC code word CS_SW[127:0] is created:

```
AW[119:0] = 120{1'b0};
AW[WORD_WIDTH-1:0] = W[WORD_WIDTH-1:0];
```

```
CW_SW[127:0] = {{8{1'b0}}, AW[119:0]};
```

Then, the 8-bit parity is calculated as the reduction XOR of the 128-bit code word CW_SW[127:0] ANDed with the following parity bit specific constants:

```
ECC_P0_SW = 128b00000001_10111111_10111011_01110101_10111110_00111010_01110010_11011100_
01000100_10000100_01001010_10001000_10010101_00101010_10101101_01011011;
ECC_P1_SW = 128b00000010_11011111_01110110_11111001_11011101_10011001_10111001_01110001_
00010001_00001000_10010011_00010001_00100110_10110011_00110110_01101101;
ECC_P2_SW = 128b00000100_11101111_11001111_10011111_10011010_11010101_11001110_10010111_
00000110_00010001_00011100_00100010_00111000_11000011_11000111_10001110;
ECC_P3_SW = 128b00001000_11110111_11101100_11110110_11101101_01100111_01001110_01101100_
10011000_00100001_11100000_01000011_11000000_11111100_00000111_11110000;
ECC_P4_SW = 128b00010000_11111011_01111011_10101111_01101011_10100110_10110101_10100110_
11100000_00111110_00000000_01111100_00000000_11111111_11111000_00000000;
ECC_P5_SW = 128b00100000_11111101_10110111_11001110_11110011_01101100_10101011_01011011_
11111111_11000000_00000000_01111111_11111111_00000000_00000000_00000000;
ECC_P6_SW = 128b01000000_11111110_11011101_01111011_01110100_11011011_01010101_10101011_
11111111_11111111_11111111_10000000_00000000_00000000_00000000_00000000;
ECC_P7_SW = 128b10000000_01111111_00000000_00000000_00000111_11111111_11111111_11111111_
11010100_01000010_00100101_10000100_01001011_10100110_01011100_10110111;
```

The parity bits are calculated as follows:

```
parity[0] = ^ (CW_SW[127:0] & ECC_P0_SW)
parity[1] = ^ (CW_SW[127:0] & ECC_P1_SW)
...
parity[7] = ^ (CW_SW[127:0] & ECC_P7_SW)
```

Follow this procedure to generate a 7-bit ECC parity for DTCM. Parity generation calculates a 7-bit Parity[6:0] over a 32-bit data word W[31:0]. First, a 64-bit ECC code word CW_SW[63:0] is created:

```
CW_SW[63:0] = 64{1'b0};
CW_SW [31:0] = W[31:0];
```


Then, the 7-bit parity is calculated as the reduction XOR of the 64-bit code word CW_SW[63:0] ANDed with the following parity bit specific constants:

```
ECC_P0_SW = 64b00000011_01111111_00110110_11011011_00100010_01010100_00101010_10101011;
ECC_P1_SW = 64b00000101_10111101_11101011_01011010_01000100_10011001_01001101_00110101;
ECC_P2_SW = 64b00001001_11011101_11011100_11101110_00001000_11100010_01110001_11000110;
ECC_P3_SW = 64b00010001_11101110_10111011_10101001_10001111_00000011_10000001_11111000;
ECC_P4_SW = 64b00100001_11110110_11010111_01110101_11110000_00000011_11111110_00000000;
ECC_P5_SW = 64b01000001_11111011_01101101_10110100_11111111_11111100_00000000_00000000;
ECC_P6_SW = 64b10000001_00000011_11111111_11111000_00010001_00101100_10010110_01011111;
```

The parity bits are calculated as follows:

```
parity[0] = ^ (CW_SW[63:0] & ECC_P0_SW)
parity[1] = ^ (CW_SW[63:0] & ECC_P1_SW)
...
parity[6] = ^ (CW_SW[63:0] & ECC_P6_SW)
```

4.6.2.7 TCM Enabling

The TCM interfaces can be enabled at reset in the system by CPUSS_CM7_x_CTL.INIT_TCM_EN.

Table 4-6. TCM Enabling

Register	Bit Field and Bit Name	Description
CPUSS_CM7_x_CTL[31:0]	INIT_TCM_EN[1:0]	TCM enable initialization after reset: Bit 0: ITCM. '0': disabled; '1': enabled. Bit 1: DTCM. '0': disabled; '1': enabled.

Note: The 'x' in the register name denotes the Cortex-M7 core number.

If they are disabled at reset, then the following code example can be used to enable both the instruction and data TCM interfaces in software:

```
CM7_ITCMCR EQU 0xE000EF90
CM7_DTCMCR EQU 0xE000EF94
LDR r11, =CM7_ITCMCR
LDR r0, [r11]
ORR r0, r0, #0x1 ; Set CM7_ITCMCR.EN field
STR r0, [r11]
LDR r11, =CM7_DTCMCR
LDR r0, [r11]
ORR r0, r0, #0x1 ; Set CM7_DTCMCR.EN field
STR r0, [r11]
DSB
ISB
```

See the Arm documentation sets for Cortex-M7 for more details.

Other masters can access ITCM/DTCM through specific address. Cortex-M7_0 can access ITCM/DTCM of Cortex-M7_1 through this address and Cortex-M7_1 can access ITCM/DTCM of Cortex-M7_0. Note that Cortex-M7 can not access its own TCM memories through this address. Such an access will result in an address decode failure and will be returned as a bus error. See the Data sheet for ITCM/DTCM address mapping.

Accessing the ITCM/DTCM by other masters is possible only when the power mode of the corresponding CM7 is ENABLED and TCMC_EN is "1".

Before changing the CM7 power mode from ENABLED to another power mode, the following steps need to be performed (not doing so and trying to access the TCM may cause hang-up):

1. Disable access to the CM7 TCM by setting the field CM7_0/1_CTL.TCMC_EN to '0'.
2. Confirm that there are no outstanding accesses to the CM7 TCM by checking that the fields CM7_0/1_STATUS.TCMC_* are '0'. Repeat step 2 if necessary.

3. Now it is safe to change the CM7 power mode.

Table 4-7. TCM Access Control and CM7 Power Mode Control

Register	Bit Field and Bit Name	Description
CM7_x_PWR_CTL[31:0]	PWR_MODE[1:0]	CM7_x Power mode: 0: Switch CM7_x off. 1: Reset CM7_x. 2: Put CM7_x in Retained mode. 3: Switch CM7_x on.
CPUSS_CM7_x_CTL[31:0]	TCMC_EN	CM7 TCMC access control: 0: Disable access to the CM7 I/D-TCM slave port (AHBS). Access attempts will get a bus error response. 1: Enable access to the CM7 I/D-TCM slave port (AHBS).

Note: The 'x' in the register name denotes the Cortex-M7 core number.

4.7 Cache ECC Fault Reporting

Correctable and non-correctable faults are reported for cache memories.

A single set of correctable and non-correctable faults is created for I-Cache and D-Cache ECC faults. If ECC error occurs on both caches at the same time, I-Cache ECC error has the priority for fault reporting.

4.8 Registers

4.8.1 Arm Specification Registers

Both CPUs have sixteen 32-bit registers, as [Table 4-8](#) shows:

- R0 to R12 - General-purpose registers. R0 to R7 can be accessed by all instructions; the other registers can be accessed by a subset of the instructions.
- R13 - Stack pointer (SP). There are two stack pointers, with only one available at a time. In thread mode, the CONTROL register indicates the stack pointer to use, Main Stack Pointer (MSP) or Process Stack Pointer (PSP). In applications with an operating system, it is recommended that the kernel should use the MSP and the threads should use the PSP.
- R14 - Link register. Stores the return program counter during function calls.
- R15 - Program counter. This register can be written to control program flow.

Table 4-8. Cortex-M7 and Cortex-M0+ Registers

Name	Type ^a	Reset Value	Description
R0 - R12	RW	Undefined	R0–R12 are 32-bit general-purpose registers for data operations.
MSP (R13) PSP (R13)	RW	[0x00000000] ^b	The SP is register R13. In thread mode, bit[1] of the CONTROL register indicates the stack pointer to use: 0 = MSP. This is the reset value. 1 = PSP. On reset, the processor loads the MSP with the value from address 0x00000000.
LR (R14)	RW	See note ^c	The link register (LR) is register R14. It stores the return information for subroutines, function calls, and exceptions.
PC (R15)	RW	[0x00000004] ^b	The program counter (PC) is register R15. It contains the current program address. On reset, the processor loads the PC with the value from address 0x00000004. Bit[0] of the value is loaded into the EPSR T-bit (see Table 4-9) at reset; it must always be 1.

Table 4-8. Cortex-M7 and Cortex-M0+ Registers (continued)

Name	Type ^a	Reset Value	Description
PSR	RW	Undefined	The program status register (PSR) combines: Application Program Status Register (APSR). Execution Program Status Register (EPSR). Interrupt Program Status Register (IPSR).
APSR	RW	Undefined	The APSR contains the current state of the condition flags from previous instruction executions.
EPSR	RO	0x01000000	On reset, the EPSR Thumb state bit is loaded with the value bit[0] of the register [0x00000004]. It must always be 1. In Cortex-M7, other bits in this register control the state of interrupt-continuable instructions and the if-then (IT) instruction.
IPSR	RO	0	The IPSR contains the current exception number.
PRIMASK	RW	0	The PRIMASK register prevents activation of all exceptions with configurable priority.
CONTROL	RW	0	The CONTROL register controls: - The privilege level in Thread mode; see 4.4 Operating Modes and Privilege Levels . - The currently active stack pointer, MSP or PSP. - Cortex-M7 only: Whether to preserve the floating-point state when processing an exception.
FAULTMASK	RW	0	Cortex-M7 only. Bit 0 = 1 prevents the activation of all exceptions except NMI.
BASEPRI	RW	0	Cortex-M7 only. When set to a nonzero value, prevents processing any exception with a priority greater than or equal to the value.

a. Describes access type during program execution in thread mode and handler mode. Debug access can differ.

b. [address] denotes the value stored at address

c. LR reset value is 0xFFFFFFFF in Cortex-M7, undefined in Cortex-M0+.

Use the MSR and MRS instructions to access the PSR, PRIMASK, CONTROL, FAULTMASK, and BASEPRI registers. [Table 4-9](#) and [Table 4-10](#) show how the PSR bits are assigned.

Table 4-9. Cortex-M7 PSR Bit Assignments

Bit	PSR Register	Name	Usage
31	APSR	N	Negative flag
30	APSR	Z	Zero flag
29	APSR	C	Carry or borrow flag
28	APSR	V	Overflow flag
27	APSR	Q	DSP overflow and saturation flag
26 – 25	EPSR	ICI/IT	Control interrupt-continuable and IT instructions
24	EPSR	T	Thumb state bit. Must always be 1. Attempting to execute instructions when the T bit is 0 results in a HardFault exception
23 – 20	–	–	Reserved

Table 4-9. Cortex-M7 PSR Bit Assignments (continued)

Bit	PSR Register	Name	Usage
19 – 16	APSR	GE	Greater than or equal flags for the SEL instruction
15 – 10	EPSR	ICI/IT	Control interrupt-continuable and IT instructions
9	–	–	Reserved
8 – 0	IPSR	ISR_NUMBER	Exception number of current ISR: 0 = Thread mode 1 = Reserved 2 = NMI 3 = HardFault 4 = MemManage 5 = BusFault 6 = UsageFault 7 - 10 = Reserved 11 = SVCcall 12 = Reserved for debug 13 = Reserved 14 = PendSV 15 = SysTick 16 = IRQ0 ... 255 = IRQ240

Table 4-10. Cortex-M0+ PSR Bit Assignments

Bit	PSR Register	Name	Usage
31	APSR	N	Negative flag
30	APSR	Z	Zero flag
29	APSR	C	Carry or borrow flag
28	APSR	V	Overflow flag
27 – 25	–	–	Reserved
24	EPSR	T	Thumb state bit. Must always be 1. Attempting to execute instructions when the T bit is 0 results in a HardFault exception
23 – 6	–	–	Reserved
5 - 0	IPSR	Exception Number	Exception number of current ISR: 0 = Thread mode 1 = Reserved 2 = NMI 3 = HardFault 4 - 10 = Reserved 11 = SVCcall 12, 13 = Reserved 14 = PendSV 15 = SysTick 16 = IRQ0 ... 47 = IRQ31

4.8.2 MCU Control Registers

Table 4-11. List of Registers

Register Name	Name	Description
CPUSS_CM0_CTL	CM0+ control	Specifies the Cortex-M0+ operation.
CPUSS_CM0_STATUS	CM0+ status	Indicates status of the Cortex-M0+, such as power mode.
CPUSS_CM7_x_CTL	CM7_x control	Specifies the operation of the Cortex-M7
CPUSS_CM7_x_STATUS	CM7_x status	Indicates status of the Cortex-M7, such as TCM access and power mode.
CPUSS_ECC_CTL	ECC control	Specifies the word address and ECC parity where an error will be injected.
CPUSS_CM7_x_PWR_CTL	CM7_x power control	Controls the CM7_x power state.
CPUSS_CM7_x_PWR_DELAY_CTL	CM7_x power delay control	Number of clock cycle delays needed after power domain powerup.

Note: The 'x' in the register name denotes the Cortex-M7 core number.

5. Inter-Processor Communication



Inter-processor communication (IPC) provides the functionality for multiple processors to communicate and synchronize their activities. IPC hardware is implemented using two register structures.

- **IPC Channel:** Communication and synchronization between processors is achieved using this structure.
- **IPC Interrupt:** Each interrupt structure configures an interrupt line, which can be triggered by a 'notify' or 'release' event of any IPC channel.

5.1 Features

The features of IPC are as follows:

- Implements locks for mutual exclusion between processors
- Allows sending messages between processors
- Supports multiple channels for communication
- Supports multiple interrupts, which can be triggered using notify or release events from the channels

5.1.1 IPC Channel

An IPC channel is implemented as six hardware registers, as shown in [Figure 5-1](#). The IPC channel registers are accessible to all processors in the system.

- **IPC_STRUCTx_ACQUIRE:** This register determines the lock feature of the IPC. The IPC channel is acquired by reading this register. If the SUCCESS field returns a '1', the read acquired the lock.

If the SUCCESS field returns a '0', the read did not acquire the lock.

Note that a single read access performs two functions:


- The attempt to acquire a lock.
- Return the result of the acquisition attempt (SUCCESS field).

The atomicity of these two functions is essential in a multi-core system with multiple CPUs.

The register also has bit fields that provide information about the processor that acquired it. When acquired, this register is released by writing any value into the IPC_STRUCTx_RELEASE register. If the register was already in an acquired state another attempt to read the register will not be able to acquire it.

- **IPC_STRUCTx_NOTIFY:** This register is used to generate an IPC notify event. Each bit in this register corresponds to an IPC interrupt structure. The notify event generated from an IPC channel can trigger multiple interrupt structures.
- **IPC_STRUCTx_RELEASE:** Any write to this register will release the IPC channel. This register also has a bit that corresponds to each IPC interrupt structure. The release event generated from an IPC channel can trigger multiple interrupt structures. To only release the IPC channel and not generate an interrupt, the user can write a zero into the IPC release register.
- **IPC_STRUCTx_DATA0 and IPC_STRUCTx_DATA1:** These two 32-bit registers are meant to hold data. They can be considered as the shared data memory for the channel. Typically, these registers will hold messages that need to be communicated between processors. If the messages are larger than the 32-bit size, the user can place a pointer in the IPC_STRUCTx_DATA0 or IPC_STRUCTx_DATA1 register.
- **IPC_STRUCTx_LOCK_STATUS:** This register provides the instantaneous lock status for the IPC channel. If the channel is acquired, this register provides details such as processor ID and protection context. The reading of lock status only provides an instantaneous status, which can be changed in the next cycle based on the activity of other processors on the channel.

Figure 5-1. IPC Channel Structure



5.1.2 IPC Interrupt

Each IPC interrupt line in the system has a corresponding IPC interrupt structure. An IPC interrupt can be triggered by a notify or a release event from any of the IPC channels in the system. The user can choose to mask any of the sources of these events using the IPC interrupt registers. Figure 5-2 shows the registers in an IPC Interrupt structure.


IPC_INTR_STRUCTx_INTR: This register provides the instantaneous status of the interrupt sources. Note that there are 16 notify and 16 release event bits in this registers. These are the notify and release events corresponding to the maximum 16 IPC channels. When a notify event is triggered in the IPC channel 0, the corresponding Notify0 bit is activated in the interrupt registers. A write of '1' to a bit will clear the interrupt.

IPC_INTR_STRUCTx_INTR_MASK: The bit in this register masks the interrupt sources. Only the interrupt sources with their masks enabled can trigger the interrupt.

IPC_INTR_STRUCTx_INTR_SET: A write of '1' into this register will set the interrupt.

IPC_INTR_STRUCTx_INTR_MASKED: This register provides the instantaneous value of the pending interrupts after they are masked. The value in this register is the result of the logical AND of registers IPC_INTR_STRUCTx_INTR and IPC_INTR_STRUCTx_INTR_MASK.

Figure 5-2. IPC Interrupt Structure




5.1.3 IPC Channels and Interrupts

The IPC block has a set of IPC interrupts associated with it. Each IPC interrupt register structure corresponds to an IPC interrupt line. This interrupt can trigger an interrupt on any of the processors in the system. The interrupt routing for processors are dependent on the device architecture.

Each IPC channel has a release and notify register, which can drive events on any of the IPC interrupts. Figure 5-3 shows an illustration of this relation between the IPC channels and the IPC interrupt structure.

Figure 5-3. IPC Channels and Interrupts



5.2 Implementing Locks

The IPC channels can be used to implement locks. Locks are typically used in multi core systems to implement some form of mutually exclusive access to a shared resource. When multiple processors share a resource, the processors are capable of acquiring and releasing the IPC channel. The processor can assume an IPC channel as a lock. The semantics of this code is that the access to the shared resource is gated by the processor's ownership of the channel. The processors need to acquire the IPC channel before they access the shared resource.

A failure to acquire the IPC channel signifies a lock on the shared resource because another processor has control of it. Note that the IPC channel will not enforce who acquires or releases the channel. All processors can acquire or release the IPC channel and the semantics of the code must make sure that the processor that acquires the channel is the one that releases it.

5.3 Message Passing

IPC channels can be used to communicate messages between processors. In this use case, the channel is used in conjunction with the interrupt structures. The IPC channel is used to lock the access to the Data register. The IPC channel is acquired by the sender and used to populate the message. The receiver reads the message and then releases the channel. Thus, between the sender placing data into the channel and receiver reading it, the channel is locked for all other tasks. The sender uses a notify event on the receiver's IPC interrupt to denote a send operation. The receiver acts on this interrupt and reads the data from the data register. After the reception is complete, the receiver releases the channel and can also generate a release event to the sender's IPC interrupt. Note that the action of locking the channel does not, in hardware, restrict access to the data register. This is a semantic that should be enforced by software.

Figure 5-4 portrays an example of a sender (Processor A) sending data to a receiver (Processor B). IPC interrupt A is configured to interrupt Processor A. IPC interrupt B is configured to interrupt Processor B.


1. The sender will attempt to acquire the IPC channel by reading the IPC_STRUCTx_ACQUIRE register until the SUCCESS bit is set. Then, the sender has ownership of the channel for data transmission.
2. After the IPC channel is acquired, the sender has control of the channel for communication and places message data up to 64 bits in the IPC_STRUCTx_DATA0 and IPC_STRUCTx_DATA1 registers.
3. Now that the message is placed in the IPC channel, the sender generates a notify event on the receiver's interrupt line. It does this by setting the corresponding bit in the IPC channel's IPC_STRUCTx_NOTIFY register. This event creates a notify event at IPC interrupt B. If the IPC channel's notify event was enabled by setting the mask bit in the IPC interrupt B, this will generate an interrupt in the receiver.
4. When it receives IPC interrupt B, the receiver can read the IPC_INTR_STRUCTx_INTR_MASKED register to

understand which IPC channel had triggered the notify event. Based on this, the receiver identifies the channel to read and reads from the IPC channel's IPC_STRUCTx_DATA0 and IPC_STRUCTx_DATA1 registers. The receiver has now received the data sent by the sender. It needs to release the channel so that other processors/processes can use it.

5. The receiver releases the channel. It also optionally generates a release event on the sender's IPC interrupt A. This will generate a release event interrupt on the sender if the corresponding channel release event was not masked.

On receiving the release interrupt, the sender can act on the event based on the application requirement. It can either try and reacquire the channel for further transmission or go on to other tasks because the transmission is complete.


Figure 5-4. Sending Messages using IPC



In the previous example, the size of the data being transmitted was just 32 bits. Larger messages can be sent as pointers. The sender can allocate a larger message structure in memory and pass the pointer in one of the 32-bit data registers. [Figure 5-5](#) illustrates this usage. Note that the user code should implement the synchronization of the message read process.

- The implementation can stall the channel until the receiver has used up all the data in the message packet and the message packet can be rewritten. This is wasteful because it will stall other inter-processor communications as the number of IPC channels is limited.
- The receiver can release the channel as soon as it receives the pointer to the message packet. It implements the synchronization logic in the message packet as a flag, which the sender sets on write complete and receiver clears on a read complete.

Figure 5-5. Communicating Larger Messages



5.4 Registers

Register	Name	Description
IPC_STRUCTx_ACQUIRE	IPC Structure Lock Acquire Register	This register is used to configure and acquire the lock
IPC_STRUCTx_RELEASE	IPC Structure Lock Release Register	This register is used to release the lock
IPC_STRUCTx_NOTIFY	IPC Structure Notification register	This register is used to generate notifications for the interrupt structure
IPC_STRUCTx_DATA0	IPC Structure Data Register 0	This field holds a 32-bit data element that is associated with the IPC structure
IPC_STRUCTx_DATA1	IPC Structure Data Register 1	This field holds a 32-bit data element that is associated with the IPC structure
IPC_STRUCTx_LOCK_STATUS	IPC Structure Lock Status Register	This register shows the status of the IPC (Lock Status, Access Mode, and so on)
IPC_INTR_STRUCTx_INTR	IPC Interrupt Status Register	This register shows the status of the interrupts
IPC_INTR_STRUCTx_INTR_SET	IPC Interrupt Set Register	Writing to this register sets the corresponding IPC_INTR_STRUCTx_INTR
IPC_INTR_STRUCTx_INTR_MASK	IPC Interrupt Mask Register	This is the mask bit for corresponding bit in IPC_INTR_STRUCTx_INTR
IPC_INTR_STRUCTx_INTR_MASKED	IPC Masked Interrupt Register	This register is the bitwise AND of INTR and INTR_MASK

Note: In IPC_STRUCTx or IPC_INTR_STRUCT, 'x' signifies the IPC instance.

6. Protection Unit



Protection units in the Traveo II series device enforce security based on different operations. A protection unit allows or restricts bus transfers on the bus infrastructure. The rules are enforced based on specific properties of a transfer.

6.1 Features


- An address range that is accessed by the transfer
 - Subregion: An address range is partitioned into eight equally-sized subregions and subregion can individual disables
- Access attributes such as:
 - Read/write attribute
 - Execute attribute to distinguish a code access from a data access
 - User/privilege attribute to distinguish access; for example, OS/kernel access from a task/thread access
 - Secure/non-secure attribute to distinguish a secure access from a non-secure access; the Arm Cortex-M CPUs do not natively support this attribute
 - A protection context attribute to distinguish accesses from different protection contexts; for Peripheral-DMA (P-DMA) and Memory-DMA (M-DMA), this attribute is extended with a channel identifier, to distinguish accesses from different channels
- Memory protection
- Provided by memory protection units (MPUs) and shared memory protection units (SMPUs)
 - MPUs distinguish user and privileged accesses from a single bus master
 - SMPUs distinguish between different protection contexts and between secure and non-secure accesses
- Peripheral protection
 - Provided by peripheral protection units (PPUs)
 - The PPU distinguish between different protection contexts; they also distinguish secure from non-secure accesses and user mode accesses from privileged mode accesses
- Protection pair structure
- Software Protection Unit (SWPU): SWPUs define flash write (or erase) permissions, and eFuse read and write permissions. An SWPU comprises of the following:
 - Flash Write Protection Unit (FWPU)
 - eFuse Read Protection Unit (ERPU)
 - eFuse Write Protection Unit (EWPU)

6.2 Configuration

6.2.1 Block Diagram

Figure 6-1 gives an overview of the location of MPUs, SMPUs, and PPUs in the system.

Figure 6-1. Protection Unit Locations



6.2.2 Protection Unit Structure

As mentioned, the MPU, SMPU, and PPU protection functionality follows the Arm MPU definition:

- Multiple protection structures are supported.
- Each structure specifies an address range in the unified memory architecture and access attributes. Address range can be as small as 32 bytes.

A bus master may have a dedicated MPU. In a CPU bus master, the MPU is typically implemented as part of the CPU and under control of the OS/kernel. In a non-CPU bus master, the MPU is typically implemented as part of the bus infrastructure and under control of the OS/kernel of the CPU that “owns/uses” the bus master. If a CPU switches tasks or if a non-CPU switches ownership, the MPU settings are typically updated by OS/kernel software. The different MPU types are:

- An MPU that is implemented as part of the CPU. This type is found in the Arm CM0+ and CM7 CPUs.
- CM7 MPU has 16 regions for each core, and CM0+ MPU has eight regions.
- An MPU that is implemented as part of the bus infrastructure. This type is found in bus masters such as test controller. The definition of this MPU type follows the Arm MPU definition (in terms of memory region and

access attribute definition) to ensure a consistent software interface.

The P-DMA, M-DMA, and cryptography component do not have an MPU. Instead, these components inherit the access control attributes of the bus transfer that programmed the channel or component.

The definition of SMPU and PPU follows the MPU definition and adds the capability to distinguish accesses from different protection contexts (the MPU does not include support for a protection context). If security is required, the SMPU and possibly PPUs registers must be controlled by a secure CPU that enforces system-wide protection.

Note that a peripheral group PPU only needs to provide access control to the peripherals within a peripheral group (peripherals with a shared bus infrastructure).

A protection violation is caused by a mismatch between a bus transfer's address region and access attributes and the protection structures' address range and access attributes.

A bus transfer that violates a protection structure results in a bus error. For AHB-Lite transfers, the address of each transfer beat is matched with the protection structure address range. The first violating beat in a transfer results in a bus error. A protection violation results in a bus error and the bus transfer will not reach its target. An MPU or SMPU

violation that targets a peripheral will not reach the associated PPU. In other words, MPU and SMPU have a higher priority over PPU.

6.2.3 Master with Missing Access Attributes

Not all masters provide all access attributes that are associated with a bus transfer. Some examples are:

- None of the bus masters have a natively fixed protection context attribute. This needs to be set dynamically based on the task being executed by the bus master.
- The Arm Cortex-M7 and Arm Cortex-M0+ CPUs provide a user/privilege attribute, but do not provide a secure/non-secure attribute natively.

To ensure system-wide restricted access, missing attributes are provided by register fields. These fields may be set during the boot process or by the secure CPU.

- The PROT_SMPU_MSx_CTL.PC_MASK_15_TO_1[], PROT_SMPU_MSx_CTL.PC_MASK_0, and PROT MPUx_MS_CTL.PC[] register fields provide protection context functionality.
- The PROT_SMPU_MSx_CTL.P register field provides the user/privileged attribute for those masters that do not provide their own attribute.
- The PROT_SMPU_MSx_CTL.NS register provides the secure/non-secure attribute for those masters that do not provide their own attribute.
- Masters that do not provide an execute attribute have the execute attribute set to '0'.

The P-DMA controller, M-DMA controller and cryptography components inherit the access control attributes of the bus transfers that programmed the channels and component.

- The PROT_SMPU_MSx_CTL and PROT MPUx_MS_CTL registers are only present for present masters.
- The PROT MPUx_MS_CTL.PC_SAVED field (and associated protection context 0 functionality, which is discussed later in the chapter) is only present for the CM0+ master.
- The PROT_SMPU_MSx_CTL.P, PROT_SMPU_MSx_CTL.NS, PROT_SMPU_MSx_CTL.PC_MASK_15_TO_1[], and PROT_SMPU_MSx_CTL.PC_MASK_0 fields are not present for P-DMA, M-DMA and cryptography masters. The bus transfer attributes are inherited: from the master that owns the P-DMA and M-DMA channels that initiated the bus transfer.
- The PROT MPUx_MS_CTL register is not present for the P-DMA and M-DMA and cryptography masters. The protection context (PC) bus transfer attribute is provided through inheritance.

6.3 Protection Context

6.3.1 Protection Context Configuration

Each bus master has an PROT MPUx_MS_CTL.PC[3:0] protection context field. This is used as the protection context attribute for all bus transfers that are initiated by the master. The SMPUs and PPUs allow or restrict bus transfers based on the protection context attribute.

Multiple masters can share a protection context. For example, a CPU and a Crypto controlled by the CPU may share a protection context (the CPU and Crypto PC[] fields are the same). Therefore, the CPU and Crypto share the SMPU and PPU access restrictions.

A bus master protection context is changed by reprogramming the master's PROT MPUx_MS_CTL PC[] field.

Each bus master has an PROT_SMPU_MSx_CTL.PC_MASK_0 and PROT_SMPU_MSx_CTL.PC_MASK_15_TO_1, or PROT_SMPU_MSx_CTL.PC_MASK_0 only protection context mask field that identifies what protection contexts can be programmed for the bus master:

- Protection context field PROT MPUx_MS_CTL.PC[3:0]. This register is controlled by the associated bus master and has the same access restrictions as the bus master's MPU registers.
- Protection context mask field PROT_SMPU_MSx_CTL.PC_MASK_15_TO_1[] and PROT_SMPU_MSx_CTL.PC_MASK_0. This register is controlled by the secure CPU and has the same access restrictions as the SMPU registers.

The PROT_SMPU_MSx_CTL.PC_MASK_15_TO_1[], PROT_SMPU_MSx_CTL.PC_MASK_0 field is a field that specifies if the PROT MPUx_MS_CTL.PC[] field can be programmed with a specific protection context. Consider an attempt to program PROT MPUx_MS_CTL.PC[] to '3':

- If PROT_SMPU_MSx_CTL.PC_MASK_15_TO_1[19] is '1', PROT MPUx_MS_CTL.PC[] is set to '3'.
- If PROT_SMPU_MSx_CTL.PC_MASK_15_TO_1[19] is '0', PROT MPUx_MS_CTL.PC[] is not changed.

As mentioned, the SMPUs and PPUs allow/restrict bus transfers based on the protection context attribute. The protection context provides an indirection between a bus master and the SMPU and PPU protection. This allows a single bus master to take on different protection roles, simply by reprogramming the protection context field PROT MPUx_MS_CTL.PC[]. A change of protection contexts has limited CPU overhead, as the SMPU and PPUs do not have to be reprogrammed.

See the PERI_PC_NR in the datasheet for the number of available PCs.

6.3.2 Protection Context 0 and 1

Traveo II supports protection contexts to isolate software execution for security and safety purposes. Protection contexts are used to restrict access to memory and peripheral resources. Traveo II supports eight protection contexts (PCs).

Out of eight PCs, two PCs are treated special: the entry to special PCs 0 and 1 is hardware controlled. For each PC i , a programmable exception handler address is provided: `CPUSS_CM0_PCx_HANDLER.ADDR[31:0]`. A CPU exception handler fetch that returns a handler address that matches the programmed `CPUSS_CM0_PCx_HANDLER.ADDR[31:0]` address value, causes the CM0+ PC to be changed to PC x by hardware. However, if the current PC is already 0 or 1, the current PC is not changed (an attempt to change the PC actually results in an AHB-Lite bus error).

This ensures that CPU execution in PC 0 or 1 cannot be interrupted/preempted by CPU execution in another PC 0 or 1. In other words, CPU execution in PC 0 or 1 requires cooperative multi-tasking between the different PCs. This means that handover between different PCs are software scheduled or controlled. A security implementation requires PC software to clear information that it wants to keep confidential from other PC software.

Note that each of the protection special PCs 0 and 1 have dedicated `CPUSS_CM0_PC_CTL.VALID[x]` field to specify that the PC's exception handler address is provided through `CPUSS_CM0_PCx_HANDLER.ADDR[31:0]`. If a PC's exception handler address is not provided, the PC is treated as an ordinary PC (PCs 2, 3, ..., 7).

Note that the current PC "pc" and a saved PC "pc_saved" implement a two entry stack. The hardware pushes the current PC to the stack upon entry of a special exception handler and hardware pops the saved PC from the stack upon entry of an ordinary exception handler. An attempt to enter a special exception handler from a special exception handler with a different PC results in an AHB-Lite bus error (which causes the CPU to enter the bus fault exception handler). This scenario should not occur in a carefully designed cooperative multi-tasking software implementation.

Of the four special PCs, PC 0 is treated differently: It is the default PC value after a DeepSleep reset. It has unrestricted access. Therefore, the Cypress boot code software always starts execution in PC 0. The boot code software initializes the protection structures and initializes the `CPUSS_CM0_PCx_HANDLER` registers.

After initialization of the protection information, the access to the protection information itself is typically restricted for all other PCs (the boot code software deploys the restrictions) and the protection information provides specific restricted access to the other special PCs and ordinary PCs (PCs 2, 3, ..., 7).

6.4 Protection Structure

The MPU, SMPU, and PPU protection structure definition follows the Arm definition. Each protection structure is defined by:

- An address region
- Access control attributes

A protection structure is always aligned on a 32-byte boundary in the memory space. Two registers define a protection structure: `ADDR` (address register) and `ATT` (attribute register). This alignment and organization allow straightforward protection of the protection structures by the protection scheme. This is elaborated upon later in this chapter.

As an `ADDR` register, MPU has `PROT_MPUx_MPU_STRUCTUREy_ADDR`, SMPU has `PROT_SMPU_SMPU_STRUCTUREx_ADDR0/1`, and PPU has `PERI_MS_PPU_PRx_SL/MS_ADDR` and `PERI_MS_PPU_FXx_SL/MS_ADDR`.

As an `ATT` register, MPU has `PROT_MPUx_MPU_STRUCTUREy_ATT`, SMPU has `PROT_SMPU_SMPU_STRUCTUREx_ATT0/1`, and PPU has `PERI_MS_PPU_PRx_SL_ATT0-3`, `PERI_MS_PPU_PRx_MS_ATT`, `PERI_MS_PPU_FXx_SL_ATT0-3`, `PERI_MS_PPU_FXx_MS_ATT`, `PERI_MS_PPU_PRx_SL/MS_SIZE`, and `PERI_MS_PPU_FXx_SL/MS_SIZE`.

6.4.1 Address Region

The address region is defined by:

- The base address of a region as specified by `ADDR.ADDR`.
- The size of a region as specified by `ATT.REGION_SIZE`.
- Individual disables for eight subregions within the region, as specified by `ADDR.SUBREGION_DISABLE`.

The `REGION_SIZE` field specifies the size of a region. The region size is a power of 2 in the range of [256 B, 4 GB]. The base address `ADDR` specifies the start of the region, which needs to be aligned to the region size. A region is partitioned into eight equally-sized subregions. The `SUBREGION_DISABLE` field specifies individual enables for the subregions within a region.

For example, a `REGION_SIZE` of '8' specifies a region size of 512 bytes. If the start address is 0x1000:5400 (512-byte aligned), the region ranges from 0x1000:5400 to 0x1000:55ff. This region is partitioned into the following eight 64-byte subregions:

subregion 0 from 0x1000:5400 to 0x1000:543f

subregion 1 from 0x1000:5440 to 0x1000:547f

...

subregion 7 from 0x1000:55c0 to 0x1000:55ff

If the SUBREGION_DISABLE is 0x82 (bit fields 1 and 7 are '1'), subregions 1 and 7 are disabled; subregions 0, 2, 3, 4, 5, and 6 are enabled.

In addition, an ATT.ENABLED field specifies if the region is enabled. Only enabled regions participate in the protection matching process. Matching identifies if a bus transfer address is contained within an enabled subregion (SUBREGION_DISABLE) of an enabled region (ENABLED).

6.4.2 Access Control Attributes

The access attributes specify access control to the region (shared by all subregions within the region). Access control is performed using a transfer's access attributes. The following access control fields are supported:

- Control for read accesses in user mode (ATT.UR field).
- Control for write accesses in user mode (ATT.UW field).
- Control for execute accesses in user mode (ATT.UX field).
- Control for read accesses in privileged mode (ATT.PR field).
- Control for write accesses in privileged mode (ATT.PW field).
- Control for execute accesses in privileged mode (ATT.PX field).
- Control for secure access (ATT.NS field).
- Control for individual protection contexts (ATT.PC_MASK_15_TO_1[] and PC_MASK_0, with PC_MASK_0 always constant at '1'). This protection context control field is present for SMPU.

The execute and read access control attributes are orthogonal. Execute transfers are typically read transfers. To allow execute or read transfers in user mode, both ATT.UR and ATT.UX need to be set to '1'. To allow data and read transfers in user mode, only ATT.UR needs to be set to '1'.

In addition, the ATT.PC_MATCH control field is supported, which controls the "matching" and "access evaluation" processes. This control field is only present for the SMPU protection structures.

For example, only protection context 2 can access a specific address range and these accesses are restricted to read and write secure accesses in privileged mode. The access control fields are programmed as follows:

- ATT.UR is '0': read accesses in user mode not allowed.
- ATT.UW is '0': write accesses in user mode not allowed.
- ATT.UX is '0': execute accesses in user mode not allowed.
- ATT.PR is '1': read accesses in privileged mode allowed.
- ATT.PW is '1': write accesses in privileged mode allowed.

- ATT.PX is '0': execute accesses in privileged mode not allowed.
- ATT.NS is '0': secure access required.
- ATT.PC_MASK_15_TO_1[10] is '1', and ATT.PC_MASK_0 is '1': protection context 0 and 2 accesses enabled (all other protection contexts are disabled).
- ATT.PC_MATCH is '0': the ATT.PC_MASK_15_TO_1[] and PC_MASK_0 field is used for access evaluation.

Three separate access evaluation subprocesses are distinguished:

- A subprocess that evaluates the access based on read/write, execute, and user/privileged access attributes.
- A subprocess that evaluates the access based on the secure/non-secure attribute.
- A subprocess that evaluates the access based on the protection context index (only used by the SMPU and PPU when ATT.PC_MATCH is '0').

If all access evaluations are successful, access is allowed. If any process evaluation is unsuccessful, access is not allowed.

Matching the bus transfer address and access evaluation of the bus transfer (based on access attributes) are two independent processes:

- Matching process. For each protection structure, the process identifies if a transfer address is contained within the address range. This identifies the matching regions.
- Access evaluation process. For each protection structure, the process evaluates the bus transfer access attributes against the access control attributes.

A protection unit typically has multiple protection structures. It evaluates the protection structures in decreasing order. The first matching structure provides the access control attributes for the evaluation of the transfer's access attributes. In other words, higher-indexed structures take precedence over lower-indexed structures.

Notes:

- If no protection structure provides a match, access is allowed.
- If multiple protection structures provide a match, the access control attributes for access evaluation are provided by the protection structure with the highest index.

As mentioned, the protection unit evaluates the protection structures in decreasing order. From a security requirements perspective, this is of importance: it should not be possible for a non-secure protection context to add protection structures that have a higher index than the protection structures that provide secure access. The protection structure with a higher index can be programmed to allow

non-secure accesses. Therefore, in a secure system, the higher programmable protection structures are protected to only allow restricted accesses. For more details, see [Protection Structure Types on page 62](#).

6.4.3 Protection Violation

If an MPU, SMPU, or PPU detects a not-allowed transfer, the bus transfer results in a bus error. Protection violations are captured in the fault report structure, and the fault report structures can generate an interrupt to indicate the occurrence of a fault. In addition, information on the violating bus transfer is communicated to the fault report structure. This is useful if the violating bus master cannot resolve the bus error by itself, but requires another CPU bus master to resolve the bus error on its behalf. Note that violating CPUs react by execution of exceptions.

For details of exceptions, see the Arm documentation sets for Cortex-M7 and Cortex-M0+.

The bus transfer does not reach its target memory location or peripheral register. For write transfers that violate PPU protection, the bus master will not see the bus error if buffering is enabled (CPUSS_BUFF_CTL.WRITE_BUFF = 1). This is because the AHB-Lite bridges in the bus infrastructure will buffer the write transfer and send the OK response to masters. In this case, the system must depend on the fault reported by PPU.

6.4.4 Protection of Protection Structures

The MPU, SMPU, and PPU-based protection architecture is consistent and provides the flexibility to implement different system-wide protection schemes. Protection structures can be set once at boot time or can be changed dynamically during device execution. For example, a CPU RTOS can change the CPU's MPU settings; a secure CPU can change the SMPU and PPUs settings. From security requirements,

it is necessary to prevent reprogramming of the protection structure from a malicious attacker.

Registers of MPU, SMPU, and PPU are the same registers as other peripherals. Furthermore, the protection structure itself can be included in the address range of the protection structure. That is, protects the protection structure by protection structure.

The first (slave) protection structure protects the resource and the second (master) protection structure protects the protection (address range of the second protection structure includes both the master and slave protection structures). We refer to the slave and master protection structures as a protection pair. Note that the address range of the master protection structure is known, that is, it is constant.


The protection architecture is flexible enough to allow for variations:

- Exclusive peripheral ownership can be shared by more than two protection contexts.
- The ability to change ownership is under control of a single protection context, and exclusive or non-exclusive peripheral ownership is shared by multiple protection contexts.

Note that in secure systems, typically a single secure CPU is used. In these systems, the ability to change ownership is assigned to the secure CPU at boot time and not dynamically changed. Therefore, it is strongly advised to assign the secure CPU its own, dedicated protection context.

Both PPU and SMPU are intended to distinguish between different protection contexts and to distinguish secure from non-secure accesses. Therefore, both PPU and SMPU protection use protection structure pairs. In the SMPU, the slave protection structure provides SMPU protection information and the master protection structure provides PPU protection information (the master and slave protection structures are registers).

Figure 6-2. Concept of Master and Slave Structure




6.4.5 MPU

The MPUs are associated to a single master. An MPU distinguishes user and privileged accesses from a single bus master. However, the capability exists to perform access control on the secure/non-secure attribute.

The MPU protection structures do not provide protection context control attributes.


Figure 6-3. MPU Functionality



6.4.6 SMPU

The SMPU is shared by all bus masters. The SMPU distinguishes between different protection contexts; it also distinguishes secure from non-secure accesses and user mode from privileged mode accesses.

Figure 6-4. SMPU Functionality



Single set of SMPU region structures provides the same protection information to all SMPUs in the systems.

The perform access control function on the user/privileged mode attribute in SMPU must not be used for accesses performed by the CM7 CPUs. AXI accesses performed in user mode of CM7 CPU may be marked as privileged accesses. See section 5.4.3 of the [Arm Cortex-M7 Technical Reference Manual](#) for more details.


Note that there are no SMPUs on the AHBP ports. Therefore, SMPU must not be used to protect the peripherals address range. Peripherals are protected by PPU.

6.4.7 PPU

The PPUs are situated in the peripheral block and are associated with a peripheral group (peripherals with a shared AHB-Lite bus infrastructure). A PPU is shared by all bus masters. The PPU distinguishes between different protection contexts; it also distinguishes secure from non-secure accesses and user mode from privileged mode accesses.

The minimum region size of the MPU and S MPU is 10 Bytes, but PPU can set the region size of at least 4 Bytes. (Register size)

Figure 6-5. PPU Functionality



Compared to the MPU and S MPU, a PPU have large number of protection regions. However, regions protected by PPU are mostly known address. Therefore, there are two types of PPU.

■ Fixed PPU structure

- To protect resources with a known address range. In other words, the ADDR.ADDR and SIZE.REGION_SIZE fields are fixed.

■ Programmable PPU structure

- To protect resources with an unknown address range, full programmability of a protection region's address range definition is used.

The programmable structure pairs slave address regions may overlap with other slave address regions. A transfer address is matched against all master and slave address regions. The master protection structures and programmable slave protection structures are given higher priority than the fixed slave protection structures. From high to low priority:

- Master protection structures
- Programmable slave protection structures
- Fixed slave protection structures

Note that programmable slave address regions have higher priority than fixed slave address regions.

The slave structures are programmed during the boot process when the PC is 0. For programmable protection structure pairs, this includes programming the slave address of the peripheral resource. After the boot process, slave address regions are not reprogrammed; only the master and slave attributes are reprogrammed. In other words, after the boot process, the protected peripheral resources are fixed and only the ownership of these resources (slave attributes) or the right to change the resource ownership (master attributes) are programmable/flexible.

Each protection structure provides ATT.NS, PW, PR, UW, UR access attributes for all PCs, except for PC 0. A PPU structure does not support PX and UX access attributes: peripheral transfer should have the Execute transfer attribute set to '0'. Note that execution from a peripheral address region is not allowed.

The programmable and fixed PPU structures are shared by all PERI master interfaces. Most of the protection information uses a single SRAM memory.

6.4.7.1 ECC for SRAM

The SRAM stored protection information is supported by ECC. This ECC supports single error correction and double error detection (SEDED). The ECC is applied to the RAM word bits and the word address that is used to access the SRAM. If correctable error (single bit) is detected during SRAM read operation, the ECC corrects the data. However, the corrected data is not updated into SRAM. If non-correctable error is detected, then the current AHB transfer is aborted. These errors are communicated through the fault reporting structure.

6.4.7.2 ECC Error Injection

The ECC faults can be debugged through an ECC parity injection mechanism.

- PERI_ECC_CTL.PARITY: ECC parity to use for ECC error injection at PERI_ECC_CTL.WORD_ADDR. Note that this field will be used by hardware only when ECC error injection is enabled by setting PERI_ECC_INJ_EN to '1'.
- PERI_ECC_CTL.ECC_INJ_EN: Enables error injection for PERI protection structure SRAM. If this is '1', the parity (ECC_CTL.PARITY) is used when a write is done to the PERI_ECC_CTL.WORD_ADDR of the SRAM.
- PERI_ECC_CTL.WORD_ADDR: Specifies the word address where the parity is injected. When write access to this SRAM address is detected and PERI_ECC_CTL.ECC_INJ_EN bit is '1', the parity (PERI_ECC_CTL.PARITY) is injected.

6.4.7.3 ECC Parity Generation by Software

To inject the ECC error for fault generation, ECC parity must be generated by software. Follow this procedure to generate 8-bit ECC parity.

1. Read present PPU attribute values of the target PPU structure ATT0-3 register.
 "x" indicates PPU_PROG or FIXED_STRUCT number.
2. Generate ACTUALWORD [74:0].
 $ACTUALWORD[74:0] = \{ \{ PC15_{\{NS, PW, PR, UW, UR\}} \} \dots \{ PC1_{\{NS, PW, PR, UW, UR\}} \} \}$
 Non-existing PC attributes values are set to '0'.
3. Calculate ADDR [10:0].

The ADDR can be calculated as follows.

PERI_MS_PPU_PRx_SL_ATT0-3: $x * 2$

PERI_MS_PPU_PRx_MS_ATT0-3: $x * 2 + 1$

PERI_MS_PPU_FXx_SL_ATT0-3: $x * 2 + 64$

PERI_MS_PPU_FXx_MS_ATT0-3: $x * 2 + 65$

The 'x' in the register name denotes the suffix number. ADDR [10:0] is set to PERI_ECC_CTL.WORD_ADDR.

4. Generate ECC parity using the following scheme.

```
CODEWORD_SW [127:0] = 128{1'b0};
CODEWORD_SW [74:0] = ACTUALWORD [74:0];
CODEWORD_SW [85: 75] = ADDR [10:0];

ECC_P0_SW = 128'b00000001_10111111_10111011_01110101_10111110_00111010_01110010_11011100
_01000100_10000100_01001010_10001000_10010101_00101010_10101101_01011011;
ECC_P1_SW = 128'b00000010_11011111_01110110_11111001_11011101_10011001_10111001_01110001
_00010001_00001000_10010011_00010001_00100110_10110011_00110110_01101101;
ECC_P2_SW = 128'b00000100_11101111_11001111_10011111_10011010_11010101_11001110_10010111
_00000110_00010001_00011100_00100010_00111000_11000011_11000111_10001110;
ECC_P3_SW = 128'b00001000_11110111_11101100_11110110_11101101_01100111_01001110_01101100
_10011000_00100001_11100000_01000011_11000000_11111100_00000111_11110000;
ECC_P4_SW = 128'b00010000_11111011_01111011_10101111_01101011_10100110_10101011_10100110
_11100000_00111110_00000000_01111100_00000000_11111111_11111000_00000000;
ECC_P5_SW = 128'b00100000_11111101_10110111_11001110_11110011_01101100_10101011_01011011
_11111111_11000000_00000000_01111111_11111111_00000000_00000000_00000000;
ECC_P6_SW = 128'b01000000_11111110_11011101_01111011_01110100_11011011_01010101_10101011
_11111111_11111111_11111111_10000000_00000000_00000000_00000000_00000000;
ECC_P7_SW = 128'b10000000_01111111_00000000_00000000_00000111_11111111_11111111_11111111
_11010100_01000010_00100101_10000100_01001011_10100110_01011100_10110111;

PARITY[0] = ^ (CW_SW[127:0] & ECC_P0_SW)
PARITY[1] = ^ (CW_SW[127:0] & ECC_P1_SW)
...
PARITY[7] = ^ (CW_SW[127:0] & ECC_P7_SW)
```

Note: “^” means reduction XOR. For example, $^(4'b0011) = 0^0^1^1$.

ECC parity is set to PERI_ECC_CTL.PARITY.

5. Set the PERI_ECC_CTL.ECC_INJ_EN to '1'.

6. Read and write back with the same value to the target PPU structure ATT0-3.

A write back will inject parity value from the ECC_CTL.PARITY[7:0] register to the SRAM PPU structure.

7. Read the target PPU structure ATT0-3.

A read will generate an ECC error.

6.4.8 Protection Structure Types

Different protection structure types are used because some resources, such as peripheral registers, have a fixed address range. For security, protection structures require pairs of neighboring protection structures.

Three types of protection structures, which have a consistent register interface are described here:

- Programmable protection structures – These structures are used by the MPUs.
- Fixed protection structure pairs. These structures are used by the PPU. Both structures have a fixed, constant address region and do not have the ATT.UX and ATT.PX attributes. In addition, PC0 is permitted with all access attributes (PC0 has unrestricted access). The master structure has the ATT.UR and ATT.PR attributes as constant '1' (reading is always allowed). See [Figure 6-6](#).
- Programmable protection structure pairs. These structures are used by the PPU and SMPU. The master structure has a fixed, constant address region. The slave structure has a programmable address region. The SMPU master structure has the ATT.UX and ATT.PX attributes as constant '0' (execution is never allowed) and the ATT.UR and ATT.PR attributes as constant '1' (reading is always allowed). The Both PPU structures do not have ATT.UX and ATT.PX attributes. In addition, PC0 is permitted with all access attributes (PC0 has unrestricted access). The PPU master structure has the ATT.UR and ATT.PR attributes as constant '1' (reading is always allowed). See [Figure 6-7](#), [Figure 6-8](#).

Note that the master protection structure in a protection structure pair is only required to address security requirements.

Figure 6-6. Fixed Protection Structure Pair




Figure 6-7. PPU Programmable Protection Structure Pair





Figure 6-8. SMPU Programmable Protection Structure Pair



6.5 SWPU

SWPU is used to implement access restrictions to flash (program/erase) and eFuse (read/write) and stored in SFlash.

This feature prevents malicious or unintended modification of flash or eFuse, and reading of sensitive eFuse data. In addition, unauthorized changes to SWPU are detected by the secure boot operation.

The SWPU is broken into two parts. The first part implements the access restrictions related to PC1 and PCx and cannot be updated. The second part is used by the application for additional access restrictions specific to the

application. ROM/flash boot reads the two parts of SWPU from SFlash and stores them in RAM.

SWPU comprises of FWPU, ERPU, EWPU. SWPU has slave and master protection structures as a protection pair, similar to the SMPU and PPU.

6.5.1 SWPU Layout

The SWPU is located at the address specified by TOC2_APP_PROTECTION_ADDR of TOC2 in SFlash. (The default address is 0x1700_7600.) FWPU has up to 16 regions, and ERPU and EWPU have up to four regions. [Table 6-1](#) lists the SWPU layout.

Table 6-1. SWPU Layout in SFlash

SWPU	Name	Size	Description
–	PU_OBJECT_SIZE	4 bytes	The total byte number of configured elements.
FWPU	N_FWPU[3:0]	4 bytes	The number of FWPU objects. FWPU has up to 16 regions.
	FWPU0_SL [3:0]	4 bytes	Configures the base address.
	FWPU0_SIZE [3:0]	4 bytes	Configures the size of protection area from FWPU_SL.
	FWPU0_SL_ATT [3:0]	4 bytes	Configures the slave attribute. This element sets the attribute for write access to Flash memory.
	FWPU0_MS_ATT [3:0]	4 bytes	Configures the master attribute. This element sets the attribute to configure the FWPU0_SL_ATT.
	:	–	Up to 16 regions.

Table 6-1. SWPU Layout in SFlash

SWPU	Name	Size	Description
ERPU	N_ERPU[3:0]	4 bytes	The number of FWPU objects. ERPU has up to four regions.
	ERPU0_SL_OFFSET_[3:0]	4 bytes	Configures the offset from eFuse base address.
	ERPU0_FUSE_SIZE_[3:0]	4 bytes	Configures the size of protection area from ERPU0_SL_OFFSET.
	ERPU0_SL_ATT_[3:0]	4 bytes	Configures the slave attribute. This element sets the attribute for read access from eFuse.
	ERPU0_MS_ATT_[3:0]	4 bytes	Configures the master attribute. This element sets the attribute to configure the ERPU0_SL_ATT.
	:	-	Up to four regions.
EWPU	N_EWPU[3:0]	4 bytes	The number of FWPU objects. EWPU has up to four regions.
	EWPU0_SL_OFFSET_[3:0]	4 bytes	Configures the offset from eFuse base address.
	EWPU0_FUSE_SIZE_[3:0]	4 bytes	Configures the size of protection area from ERPU0_SL_OFFSET.
	EWPU0_SL_ATT_[3:0]	4 bytes	Configures the slave attribute. This element sets the attribute for write access to eFuse.
	EWPU0_MS_ATT_[3:0]	4 bytes	Configures the master attribute. This element sets the attribute to configure the EWPU0_SL_ATT.
	:	-	Up to four regions.

Each element is described here. The suffix 'x' indicates FWPU region number.


- PU_OBJECT_SIZE: This element defines the total byte number of configured elements, which includes 4 bytes of PU_OBJECT_SIZE. Note that SWPU consists of up to 512 bytes. Blanks cannot be inserted between elements of each protection unit.
- N_FWPU[3:0], N_ERPU[3:0], N_EWPU[3:0]: These elements define the number of each protection unit. There are no EWPU objects if set to '0'. Note that the maximum number of areas for each unit cannot be exceeded.
- FWPUx_SL_[3:0]: This element sets the base address of the Flash memory to be protected by FWPU. The absolute 32-bit address needs to be specified. Also, the last two bits should be 0 for alignment purposes.
- ERPUx_SL_OFFSET_[3:0], EWPUx_SL_OFFSET_[3:0]: These elements set the offset from the eFuse base address to be protected by ERPU or EWPU.
- FWPUx_SIZE_[3:0], ERPUx_FUSE_SIZE_[3:0], EWPUx_FUSE_SIZE_[3:0]: These elements set the area size to be protected by each protection unit. The MSb indicates that the region is enabled when set to '1'. [Figure 6-9](#) shows the composition of each element.

Figure 6-9. Composition of Size Elements



- FWPUx_SL_ATT_[3:0], ERPUx_SL_ATT_[3:0], EWPUx_SL_ATT_[3:0]: These elements set the attribute to access Flash memory or eFuse. [Figure 6-10](#) shows the composition of each attribute setting element.

Figure 6-10. Composition of Attribute Elements



- UR: Read accesses in user mode are allowed for ERPU, when this bit sets to '1'.
- PR: Read accesses in privileged mode are allowed for ERPU, when this bit set to '1'.
- UW: Write accesses in user mode are allowed for FWPU and EWPU, when this bit sets to '1'.
- PW: Write accesses in privileged mode are allowed for FWPU and EWPU, when this bit set to '1'.
- NS: Non-secure accesses are allowed, when this bit set to '1'.
- PC_MASK: Accesses with protection context (PC) are allowed, when the corresponding bit is set to '1'. Traveo II has eight PCs. Therefore, MC_MASK[15:8] is invalid.
- FWPUx_MS_ATT_[3:0], ERPUx_MS_ATT_[3:0], EWPUx_MS_ATT_[3:0]: These elements set the attribute to access slave elements. Composition of each attribute setting element is same as [Figure 6-10](#).

6.5.2 SWPU Configuration

The following example shows an SWPU configuration of two FWPU, one ERPU, and one EWPU.

In the first FWPU (Region0), base address is 0x10000000, size is 0x1000, all PCs allow full write access. In the second FWPU (Region1), base address is 0x10008000, size is 0x8000, all PCs allow only write access with privileged. The master attribute for both FWPU is that all PCs allow full access. The configuration of each FWPU element is as follows:

- N_FWPU[3:0]: When two FWPU are added, N_FWPU0 should be 0x02. N_FWPU1, N_FWPU2, and N_FWPU3 should be 0x00.
- FWPU0_SL[3:0]: The first FWPU base address is 0x10000000. Therefore, FWPU0_SL_0, FWPU0_SL_1, FWPU0_SL_2, and FWPU0_SL_3 correspond to 0x00 0x00 0x00 0x10.
- FWPU0_SIZE[3:0]: The first FWPU size is 0x1000. Therefore, FWPU0_SIZE_0, FWPU0_SIZE_1, FWPU0_SIZE_2, and FWPU0_SIZE_3 correspond to 0x00 0x10 0x00 0x80. Note that the FWPU0_SIZE_3 is 0x80 because the MSb indicates that the region is enabled.
- FWPU0_SL_ATT[3:0]: The first FWPU attribute is all PCs allow full access. Therefore, FWPU0_SL_ATT_0, FWPU0_SL_ATT_1, FWPU0_SL_ATT_2, and FWPU0_SL_ATT_3 correspond to 0x07 0x00 0xFF 0x00.
- FWPU0_MS_ATT[3:0]: The first FWPU master attribute is all PCs allow full access. Therefore, FWPU0_MS_ATT_0, FWPU0_MS_ATT_1, FWPU0_MS_ATT_2, and FWPU0_MS_ATT_3 correspond to 0x07 0x00 0xFF 0x00.

The second FWPU is configured as follows:

- FWPU1_SL[3:0]: The first FWPU base address is 0x10008000. Therefore, FWPU1_SL_0, FWPU1_SL_1, FWPU1_SL_2, and FWPU1_SL_3 correspond to 0x00 0x80 0x00 0x10.
- FWPU1_SIZE[3:0]: The first FWPU size is 0x8000. Therefore, FWPU1_SIZE_0, FWPU1_SIZE_1, FWPU1_SIZE_2, and FWPU1_SIZE_3 correspond to 0x00 0x80 0x00 0x80. Note that the FWPU0_SIZE_3 is 0x80 because the MSb indicates that the region is enabled.
- FWPU1_SL_ATT[3:0]: The first FWPU attribute is all PCs allow only access with privileged. Therefore, FWPU1_SL_ATT_0, FWPU1_SL_ATT_1, FWPU1_SL_ATT_2, and FWPU1_SL_ATT_3 correspond to 0x06 0x00 0xFF 0x00.
- FWPU1_MS_ATT[3:0]: The first FWPU master attribute is all PCs allow full access. Therefore, FWPU1_MS_ATT_0, FWPU1_MS_ATT_1,

FWPU1_MS_ATT_2, and FWPU1_MS_ATT_3 correspond to 0x07 0x00 0xFF 0x00.

In this example, ERPU protects customer data in eFuse. The customer data is located at offset 0x68 from eFuse base address, and the size is 0x18, all PCs allow full read access. The master attribute for ERPU is all PCs allow full access. The configuration of ERPU element is as follows:

- N_ERPU[3:0]: When one ERPU is added, N_ERPU0 should be 0x01. N_ERPU1, N_ERPU2, and N_ERPU3 should be 0x00.
- ERPU0_SL_OFFSET[3:0]: The offset is 0x68. Therefore, ERPU0_SL_OFFSET_0, ERPU0_SL_OFFSET_1, ERPU0_SL_OFFSET_2, and ERPU0_SL_OFFSET_3 correspond to 0x68 0x00 0x00 0x00.
- ERPU0_FUSE_SIZE[3:0]: The size is 0x18. Therefore, ERPU0_SIZE_0, ERPU0_SIZE_1, ERPU0_SIZE_2, and ERPU0_SIZE_3 correspond to 0x18 0x00 0x00 0x80. Note that the ERPU0_SIZE_3 is 0x80 because the MSb indicates that the region is enabled.
- ERPU0_SL_ATT[3:0]: The attribute is all PCs allow full access. Therefore, ERPU0_SL_ATT_0, ERPU0_SL_ATT_1, ERPU0_SL_ATT_2, and ERPU0_SL_ATT_3 correspond to 0x07 0x00 0xFF 0x00.
- ERPU0_MS_ATT[3:0]: The master attribute is all PCs allow full access. Therefore, ERPU0_MS_ATT_0, ERPU0_MS_ATT_1, ERPU0_MS_ATT_2, and ERPU0_MS_ATT_3 correspond to 0x07 0x00 0xFF 0x00.

In the case of EWPU configuration, EWPU protects customer data in eFuse. Offset is 0x68, size is 0x18, and all PCs allow full read access. The master attribute for EWPU is all PCs allow full access. The configuration of the EWPU element is as follows:

- N_EWPU[3:0]: When one EWPU is added, N_EWPU0 should be 0x01. N_EWPU1, N_EWPU2, and N_EWPU3 should be 0x00.
- EWPU0_SL_OFFSET[3:0]: The offset is 0x68. Therefore, EWPU0_SL_OFFSET_0, EWPU0_SL_OFFSET_1, EWPU0_SL_OFFSET_2, and EWPU0_SL_OFFSET_3 correspond to 0x68 0x00 0x00 0x00.
- EWPU0_FUSE_SIZE[3:0]: The size is 0x18. Therefore, EWPU0_SIZE_0, EWPU0_SIZE_1, EWPU0_SIZE_2, and EWPU0_SIZE_3 correspond to 0x18 0x00 0x00 0x80. Note that the EWPU0_SIZE_3 is 0x80 because the MSb indicates that the region is enabled.
- EWPU0_SL_ATT[3:0]: The attribute is all PCs allow full access. Therefore, EWPU0_SL_ATT_0, EWPU0_SL_ATT_1, EWPU0_SL_ATT_2, and EWPU0_SL_ATT_3 correspond to 0x07 0x00 0xFF 0x00.

- EWPU0_MS_ATT_[3:0]: The master attribute is all PCs allow full access. Therefore, EWPU0_MS_ATT_0, EWPU0_MS_ATT_1, EWPU0_MS_ATT_2, and EWPU0_MS_ATT_3 correspond to 0x07 0x00 0xFF 0x00.

Table 6-2 lists the SWPU layout in SFlash in the above configuration.

Table 6-2. SWPU Layout After Configuration

Address	Element Name	Setting Value
0x17007600	PU_OBJECT_SIZE	0x50
0x17007601		0x00
0x17007602		0x00
0x17007603		0x00
0x17007604	N_FWPU0	0x02
0x17007605	N_FWPU1	0x00
0x17007606	N_FWPU2	0x00
0x17007607	N_FWPU3	0x00
0x17007608	FWPU0_SL_0	0x00
0x17007609	FWPU0_SL_1	0x00
0x1700760A	FWPU0_SL_2	0x00
0x1700760B	FWPU0_SL_3	0x10
0x1700760C	FWPU0_SIZE_0	0x00
0x1700760D	FWPU0_SIZE_1	0x10
0x1700760E	FWPU0_SIZE_2	0x00
0x1700760F	FWPU0_SIZE_3	0x80
0x17007610	FWPU0_SL_ATT_0	0x07
0x17007611	FWPU0_SL_ATT_1	0x00
0x17007612	FWPU0_SL_ATT_2	0xFF
0x17007613	FWPU0_SL_ATT_3	0x00
0x17007614	FWPU0_MS_ATT_0	0x07
0x17007615	FWPU0_MS_ATT_1	0x00
0x17007616	FWPU0_MS_ATT_2	0xFF
0x17007617	FWPU0_MS_ATT_3	0x00
0x17007618	FWPU1_SL_0	0x00
0x17007619	FWPU1_SL_1	0x80
0x1700761A	FWPU1_SL_2	0x00
0x1700761B	FWPU1_SL_3	0x10
0x1700761C	FWPU1_SIZE_0	0x00
0x1700761D	FWPU1_SIZE_1	0x80
0x1700761E	FWPU1_SIZE_2	0x00
0x1700761F	FWPU1_SIZE_3	0x80
0x17007620	FWPU1_SL_ATT_0	0x06
0x17007621	FWPU1_SL_ATT_1	0x00
0x17007622	FWPU1_SL_ATT_2	0xFF
0x17007623	FWPU1_SL_ATT_3	0x00
0x17007624	FWPU1_MS_ATT_0	0x07
0x17007625	FWPU1_MS_ATT_1	0x00

Table 6-2. SWPU Layout After Configuration

Address	Element Name	Setting Value
0x17007626	FWPU1_MS_ATT_2	0xFF
0x17007627	FWPU1_MS_ATT_3	0x00
0x17007628	N_ERPU0	0x01
0x17007629	N_ERPU1	0x00
0x1700762A	N_ERPU2	0x00
0x1700762B	N_ERPU3	0x00
0x1700762C	ERPU0_SL_OFFSET_0	0x68
0x1700762D	ERPU0_SL_OFFSET_1	0x00
0x1700762E	ERPU0_SL_OFFSET_2	0x00
0x1700762F	ERPU0_SL_OFFSET_3	0x00
0x17007630	ERPU0_FUSE_SIZE_0	0x18
0x17007631	ERPU0_FUSE_SIZE_1	0x00
0x17007632	ERPU0_FUSE_SIZE_2	0x00
0x17007633	ERPU0_FUSE_SIZE_3	0x80
0x17007634	ERPU0_SL_ATT_0	0x07
0x17007635	ERPU0_SL_ATT_1	0x00
0x17007636	ERPU0_SL_ATT_2	0xFF
0x17007637	ERPU0_SL_ATT_3	0x00
0x17007638	ERPU0_MS_ATT_0	0x07
0x17007639	ERPU0_MS_ATT_1	0x00
0x1700763A	ERPU0_MS_ATT_2	0xFF
0x1700763B	ERPU0_MS_ATT_3	0x00
0x1700763C	N_EWPU0	0x01
0x1700763D	N_EWPU1	0x00
0x1700763E	N_EWPU2	0x00
0x1700763F	N_EWPU3	0x00
0x17007640	EWPU0_SL_OFFSET_0	0x68
0x17007641	EWPU0_SL_OFFSET_1	0x00
0x17007642	EWPU0_SL_OFFSET_2	0x00
0x17007643	EWPU0_SL_OFFSET_3	0x00
0x17007644	EWPU0_FUSE_SIZE_0	0x18
0x17007645	EWPU0_FUSE_SIZE_1	0x00
0x17007646	EWPU0_FUSE_SIZE_2	0x00
0x17007647	EWPU0_FUSE_SIZE_3	0x80
0x17007648	EWPU0_SL_ATT_0	0x07
0x17007649	EWPU0_SL_ATT_1	0x00
0x1700764A	EWPU0_SL_ATT_2	0xFF
0x1700764B	EWPU0_SL_ATT_3	0x00
0x1700764C	EWPU0_MS_ATT_0	0x07
0x1700764D	EWPU0_MS_ATT_1	0x00
0x1700764E	EWPU0_MS_ATT_2	0xFF
0x1700764F	EWPU0_MS_ATT_3	0x00
:	-	Blank
0x170077FF		

6.6 Registers

Table 6-3. List of MPU Registers

Register	Name	Description
PROT_MPUx_MS_CTL	Master control register	Specify the protection context of the bus transfer
PROT_MPUx_MPU_STRUCTy_ADDR	MPU region address register	Defines a MPU address region.
PROT_MPUx_MPU_STRUCTy_ATT	MPU region attributes register	Defines a MPU access control register.

Note: The 'x' in the registers name denotes the master number and the 'y' denotes the region number.

Table 6-4. List of SMPU Registers

Register	Name	Description
PROT_SMPU_MSx_CTL	Protection context control register	Specify the protection context of the bus transfer.
PROT_SMPU_SMPU_STRUCTy_ADDR0	SMPU region address 0 (slave structure) register	Defines a SMPU address region (slave structure).
PROT_SMPU_SMPU_STRUCTy_ATT0	SMPU region attributes 0 (slave structure) register	Defines SMPU access control (slave structure).
PROT_SMPU_SMPU_STRUCTy_ADDR1	SMPU region address 1 (master structure) register	Defines a SMPU address region (master structure).
PROT_SMPU_SMPU_STRUCTy_ATT1	SMPU region attributes 1 (master structure) register	Defines SMPU access control (master structure).

Note: The 'x' in the registers name denotes the master number and the 'y' denotes the region number.

Table 6-5. List of PPU Registers

Register	Name	Description
PERI_MS_PPU_PRx_SL_ADDR	Programmable PPU slave region, base address register	Specifies the base address of the slave region.
PERI_MS_PPU_PRx_SL_SIZE	Programmable PPU slave region, size register	Specifies the size of the slave region and sets region enable. Typically, it is programmed by the boot process with protection context.
PERI_MS_PPU_PRx_SL_ATT0,1,2,3	Programmable PPU slave attributes 0, 1, 2, 3 register	Defines access control (slave structure).
PERI_MS_PPU_PRx_MS_ADDR	Programmable PPU master region, base address register	Specifies the base address of the master region. This register is fixed (non-programmable).
PERI_MS_PPU_PRx_MS_SIZE	Programmable PPU master region, size register	Specifies the size of the master region. This register is fixed (non-programmable).
PERI_MS_PPU_PRx_MS_ATT0,1,2,3	Programmable PPU master attributes 0, 1, 2, 3 register	Defines access control (master structure).
PERI_MS_PPU_FXx_SL_ADDR	Fixed PPU slave region, base address register	Specifies the base address of the slave region.
PERI_MS_PPU_FXx_SL_SIZE	Fixed PPU slave region, size register	Specifies the size of the slave region and sets region enable. Typically, it is programmed by the boot process with protection context.
PERI_MS_PPU_FXx_SL_ATT0,1,2,3	Fixed PPU slave attributes 0, 1, 2, 3 register	Defines access control (slave structure).

Table 6-5. List of PPU Registers

Register	Name	Description
PERI_MS_PPU_FXx_MS_ADDR	Fixed PPU master region, base address register	Specifies the base address of the master region. This register is fixed (non-programmable).
PERI_MS_PPU_FXx_MS_SIZE	Fixed PPU master region, size register	Specifies the size of the master region. This register is fixed (non-programmable).
PERI_MS_PPU_FXx_MS_ATT0,1,2,3	Fixed PPU master attributes 0,1,2,3 register	Defines access control (master structure).
PERI_ECC_CTL	ECC control register	Provides ECC support for the SRAM protection structures in the master interface peripherals (peripheral group 0, peripheral 1).

Note: The 'x' in the register name denotes the master number.

Table 6-6. List of Buffer Control Registers

Register	Name	Description
CPUSS_BUFF_CTL	Buffer control register	Specifies if write transfer can be buffered in the bus infrastructure bridges.

7. Direct Memory Access



The Traveo II device supports two kinds of DMA controllers: Peripheral DMA (P-DMA) and Memory DMA (M-DMA). P-DMA is used for peripheral-to-memory and memory-to-peripheral data transfers and provides low latency for a large number of channels. P-DMA controller uses a single data transfer engine that is shared by the associated channels. It supports independent accesses to peripherals using the AHB multi-layer bus. M-DMA is used for memory-to-memory data transfers and provides high memory bandwidth for a small number of channels. M-DMA uses a dedicated data transfer engine for each channel.

P-DMA and M-DMA have a similar register interface and are compared as follows:

- P-DMA focuses on peripheral-to-memory and memory-to-peripheral data transfers (but it can also perform memory-to-memory data transfers). M-DMA focuses on memory-to-memory data transfers (but it can also perform peripheral-to-memory and memory-to-peripheral data transfers).
- P-DMA focuses on achieving low latency for a large number of channels. M-DMA focuses on achieving high memory bandwidth for a small number of channels.
- P-DMA uses a single data transfer engine that is shared by all channels. M-DMA uses a dedicated data transfer engine for each channel.

Note: DW and P-DMA have the same meaning in this DMA chapter. Also, DMAC and M-DMA are the same. Register names are labeled DW and DMAC.

7.1 Peripheral DMA (P-DMA)

P-DMA is used to transfer data between memory and peripherals without CPU involvement: the CPU configures/programs the P-DMA but the actual transfer is done by the P-DMA controller. The primary design target is P-DMA functionality at limited area overhead to the platform. Functionally, the P-DMA controller is similar to a general-purpose DMA controller.

7.1.1 Overview

The P-DMA controller is part of the CPUSS and controls data transfer between peripherals and memory. This controller can be configured/programmed to perform multiple independent data transfers. Each data transfer is managed by a channel. The number of channels varies for different part numbers; more details are available in the device datasheet.

A channel has an associated priority and is scheduled according to its priority.

A data transfer is initiated by an input trigger. This trigger may originate from the source of the transfer, destination of the transfer, CPU software, or from another SoC component. Triggers provide Active/Sleep functionality and are not available in DeepSleep and Hibernate power modes.

The data transfer specifics are specified by a descriptor. This descriptor specifies (among other things):

- The source and destination address locations and the size of the transfer.
- The actions of a channel; for example, generation of output triggers and interrupts.
- Data transfer types can be single, 1D, 2D, or CRC as defined in the descriptor structure. These types essentially define the address sequences generated for source and destination. 1D and 2D transfers are used for “scatter gather” and other useful transfer operations.

A channel's descriptor state is encoded as part of the channel's register state (and not as part of the descriptor). The following registers provide a channel's descriptor state:

- DWx_CH_STRUCTy_CH_CTL – This register provides generic channel control information.

- **DWx_CH_STRUCTy_CH_CURR_PTR** – This register provides the address of the memory location where the current descriptor is located. Software needs to initialize this register. Hardware sets this register to the current descriptor's next descriptor pointer, when advancing from the current descriptor to the next descriptor in a descriptor list.
- **DWx_CH_STRUCTy_CH_IDX** – This register provides the current X and Y indices of the channel into the current descriptor. Software needs to initialize this register. Hardware sets the X and Y indices to 0, when advancing from the current descriptor to the next descriptor in a descriptor list.

Note that channel state is retained in DeepSleep power mode.

The P-DMA controller is an Active/Sleep power mode functionality. Software should not initiate DeepSleep system power mode entry if there are any active P-DMA controller channels transferring data. Note that there is no way of capturing the active channel data while transitioning to DeepSleep system power mode.

7.1.2 Channels

P-DMA controller supports multiple independent data transfers that are managed by a channel. Each channel

connects to a specific system trigger through a trigger multiplexer that is outside the P-DMA controller. See the [Trigger Multiplexer chapter on page 493](#) for more details.

Channel priority. A channel is assigned a priority (DWx_CH_STRUCTy_CH_CTL.PRIO) between 0 and 3, with 0 being the highest priority and 3 being the lowest priority. Channels with the same priority constitute a priority group. Priority decoding determines the highest priority pending channel. This channel is determined as follows.

- The highest priority group with pending channels is identified first.
- Within this priority group, the following “round-robin” arbitration is applied. A “round” consists of a contiguous sequence of channel activations, within this priority group, without any repetition. Within a round, higher priority is given to the lower channel indices. The notion of a round guarantees that within a group, higher channel indices do not yield to lower indices indefinitely.

Channel state. At any given time, there is at most one channel that is actively performing a data transfer. This channel is called the active channel. A channel can be in one of four channel states.

Table 7-1. P-DMA Channel States

Channel State	Description
Disabled	The channel is disabled by setting DWx_CH_STRUCTy_CH_CTL.ENABLED to 0. The channel trigger is ignored in this state. Note: If an active channel is disabled by software, there should be no assumptions made about the state of the channel (current position of the transfer as reflected by the registers or descriptors). A software channel re-enable should prepare the new descriptors and reconfigure the channel.
Blocked	The channel is enabled and is waiting for a trigger to initiate a data transfer.
Pending	The channel is enabled and has received an active trigger. In this state, the channel is ready to initiate a data transfer but waiting for it to be scheduled.
Active	The channel is enabled, has received an active trigger, and has been scheduled. It is actively performing data transfers. If there are multiple channels pending, the highest priority pending channel is scheduled.

The data transfer associated with a trigger is made up of one or more “atomic transfers” or “single transfers”. For example a 1D transfer consists of X_COUNT+1 single transfers.

A channel may be marked preemptable (DWx_CH_STRUCTy_CH_CTL.PREEMPTABLE). If preemptable, and there is a higher priority pending channel, then that channel can preempt the current channel between single transfers.

A channel has two access control attributes that are SMPU and PPU for access control:

- Privileged Mode (DWx_CH_STRUCTy_CH_CTL.P) attribute can be set to privileged or user.
- Non-secure (DWx_CH_STRUCTy_CH_CTL.NS) attribute can be set to secure or non-secure.

A descriptor associated with each channel describes the data transfer. The descriptor is stored in memory and DWx_CH_STRUCTy_CH_CURR_PTR contains the descriptor address associated with channel “y”.

7.1.3 Descriptors

A descriptor is stored in memory and describes a data transfer. The descriptor is read-only for the P-DMA controller.

Descriptor Type (DESCR_TYPE) – There are four types of descriptors.

Table 7-2. P-DMA Descriptor Types

Descriptor Type	Description
Single transfer	This transfers a single data element (8-bit, 16-bit, or 32 bit) as shown in Figure 7-1 . The descriptor size is four 32-bit words: DESCR_CTL, DESCR_SRC, DESCR_DST, and DESCR_NEXT_PTR.
1D transfer	This performs a one-dimensional “for loop” (described in C) as shown Figure 7-2 . A 1D transfer is made up of X_COUNT+1 single transfers. The descriptor size is five 32-bit words: DESCR_CTL, DESCR_SRC, DESCR_DST, DESCR_X_CTL, and DESCR_NEXT_PTR.
2D transfer	This performs a two-dimensional “for loop” (described in C) as shown in Figure 7-3 . A 2D transfer is made up of (Y_COUNT+1) 1D transfers. The descriptor size is six 32-bit words: DESCR_CTL, DESCR_SRC, DESCR_DST, DESCR_X_CTL, DESCR_Y_CTL, and DESCR_NEXT_PTR.
CRC transfer	This performs a one-dimensional “for loop” similar to the 1D transfer. However, the source data is not transferred to a destination. Instead, a CRC is calculated over the source data as shown in Figure 7-4 . The CRC configuration is provided through a set of registers that is shared by all P-DMA channels and the assumption is that the P-DMA channels use the CRC functionality mutually exclusive in time. These registers are: DWx_CRC_CTL0, DWx_CRC_DATA_CTL0, DWx_CRC_POL_CTL0, DWx_CRC_LFSR_CTL0, DWx_CRC_REM_CTL0, and DWx_CRC_REM_RESULT0. Note that the CRC configuration is the same as the Crypto CRC configuration.

Figure 7-1. Single Transfer

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
DST_ADDR[0] = (t_DATA_SIZE) SRC_ADDR[0];
```

Figure 7-2. 1D Transfer

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (X_IDX = 0; X_IDX ≤ X_COUNT; X_IDX++) {
    DST_ADDR[X_IDX * DST_X_INCR] =
        (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR];
}
```


Figure 7-3. 2D Transfer

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (Y_IDX = 0; Y_IDX ≤ Y_COUNT; Y_IDX++) {
    for (X_IDX = 0; X_IDX ≤ X_COUNT; X_IDX++) {
        DST_ADDR[X_IDX * DST_X_INCR + Y_IDX * DST_Y_INCR] =
            (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR + Y_IDX * SRC_Y_INCR];
    }
}
```

Figure 7-4. CRC Transfer

```
// DST_ADDR is a pointer to an address location where the calculated CRC is stored.
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
CRC_STATE = CRC_LFSR_CTL;
for (X_IDX = 0; X_IDX ≤ X_COUNT; X_IDX++) {
    Update_CRC (CRC_STATE, (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR];
}
DST_ADDR = CRC_STATE;
```

The variables X_IDX and Y_IDX are stored in the channel register state (DWx_CH_STRUCTy_CH_IDX register).

The parameters X_COUNT, Y_COUNT, SRC_X_INCR, SRC_Y_INCR, DST_X_INCR, DST_Y_INCR, SRC_ADDR, DST_ADDR, SRC_TRANSFER_SIZE, DST_TRANSFER_SIZE, and DATA_SIZE are stored in the descriptor.

Descriptor Size – The size of a descriptor depends on its descriptor type. Only relevant parameters are stored. For example, a 1D descriptor does not contain the Y_COUNT, SRC_Y_INCR, and DST_Y_INCR parameters.

Transfer Size (SRC_TRANSFER_SIZE and DST_TRANSFER_SIZE) – In a data transfer, the source data is cast into the type specified by DATA_SIZE and assigned to the destination. The source type is determined by SRC_TRANSFER_SIZE and the destination type is determined by DST_TRANSFER_SIZE. All types are unsigned. All address computations use C semantics based on the transfer size.

Descriptor Chaining – Descriptors chained together. DESCR_NEXT_PTR field contains a pointer to the next descriptor in the chain. A channel executes the next descriptor in the chain when it completes executing the current descriptor. The last descriptor in the chain has DESCR_NEXT_PTR set to '0' (null pointer). A descriptor chain is also referred to as a descriptor list. It is possible to have a circular list in which case the execution continues indefinitely until there is an error or the channel or the controller is disabled by software.

Trigger-in Type (TR_IN_TYPE) – An input trigger initiates a data transfer and the TR_IN_TYPE defines the action on a trigger.

Table 7-3. P-DMA Trigger-in Types

Trigger Type	Description
Type 0	Trigger results in the execution of a single transfer. In a 1D or 2D transfer, this will execute a single transfer in the loop.
Type 1	Trigger results in the execution of a single 1D transfer. If the descriptor type is “single transfer” this behaves similar to type 0. If the descriptor type is 2D, it results in executing the inner loop once.
Type 2	Trigger results in the execution of the current descriptor.
Type 3	Trigger results in the execution of a descriptor list.

Trigger-out Type (TR_OUT_TYPE) – This defines when an output trigger is generated.

Table 7-4. P-DMA Trigger-out Types

Trigger Type	Description
Type 0	Output trigger is generated after a single transfer. In a 1D or 2D transfer, an output trigger is generated after each transfer in the loop.
Type 1	Output trigger is generated after a single 1D transfer. If the descriptor type is “single transfer”, this behaves similar to type 0. If the descriptor type is 2D, an output trigger is generated after each execution of the inner loop.
Type 2	Output trigger is generated after the execution of the current descriptor.
Type 3	Output trigger is generated after the execution of a descriptor list.

Interrupt Type (INTR_TYPE) – This defines when a completion interrupt is generated.

Table 7-5. P-DMA Interrupt Types

Trigger Type	Description
Type 0	Interrupt is generated after a single transfer. In a 1D or 2D transfer, an interrupt is generated after each transfer in the loop.
Type 1	Interrupt is generated after a single 1D transfer. If the descriptor type is single transfer, this behaves similar to type 0. If the descriptor type is 2D, an interrupt is generated after each execution of the inner loop.
Type 2	Interrupt is generated after the execution of the current descriptor.
Type 3	Interrupt is generated after the execution of a descriptor list.

Wait for Deactivation (WAIT_FOR_DEACT) – Specifies whether the P-DMA controller should wait for the input trigger to be deactivated after it has completed the data transfer corresponding to the current trigger. This field is used for level-sensitive triggers to give sufficient time for the triggering agent to deactivate the trigger. The wait specified can be 0, up to four cycles, up to 16 cycles, or indefinite. Pulse-sensitive triggers should have this field set to 0.

7.1.4 Interrupts

P-DMA can generate interrupts on completion and on various error conditions.

- The INTR_TYPE descriptor control defines when a completion condition (COMPLETION) is activated.
- The error conditions include SRC_BUS_ERROR, DST_BUS_ERROR, SRC_MISAL, DST_MISAL, CURR_PTR_NULL, ACTIVE_CH_DISABLED, and DESCR_BUS_ERROR.

The source of the interrupt is stored in DWx_CH_STRUCTy.CH_STATUS.INTR_CAUSE. INTR_TYPE defined in the descriptor controls when a completion interrupt is generated. Each channel has four interrupt related registers.

DWx_CH_STRUCTy_INTR – Each channel has an interrupt request register. Bit 0 is set 1 when interrupt event (completion or error) is detected. Software can clear this by writing to this bit.

DWx_CH_STRUCTy_INTR_SET – Each channel has an interrupt set register. Software can write 1 to this register to set the corresponding DWx_CH_STRUCTy_INTR register.

DWx_CH_STRUCTy_INTR_MASK – Each channel has an interrupt mask register. The corresponding interrupt is enabled by writing 1 to this register.

DWx_CH_STRUCTy_INTR_MASKED – Each channel has an interrupt masked register. When read, this register

reflects a bitwise AND between the interrupt request and mask registers.

The P-DMA is an Active power mode peripheral; this means, it uses Active functionality interrupts. Therefore, DWx_CH_STRUCTy_INTR and DWx_CH_STRUCTy_INTR_SET are not retained in DeepSleep power mode (DWx_CH_STRUCTy_INTR_MASK is retained).

7.1.5 P-DMA Controller Status Registers

The controller DWx_STATUS0 register contains the following information.

- DWx_STATUS0.ACTIVE - Active channel present, no/yes
- DWx_STATUS0.P - Active channel access control user/privileged
- DWx_STATUS0.NS - Active channel access control secure/non-secure


- DWx_STATUS0.B - Active channel access control non-bufferable/bufferable
- DWx_STATUS0.PC - Active channel protection context
- DWx_STATUS0.CH_IDX - Active channel index if there is an active channel
- DWx_STATUS0.PRIO - Active channel priority
- DWx_STATUS0.PREEMPTABLE - Active channel preemptable
- DWx_STATUS0.STATE - One of inactive, loading descriptor, loading data element, storing data element, or waiting for trigger deactivation

The DWx_CH_STRUCTy_CH_STATUS.PENDING register bit specifies whether the channel is currently pending or not.

7.1.6 P-DMA Controller Design

The following figure gives an overview of the P-DMA controller design.

Figure 7-5. P-DMA Controller Design



In this figure, the P-DMA controller output triggers are feedback as input triggers to the component. This feedback is accomplished outside of the component.

The following design components are distinguished:

- **Trigger selection.** This component is outside the P-DMA controller and connects each channel to one specific system trigger. This multiplexer layer allows a controller with a limited number of channels to support a larger number of system triggers. This is an important function as the controller's area scales with the number of channels (and to a lesser degree with the number of system triggers). This is because each channel requires a channel structure. Furthermore, although the number of system triggers is large, typical use cases only use a limited subset of system triggers and as a result only a limited number of channels is required. A logical 1 on a selected trigger line indicates an activated trigger and results in a channel data transfer.
- **Pending triggers** keeps track of activated triggers by locally storing them in pending bits. This is essential because multiple channel trigger may be activated simultaneously, whereas only one channel can be served by the data transfer engine at a time. This component enables the use of both level sensitive and pulse sensitive triggers.
 - Level-sensitive triggers are associated to a certain state, such as a FIFO being full. These triggers remain active as long as the state is maintained. For these triggers, keeping track of pending triggers in the P-DMA controller is not absolutely required, as the triggers are maintained outside of the controller.

- Pulse-sensitive triggers are associated to a certain event, such as an ADC sample becoming available. For these triggers, it is essential to keep track of them in the P-DMA controller as the trigger pulse may disappear before it is served by the data transfer engine.
- **Priority decoder** determines the highest priority channel with an active trigger. Within a priority group, triggers are decoded on a round-robin basis.
- **Data transfer engine** is responsible for the data transfer from a source location to a destination location. When idle, the data transfer engine is ready to accept the highest priority activated channel. It is also responsible for reading the channel descriptor from memory.
- **Master I/F** is an AHB bus master, which allows the controller to initiate AHB data transfers to the source and destination locations as well as to read the descriptor from memory.
- **Registers** - A description of the registers is found in the memory map. Each channel has a


DWx_CH_STRUCTy_CH_CURR_PTR that points to a descriptor structure in memory that specifies the data transfer.

- **Slave I/F** is an AHB bus slave, which allows the main CPU to access controller control/status registers.
- **Interrupt logic** includes interrupt status for each of the channels.

A note on output triggers. Each channel has an output trigger, tr_out. The trigger is generated as defined by TR_OUT_TYPE in the descriptor. At the system level, these output triggers can be connected to the trigger multiplexer component. This connection allows a P-DMA controller output trigger to be connected to a P-DMA controller input trigger. In other words, the completion of a specific transfer of one channel can activate another channel.

As described, each design component performs a specific function, which is best illustrated by a specific example. The following figure shows the same controller design, with a trigger/data/interrupt flow superimposed on it.

Figure 7-6. P-DMA Controller Flow



The flow exemplifies the steps that are involved in a P-DMA controller data transfer:

1. The main CPU programs the descriptor chain (in memory) and associates it with a specific channel. Further it programs the channel registers to set the desired attributes for the channel. It also programs the register that selects a specific system trigger for the channel.
2. The channel's system trigger is activated.
3. Priority decoding determines the highest priority pending channel.
4. The data transfer engine accepts the activated channel, and uses the channel identifier to load the channel's descriptor structure from memory using the master I/F.

The descriptor structure specifies the channel's data transfers.

5. The data transfer engine uses the master I/F to load data from the source location.
6. The data transfer engine uses the master I/F to store data to the destination location. The amount of data transferred depends on the TR_IN_TYPE.
7. The data transfer engine updates the active descriptor registers and the channel structure registers to reflect the data transfers. There is no update to the descriptors in memory.
8. Output trigger generation and interrupt generation are determined by TR_OUT_TYPE and INTR_TYPE respectively. P-DMA can generate an error interrupt if it encounters an error. A channel gets disabled on error.

A note on throughput. The P-DMA controller data transfer steps can be classified as either: initialization, concurrent, or sequential step:

Initialization. This includes step 1, which programs the descriptor structures. This step is done for each descriptor structure. It is performed by the main CPU, and is not initiated by an activated channel trigger.

Concurrent. This includes steps 2 and 3. These steps are performed in parallel for each channel.

Sequential. This includes steps 4 through 8. These steps are performed sequentially for each activated channel. As a result, the P-DMA controller throughput is determined by the time it takes to perform these steps. This time consists of two parts: the time spent by the controller and the time spent on the bus infrastructure. The latter time is dependent on the latency of the bus (determined by arbiter and bridge components) and the target memories/peripherals. If no wait states are incurred when accessing the target memories/peripherals, the sequential steps take 12 cycles (excluding trigger synchronization and activation covered in steps 2 and 3). In other words, the P-DMA controller can sustain 100 MHz/12 cycles = 8.33 M data transfers per second.

7.1.6.1 P-DMA channel configuration SRAMs

The P-DMA controller uses SRAM memory to store some fields of the channel configuration. The following fields of the channel configuration are part of the SRAM memory.

- DWx_CH_STRUCTy_CH_CTL.P,
- DWx_CH_STRUCTy_CH_CTL.NS,
- DWx_CH_STRUCTy_CH_CTL.B,
- DWx_CH_STRUCTy_CH_CTL.PC,
- DWx_CH_STRUCTy_CH_CTL.PREEMPTABLE
- DWx_CH_STRUCTy_CH_IDX.X_IDX,
- DWx_CH_STRUCTy_CH_IDX.Y_IDX
- DWx_CH_STRUCTy_CH_CURR_PTR.ADDR

7.1.6.2 ECC for P-DMA Channel Configuration SRAMs

The P-DMA SRAM memory uses 7-bit SECDED parity for each 32 bits of data. Address coverage is not included. ECC functionality can be enabled or disabled through the DWx_CTL0.ECC_EN register field.

Both the correctable and non-correctable ECC errors are reported to the central fault structure.

Note that P-DMA channel configuration SRAM should be initialized before reading from the SRAM to avoid unwanted ECC faults.

ECC Error Injection

The P-DMA SRAM ECC supports error injection through the following registers:

- DWx_CTL0.ECC_INJ_EN
- DWx_ECC_CTL0
- DWx_CH_STRUCTy_SRAM_DATA0
- DWx_CH_STRUCTy_SRAM_DATA1

DWx_CH_STRUCTy_SRAM_DATA0 and DWx_CH_STRUCTy_SRAM_DATA1 are provided for ECC fault injection functionality. These registers should not be used to control regular functionality (except that they can be used for initialization of P-DMA SRAMs).

For ECC fault injection, update a complete 32-bit SRAM data word with a user-provided ECC parity (specified by DWx_ECC_CTL0.PARITY) at a specific SRAM location (specified by DWx_ECC_CTL0.WORD_ADDR).

ECC Parity Generation by Software

To inject the ECC error for fault generation, ECC parity must be generated by software.

Follow this procedure to generate a 7-bit ECC parity for P-DMA SRAM. Parity generation calculates a 7-bit Parity[6:0] over a 32-bit data word W[31:0]. First, a 64-bit ECC code word CW_SW[63:0] is created:

```
CW_SW[63:0] = 64{1'b0};
CW_SW[31:0] = W[31:0];
```

Then, the 7-bit parity is calculated as the reduction XOR of the 64-bit code word CW_SW [63:0] ANDed with the following parity bit specific constants:

```

ECC_P0_SW = 64b00000011_01111111_00110110_11011011_00100010_01010100_00101010_10101011;
ECC_P1_SW = 64b00000101_10111101_11011011_01011010_01000100_10011001_01001101_00110101;
ECC_P2_SW = 64b00001001_11011101_11011100_11101110_00001000_11100010_01110001_11000110;
ECC_P3_SW = 64b00010001_11101110_10111011_10101001_10001111_00000011_10000001_11111000;
ECC_P4_SW = 64b00100001_11110110_11010111_01110101_11110000_00000011_11111110_00000000;
ECC_P5_SW = 64b01000001_11111011_01101101_10110100_11111111_11111100_00000000_00000000;
ECC_P6_SW = 64b10000001_00000011_11111111_11111000_00010001_00101100_10010110_01011111;
  
```

The parity bits are calculated as follows:

```

parity[0] = ^ (CW_SW[63:0] & ECC_P0_SW)
parity[1] = ^ (CW_SW[63:0] & ECC_P1_SW)
...
parity[6] = ^ (CW_SW[63:0] & ECC_P6_SW)
  
```


7.1.7 Functionality

This section illustrates the descriptor features and P-DMA functionality through two examples.

Example 1 - Single transfer. This example illustrates how a trigger initiates a transfer of a 16-bit sample from an ADC source to a memory destination. The ADC source has a bus interface that only supports 32-bit transfers. The memory has a bus interface that supports 8-bit, 16-bit, and 32-bit transfers. The transferred sample should be written to 16 memory bits.

The ADC sample location is at address 0x4000:0000. The memory location is at address 0x0000:0000.

Figure 7-7. Single Transfer



Setup. Let us assume that P-DMA0/DW0 channel 3 is used for this data transfer and the trigger from ADC is connected to this channel.

Initialize the channel registers.

1. DW0_CH_STRUCT3_CH_IDX.X_IDX = 0 and DW0_CH_STRUCT3_CH_IDX.Y_IDX = 0.
2. DW0_CH_STRUCT3_CH_CURR_PTR = address of the descriptor in memory.
3. Set the descriptor as follows:
 - DESC_SRC = 0x4000:0000
 - DESC_DST = 0x0000:0000
 - DESC_CTL.DESCR_TYPE = 0 (single transfer)
 - DESC_CTL.WAIT_FOR_DEACT = 0
 - DESC_CTL.INTR_TYPE = 2 (interrupt is generated after the execution of the current descriptor). The CPU is interrupted after the execution of the current descriptor. Because this is a single transfer, INTR_TYPE can be 0 or 1 and will have the same effect.
 - DESC_CTL.TR_IN_TYPE = 0 (trigger results in the execution of a single transfer). Setting it to 1 or 2 will have the same effect as the descriptor type is single transfer.
 - DESC_CTL.DATA_SIZE = 1 (16 bits).
 - DESC_CTL.SRC_TRANSFER_SIZE = 1 (32 bits)

- DESC_CTL.DST_TRANSFER_SIZE = 0 (DATA_SIZE = 16 bits)
- DESC_X_CTL.SRC_X_INCR = 0 (FIFO)
- DESC_NEXT_PTR = NULL.
- 4. DW0_CH_STRUCT3_CH_CTL.ENABLED = 1.


Transfer. When the trigger is received, the transfer engine will load 32 bits from the ADC location and will store the lower 16 bits to the 0x0000:0000 memory location. Successive triggers will have no impact on the transfer because the link pointer is set to NULL. If the link pointer points to itself, then successive triggers will result in the same behavior as the original single transfer.

Example 2 - 2D transfer. In this example, the data transfer is from a peripheral that gathers input from two channels. The transfer is to a buffer in memory with the following constraints. Figure 7-8 is a pictorial representation of the transfer.

1. The two channel data elements are interleaved in the peripheral FIFO.
2. Each data element is 2 bytes X0 and X1. The first data element of channel 1 is (X1, X0). The first data element of channel 2 is (X3, X2).
3. Destination is a double buffer: the CPU processes one buffer while P-DMA fills in the other buffer.

4. Destination only considers channel 1. So the buffer contains (X1, X0), (X5, X4), (X9, X8), ...
5. For each trigger from the peripheral, we want to transfer eight data elements or 16 B (X1, X0), (X5, X4), (X9, X8), ... , (X29, X28) to the destination.
6. Each buffer is 128 data elements or 256 B. When a buffer is full, the transfer switches to the other buffer.
7. When a buffer is full, an interrupt is generated.
8. The data transfer must continue indefinitely until the CPU disables the channel.

Figure 7-8. 2D Transfer



```
// DST_ADDR is a pointer to the buffer in memory of type uint16_t
// SRC_ADDR is a pointer to the FIFO of type uint32_t
for (Y_IDX = 0; Y_IDX ≤ 15; Y_IDX++) {
    for (X_IDX = 0; X_IDX ≤ 7; X_IDX++) {
        DST_ADDR[X_IDX * 1 + Y_IDX * 8] = (uint16_t) SRC_ADDR[0];
    }
}
```

Double buffering requires chaining of two descriptors. The basic data transfer is achieved by a 2D loop as shown in above pseudo code.

Set the first descriptor as follows. The two channel data elements are interleaved in the peripheral FIFO.

1. DESC_SRC = address of the peripheral FIFO
2. DESC_DST = address of the first buffer in memory
3. DESC_CTL.DESCR_TYPE = 2 (2D transfer)
4. DESC_CTL.INTR_TYPE = 2 (interrupt is generated after the execution of the current descriptor). The CPU is interrupted when one of the buffers is completely filled.
5. DESC_CTL.TR_IN_TYPE = 1 (trigger results in the execution of a single 1D transfer). Here, eight elements are transferred.
6. DESC_CTL.DATA_SIZE = 1 (16 bits)
7. DESC_CTL.SRC_TRANSFER_SIZE = 1 (32 bits)
8. DESC_CTL.DST_TRANSFER_SIZE = 0 (DATA_SIZE = 16 bits)
Note: The SRC_TRANSFER_SIZE of 32 bits and the DATA_SIZE of 16 bits effectively suppress the transfer of the second channel data elements.
9. DESC_X_CTL.SRC_X_INCR = 0 (FIFO)
10. DESC_X_CTL.DST_X_INCR = 1
11. DESC_X_CTL.X_COUNT = 7: 8 data elements
12. DESC_Y_CTL.SRC_X_INCR = 0 (FIFO)
13. DESC_Y_CTL.DST_Y_INCR = 8
14. DESC_X_CTL.Y_COUNT = 15. The buffer size is $(15+1) \times (7+1) = 128$ data elements.
15. DESC_NEXT_PTR = address of the second descriptor in memory

Set the second descriptor same as the first except:

1. DESC_DST = address of the second buffer in memory
2. DESC_NEXT_PTR = address of the first descriptor in memory

This setting results in the required data transfer. Only channel 1 data elements are transferred according to the requirement. An interrupt is generated when a buffer is full. The destination buffers are alternated because of the chaining.

7.1.8 P-DMA Descriptor Structure

The P-DMA descriptor is stored in memory and it consists of six fields as follows:

Table 7-6. P-DMA Descriptor Structure

Offset	Name	Description
0x00	DESCR_CTL	Descriptor control
0x04	DESCR_SRC	Descriptor source
0x08	DESCR_DST	Descriptor destination
0x0c	DESCR_X_CTL	Descriptor X loop control
0x10	DESCR_Y_CTL	Descriptor Y loop control
0x14	DESCR_NEXT_PTR	Descriptor next pointer

The offset is based on the descriptor pointer position for each channel, which is stored in the register (DWx_CH_STRUCTy_CH_CURR_PTR).

The structure and explanation of each field are as follows:

DESCR_CTL - Descriptor control

Table 7-7. P-DMA Descriptor Control

Bit	Name	Description
1:0	WAIT_FOR_DEACT	Specifies whether the controller should wait for the input trigger to be deactivated; that is, the selected system trigger is not active. This field is used to synchronize the controller with the agent that generated the trigger. This field is used only on completion of the transfer as specified by TR_IN. For example, a TX FIFO indicates that it is empty and needs a new data sample. The agent removes the trigger only when the data sample has been written by the controller and received by the agent. Furthermore, the agent's trigger may be delayed by a few cycles before it reaches the controller. This field is used for a level-sensitive trigger, which reflects the state (pulse sensitive triggers should have this field set to '0'). The wait cycles incurred by this field reduce P-DMA controller performance. 0: Do not wait for trigger deactivation (for pulse sensitive triggers). 1: Wait for up to 4 cycles. 2: Wait for up to 16 cycles. 3: Wait indefinitely. This option may result in controller lockup if the trigger is not deactivated.
3:2	INTR_TYPE	Specifies when a completion interrupt is generated (CH_STATUS.INTR_CAUSE is set to COMPLETION): 0: An interrupt is generated after a single transfer. 1: An interrupt is generated after a single 1D transfer. ☐ If the descriptor type is 'single', the interrupt is generated after a single transfer. ☐ If the descriptor type is 1D, CRC, or 2D, the interrupt is generated after the execution of a 1D transfer. 2: An interrupt is generated after the execution of the current descriptor. Independent of the value of DESCR_NEXT_PTR.ADDR of the current descriptor. 3: An interrupt is generated after the execution of the current descriptor. The value of DESCR_NEXT_PTR.ADDR of the current descriptor must be 0.
5:4	TR_OUT_TYPE	Specifies when an output trigger is generated: 0: An output trigger is generated after a single transfer. 1: An output trigger is generated after a single 1D transfer. ☐ If the descriptor type is 'single', the output trigger is generated after a single transfer. ☐ If the descriptor type is 1D, CRC, or 2D, the output trigger is generated after the execution of a 1D transfer. 2: An output trigger is generated after the execution of the current descriptor. 3: An output trigger is generated after the execution of a descriptor list: after the execution of the current descriptor and the current descriptor DESCR_NEXT_PTR.ADDR is 0.
7:6	TR_IN_TYPE	Specifies the input trigger type (not to be confused with the descriptor type): 0: A trigger results in the execution of a single transfer. The descriptor type can be single, 1D, or 2D. 1: A trigger results in the execution of a single 1D transfer. ☐ If the descriptor type is 'single', the trigger results in the execution of a single transfer. ☐ If the descriptor type is 1D or 2D, the trigger results in the execution of a 1D transfer. 2: A trigger results in the execution of the current descriptor. 3: A trigger results in the execution of the current descriptor and continues (without requiring another input trigger) with the execution of the next descriptor using the next descriptor's information.
24	CH_DISABLE	Specifies whether the channel is disabled after completion of the current descriptor (independent of the value of the DESCR_NEXT_PTR value): 0: Channel is not disabled. 1: Channel is disabled. Note: A disabled channel will ignore its input trigger.

Table 7-7. P-DMA Descriptor Control

Bit	Name	Description
26	SRC_TRANSFER_SIZE	Specifies the bus transfer size to the source location: 0: As specified by DATA_SIZE. 1: Word (32 bits). Distinguishing bus transfer size from data element size allows for source components with data elements that are smaller than their 32-bit bus interface width. For example, an ADC source has a 32-bit bus transfer size, but only provides a 16-bit data element.
27	DST_TRANSFER_SIZE	Specifies the bus transfer size to the destination location: 0: As specified by DATA_SIZE. 1: Word (32 bits). Distinguishing bus transfer size from data element size allows for destination components with data elements that are smaller than their 32-bit bus interface width. For example, a DAC destination has a 32-bit bus transfer size, but only requires a 16-bit data element.
29:28	DATA_SIZE	Specifies the data element size: 0: Byte (8 bits). 1: Halfword (16 bits). 2: Word (32 bits). DATA_SIZE, SRC_TRANSFER_SIZE, and DST_TRANSFER_SIZE together determine how data elements are transferred. The following are the nine legal settings: ☐ DATA is 8 bit, SRC is 8 bit, DST is 8 bit. ☐ DATA is 8 bit, SRC is 32 bit (higher 24 bits are dropped), DST is 8 bit. ☐ DATA is 8 bit, SRC is 8 bit, DST is 32 bit (higher 24 bits are made "0"). ☐ DATA is 8 bit, SRC is 32 bit (higher 24 bits are dropped), DST is 32 bit (higher 24 bits are made "0"). ☐ DATA is 16 bit, SRC is 16 bit, DST is 16 bit. ☐ DATA is 16 bit, SRC is 32 bit (higher 16 bits are dropped), DST is 16 bit. ☐ DATA is 16 bit, SRC is 16 bit, DST is 32 bit (higher 16 bits are made "0"). ☐ DATA is 16 bit, SRC is 32 bit (higher 16 bits are dropped), DST is 32 bit (higher 16 bits are made "0"). ☐ DATA is 32 bit, SRC is 32 bit, DST is 32 bit.
31:30	DESCR_TYPE	Specifies the descriptor type (not to be confused with the trigger type): 0: Single transfer. The DESCR_X_CTL and DESCR_Y_CTL registers are not present and DESCR_NEXT_PTR is at offset 0x0c. 1: 1D transfer. The DESCR_X_CTL register is present, the DESCR_Y_CTL is not present, and DESCR_NEXT_PTR is at offset 0x10. A 1D transfer consists of DESCR_X_CTL.X_COUNT single transfers. 2: 2D transfer. The DESCR_X_CTL and DESCR_Y_CTL registers are present and DESCR_NEXT_PTR is at offset 0x14. A 2D transfer consists of DESCR_X_CTL.X_COUNT*DESCR_Y_CTL.Y_COUNT single transfers. 3: CRC transfer. The DESCR_X_CTL register is present, the DESCR_Y_CTL is not present and DESCR_NEXT_PTR is at offset 0x10. A CRC transfer consists of DESCR_X_CTL.X_COUNT single transfers. After the execution of the current descriptor, the DESCR_NEXT_PTR address is copied to the channel's DWx_CH_STRUCTy_CH_CURR_PTR address and DWx_CH_STRUCTy_CH_IDX.X_IDX and DWx_CH_STRUCTy_CH_IDX.Y_IDX are set to 0.

DESCR_SRC - Descriptor source

Table 7-8. P-DMA Descriptor Source

Bit	Name	Description
31:0	SRC_ADDR	Base address of source location.

DESCR_DST - Descriptor destination

Table 7-9. P-DMA Descriptor Destination

Bit	Name	Description
31:0	DST_ADDR	Base address of destination location. Note: For a CRC transfer descriptor, the calculated CRC value is stored at this address location. It is not subjected to post processing specified by the DWx_CRC_CTL0/DWx_CRC_REM_CTL0 registers. The CRC result after post processing is only available in the DWx_CRC_REM_RESULT0 register.

DESCR_X_CTL - Descriptor X loop control

Table 7-10. P-DMA Descriptor X Loop Control

Bit	Name	Description
11:0	SRC_X_INCR	Specifies increment of source address for each X loop iteration (in multiples of SRC_TRANSFER_SIZE). This field is a signed number in the range [−2048, 2047]. If this field is 0, the source address is not incremented. This is useful for reading from RX FIFO structures.
23:12	DST_X_INCR	Specifies increment of destination address for each X loop iteration (in multiples of DST_TRANSFER_SIZE). This field is a signed number in the range [−2048, 2047]. If this field is 0, the destination address is not incremented. This is useful for writing to TX FIFO structures. Note: This field is not used for CRC transfer descriptors and must be set to '0'.
31:24	X_COUNT	Number of iterations (minus 1) of the X loop (X_COUNT+1 is the number of single transfers in a 1D transfer). This field is an unsigned number in the range [0, 255], representing 1 through 256 iterations.

DESCR_Y_CTL - Descriptor Y loop control

This register is not present for the single, 1D, and CRC transfer descriptor types.

Table 7-11. P-DMA Descriptor Y Loop Control

Bit	Name	Description
11:0	SRC_Y_INCR	Specifies increment of source address for each Y loop iteration (in multiples of SRC_TRANSFER_SIZE). This field is a signed number in the range [−2048, 2047].
23:12	DST_Y_INCR	Specifies increment of destination address for each Y loop iteration (in multiples of DST_TRANSFER_SIZE). This field is a signed number in the range [−2048, 2047].
31:24	Y_COUNT	Number of iterations (minus 1) of the Y loop (Y_COUNT+1) × (X_COUNT+1) is the number of single transfers in a 2D transfer). This field is an unsigned number in the range [0, 255], representing 1 through 256 iterations.

DESCR_NEXT_PTR - Descriptor next pointer

Note: For a single transfer descriptor type, this register is at offset 0x0c. For a 1D transfer descriptor type, this register is at offset 0x.

Table 7-12. P-DMA Descriptor Next Pointer

Bit	Name	Description
31:2	ADDR	Address of next descriptor in the descriptor list. When this field is 0, this is the last descriptor in the descriptor list.

7.2 Memory DMA (M-DMA)

The M-DMA controller is used to transfer data between memory and peripherals without CPU involvement:

- The CPU configures/programs the M-DMA controller.
- The M-DMA controller performs the data transfers.

The primary design target is to achieve high memory bandwidth with limited area overhead to the platform.

The main difference between the M-DMA and P-DMA controllers is that the M-DMA controller has dedicated channel logic (with channel state) for each channel, whereas the P-DMA reuses the channel logic for all channels. Furthermore, the M-DMA channel logic includes a 16-byte FIFO for temporary storage of data. This results in increased memory bandwidth, but comes at the cost of significant silicon area overhead for each channel. M-DMA supports an additional descriptor type called “Memory Copy”. This is a special 1D transfer; `DESCR_X_INCR.SRC_X_INCR` and `DESCR_X_INCR.DST_X_INCR` are implicitly set to '1' and not part of the descriptor. This descriptor makes it possible to achieve higher bandwidth for certain class of transfers.

7.2.1 Overview

The M-DMA controller can be configured/programmed to perform multiple independent data transfers. Each data transfer is managed by a channel. The number of channels varies for different part numbers; more details are available in the device datasheet.

A channel has an associated priority. When there are multiple bus transfer requests, the priority decoder determines the highest priority channel for the request.

A data transfer is initiated by an input trigger. This trigger may originate from the source of the transfer, destination of the transfer, CPU software, or from another SoC component. Triggers provide Active/Sleep functionality and are not available in DeepSleep and Hibernate power modes.

The data transfer specifics are specified by a descriptor. This descriptor specifies (among other things):

- The source and destination address locations and the size of the transfer.
- The actions of a channel; for example, generation of output triggers and interrupts.
- Data transfer types can be single, 1D, or 2D as defined in the descriptor structure. These types essentially define the address sequences generated for source and destination. 1D and 2D transfers are used for “scatter gather” and other useful transfer operations.

A channel's descriptor state is encoded as part of the channel's register state (and not as part of the descriptor).

7.2.2 Channels

M-DMA supports multiple independent data transfers that are managed by different channels. Each channel connects to a specific system trigger through a trigger multiplexer that is outside the M-DMA controller. See the [Trigger Multiplexer chapter on page 493](#) for details.

Channel priority. A channel is assigned a priority (`DMAC_CHx_CTL.PRIO`) between 0 and 3, with 0 being the highest priority and 3 being the lowest priority. Priority decoding determines the highest priority pending channel. Channels with the same priority constitute a priority group and within this priority group, the following round-robin arbitration is applied.

A “round” consists of a contiguous sequence of channel activations, within this priority group, without any repetition. Within a round, higher priority is given to the lower channel indices. The notion of a round guarantees that within a group, higher channel indices do not yield to lower indices indefinitely.

The data transfer associated with a trigger is made up of one or more atomic transfers or single transfers. For example, a 1D transfer consists of `X_COUNT+1` single transfers.

Channel registers. A channel has three access control attributes that are used by the SMPUs and PPU's for access control:

- Privileged Mode (`DMAC_CHx_CTL.P`) attribute can be set to privileged or user.
- Non-secure (`DMAC_CHx_CTL.NS`) attribute can be set to secure or non-secure.
- PC (`DMAC_CHx_CTL.PC`) can be set to one of the protection contexts.

These three fields are inherited from the write transaction and not specified by the transaction write data.

Channel registers. The following registers provide a channel's descriptor state:

- `DMAC_CHx_CTL`. This register provides generic channel control information.
- `DMAC_CHx_CURR`. This register provides the address of the memory location where the current descriptor is located. Software needs to initialize this register. Hardware sets this register to the current descriptor's next descriptor pointer, when advancing from the current descriptor to the next descriptor in a descriptor list. When this field is 0, there is no valid descriptor.
- `DMAC_CHx_IDX`. This register provides the current X and Y indices of the channel into the current descriptor. Software needs to initialize this register. Hardware sets the X and Y indices to 0, when advancing from the current descriptor to the next descriptor in a descriptor list.

- **DMAC_CHx_SRC.** This register provides the current address of source location.
- **DMAC_CHx_DST.** This register provides the current address of destination location.
- **DMAC_CHx_DESCR_STATUS.** This register provides the validity of other DMAC_CHx_DESCR registers.
- **DMAC_CHx_DESCR_CTL.** This register contains a copy of DESCR_CTL of the currently active descriptor.
- **DMAC_CHx_DESCR_SRC.** This register contains a copy of DESCR_SRC of the currently active descriptor.
- **DMAC_CHx_DESCR_DST.** This register contains a copy of DESCR_DST of the currently active descriptor.
- **DMAC_CHx_DESCR_X_INCR.** This register contains a copy of DESCR_X_INCR of the currently active descriptor.
- **DMAC_CHx_DESCR_Y_SIZE.** This register contains a copy of DESCR_X_SIZE of the currently active descriptor.
- **DMAC_CHx_DESCR_Y_INCR.** This register contains a copy of DESCR_Y_INCR of the currently active descriptor.
- **DMAC_CHx_DESCR_Y_SIZE.** This register contains a copy of DESCR_Y_SIZE of the currently active descriptor.
- **DMAC_CHx_DESCR_NEXT.** This register contains a copy of DESCR_NEXT_PTR of the currently active descriptor.
- **DMAC_CHx_INTR.** This register contains the interrupts that are currently activated for this channel.
- **DMAC_CHx_INTR_SET.** Writing '1' to the appropriate bit in this register sets the corresponding DMAC_CHx_INTR field to 1.
- **DMAC_CHx_INTR_MASK.** Mask for corresponding field in DMAC_CHx_INTR register.
- **DMAC_CHx_INTR_MASKED.** Logical and of corresponding DMAC_CHx_INTR and DMAC_CHx_INTR_MASK fields.
- **DMAC_CHx_TR_CMD.** This register allows the channel to be triggered through software. This is in addition to the software trigger control available in the trigger multiplexer.

Note that channel state is retained in DeepSleep power mode.

The M-DMA controller is an Active/Sleep power mode functionality. Software should not initiate DeepSleep system power mode entry if there are any active M-DMA controller channels transferring data. Note that there is no way of capturing the active channel data while transitioning to DeepSleep system power mode.

7.2.3 Descriptors

A descriptor is stored in memory and describes a data transfer. The descriptor is read-only for the M-DMA controller.

Descriptor Type (DESCR_TYPE) - There are five types of descriptors.

Table 7-13. M-DMA Descriptor Types

Descriptor Type	Description
Single transfer	This transfers a single data element (8-bit, 16-bit, or 32-bit) as shown in Figure 7-9 . The descriptor size is four 32-bit words: DESCR_CTL, DESCR_SRC, DESCR_DST, and DESCR_NEXT_PTR.
1D transfer	This performs a one-dimensional "for loop" (described in C) as shown in Figure 7-10 . A 1D transfer is made up of X_COUNT+1 single transfers. The descriptor size is six 32-bit words: DESCR_CTL, DESCR_SRC, DESCR_DST, DESCR_X_INCR, DESCR_X_SIZE, and DESCR_NEXT_PTR.
2D transfer	This performs a two-dimensional "for loop" (described in C) as shown in Figure 7-11 . A 2D transfer is made up of (Y_COUNT+1) 1D transfers. The descriptor size is eight 32-bit words: DESCR_CTL, DESCR_SRC, DESCR_DST, DESCR_X_INCR, DESCR_X_SIZE, DESCR_Y_INCR, DESCR_Y_SIZE, and DESCR_NEXT_PTR.
Memory Copy	This is a special case of 1D transfer as shown in Figure 7-12 ; DESCR_X_INCR.SRC_X_INCR and DESCR_X_INCR.DST_X_INCR are implicitly set to 1 and not part of the descriptor. The size of the descriptor is five 32-bit words. The M-DMA is optimized for performance.
Scatter	This descriptor type is intended to write a set of 32-bit data elements as shown in Figure 7-13 , whose addresses are "scattered" around the address space. The size of the descriptor is four 32-bit words. DESCR_CTL, DESCR_SRC, DESCR_X_SIZE, and DESCR_NEXT_PTR.

The functionality of these five descriptor types is described by the following pseudo code.

Figure 7-9. Single Transfer

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
DST_ADDR[0] = (t_DATA_SIZE) SRC_ADDR[0];
```

Figure 7-10. 1D Transfer

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (X_IDX = 0; X_IDX ≤ X_COUNT; X_IDX++) {
    DST_ADDR[X_IDX * DST_X_INCR] =
        (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR];
}
```

Figure 7-11. 2D Transfer

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (Y_IDX = 0; Y_IDX ≤ Y_COUNT; Y_IDX++) {
    for (X_IDX = 0; X_IDX ≤ X_COUNT; X_IDX++) {
        DST_ADDR[X_IDX * DST_X_INCR + Y_IDX * DST_Y_INCR ] =
            (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR + Y_IDX * SRC_Y_INCR];
    }
}
```

Figure 7-12. Memory Copy

```
// DST_ADDR is a pointer to an object of type uint8_t
// SRC_ADDR is a pointer to an object of type uint8_t
// This transfer type uses 8-bit, 16-bit an 32-bit transfers. The hardware ensures that
// alignment requirements are met.
for (X_IDX = 0; X_IDX ≤ X_COUNT; X_IDX++) {
    DST_ADDR[X_IDX] = SRC_ADDR[X_IDX];
}
```

Figure 7-13. Scatter

```
// SRC_ADDR is a pointer to an object of type uint32_t
for (X_IDX = 0; X_IDX ≤ X_COUNT; X_IDX += 2) {
    address = SRC_ADDR[X_IDX];
    data    = SRC_ADDR[X_IDX + 1];
    *address = data;
}
```

Descriptor size – The size of a descriptor depends on its type. Only relevant parameters are stored. For example, a 1D descriptor does not contain the Y_SIZE and Y_INCR parameters.

Transfer size (SRC_TRANSFER_SIZE and DST_TRANSFER_SIZE) – In a data transfer, the source data is cast into the type specified by DATA_SIZE and assigned to the destination. The source type is determined by SRC_TRANSFER_SIZE and the destination type is determined by DST_TRANSFER_SIZE. All types are unsigned. All address computations use C semantics based on the transfer size.

Descriptor chaining – Descriptors chained together. DESCR_NEXT_PTR field contains a pointer to the next descriptor in the chain. A channel executes the next descriptor in the chain when it completes executing the current descriptor. The last descriptor in the chain has DESCR_NEXT_PTR set to '0' (NULL pointer). A descriptor chain is also referred to as a descriptor list. It is possible to have a circular list in which case the execution continues indefinitely until there is an error or the channel or the controller is disabled by software.

Trigger-in type (TR_IN_TYPE) – An input trigger initiates a data transfer and the TR_IN_TYPE defines the action on a trigger.

Table 7-14. M-DMA Trigger-in Types

Trigger Type	Description
Type 0	Trigger results in the execution of a single transfer. In a 1D or 2D transfer, this will execute a single transfer in the loop.
Type 1	Trigger results in the execution of a single 1D transfer. If the descriptor type is single transfer this behaves similar to type 0. If the descriptor type is 2D, it results in executing the inner loop once.
Type 2	Trigger results in the execution of the current descriptor.
Type 3	Trigger results in the execution of a descriptor list.

Trigger-out Type (TR_OUT_TYPE) – This defines when an output trigger is generated.

Table 7-15. M-DMA Trigger-out Types

Trigger Type	Description
Type 0	Output trigger is generated after a single transfer. In a 1D or 2D transfer, an output trigger is generated after each transfer in the loop.
Type 1	Output trigger is generated after a single 1D transfer. In a single transfer descriptor type, this behaves similar to type 0. If the descriptor type is 2D, an output trigger is generated after each execution of the inner loop.
Type 2	Output trigger is generated after the execution of the current descriptor.
Type 3	Output trigger is generated after the execution of a descriptor list.

Interrupt Type (INTR_TYPE) – This defines when a completion interrupt is generated.

Table 7-16. M-DMA Interrupt Types

Trigger Type	Description
Type 0	Interrupt is generated after a single transfer. In a 1D or 2D transfer, an interrupt is generated after each transfer in the loop
Type 1	Interrupt is generated after a single 1D transfer. In a single transfer descriptor type, this behaves similar to type 0. If the descriptor type is 2D, an interrupt is generated after each execution of the inner loop.
Type 2	Interrupt is generated after the execution of the current descriptor.
Type 3	Interrupt is generated after the execution of a descriptor list.

Wait for Deactivation (WAIT_FOR_DEACT) – Specifies whether the M-DMA controller should wait for the input trigger to be deactivated after it has completed the data transfer corresponding to the current trigger. This field is used for level-sensitive triggers to give sufficient time for the triggering agent to deactivate the trigger. The wait specified can be 0, up to four cycles, up to 16 cycles, or indefinite. Pulse-sensitive triggers should have this field set to 0.

Data Prefetch – If this bit is set, source data transfers are initiated as soon as the channel is enabled, the current descriptor pointer is not 0, and there is space available in the channel's data FIFO. When the input trigger is activated, the trigger can initiate destination data transfers with data that is already in the channel's data FIFO. This effectively shortens the initial delay of the data transfer. Data prefetch should be used with care, to ensure that data synchronization is not violated.

7.2.4 Interrupts

M-DMA can generate interrupts on completion and on error conditions:

- The INTR_TYPE descriptor control defines when a completion condition (COMPLETION) is activated.
- The error conditions include SRC_BUS_ERROR, DST_BUS_ERROR, SRC_MISAL, DST_MISAL, CURR_PTR_NULL, ACTIVE_CH_DISABLED, and DESCR_BUS_ERROR.

DMAC_CHx_INTR – Each channel has an interrupt request register. There are eight possible causes that can generate an interrupt. These causes are encoded in bits 0 to 7. Software can clear these by writing to these bits.

DMAC_CHx_INTR_SET – Each channel has an interrupt set register. There are eight bits (same as DMAC_CHx_INTR) and software can write 1 to any of these bits to set the corresponding DMAC_CHx_INTR bit.

DMAC_CHx_INTR_MASK – Each channel has an interrupt mask register. There are eight bits (same as DMAC_CHx_INTR) and they can be selectively enabled by writing 1 to the corresponding bits.

DMAC_CHx_INTR_MASKED – Each channel has an interrupt masked register. When read, this register reflects a

bitwise “and” between the interrupt request and mask registers.

The M-DMA is an Active power mode peripheral; this means, it uses Active functionality interrupts. Therefore, DMAC_CHx_INTR and DMAC_CHx_INTR_SET are not retained in DeepSleep power mode (DMAC_CHx_INTR_MASK is retained).

7.2.5 Control and Active Registers


DMAC_CHx_CTL.ENABLED indicates whether the M-DMA is enabled. Software writes to this register to enable the controller.

The ACTIVE register indicates which channels are currently active – enabled channels whose trigger is activated.

7.2.6 M-DMA Controller Design

The following figure gives an overview of the M-DMA controller design.

Figure 7-14. M-DMA Controller Design



The following components are distinguished:

Channel logic. Each M-DMA controller channel has its own dedicated channel logic. This logic tracks the channel's input trigger and maintains the channel state (channel registers and a copy of the current descriptor from memory) and a data transfer engine. The data transfer engine transfers data elements from a source location to a destination location as specified by the channel state. The channels transfer requests are arbitrated by the priority decoder using channel specific priorities.

Each channel consists of two state machines that are connected through a 16-byte FIFO. The first state machine reads the descriptors from memory and data from the source location. When the current descriptor is read from memory, it is part of the channel's state. Source location data is temporarily buffered in the FIFO. The second state machine writes the buffered data in the FIFO to the source location.

Priority decoder determines the highest priority channel with a bus transfer request.

registers. A description of the registers is available in the memory map. This memory map also describes the descriptors.


Master I/F is an AHB-Lite bus master, which allows the controller to initiate AHB-Lite data transfers to the source and destination locations as well as to read the descriptor from memory.

Slave I/F is an AHB-Lite bus slave, which allows the main CPU to access M-DMA controller control/status registers.

7.2.7 Examples


Example: The source is a 32-bit word addressable peripheral; the destination is regular memory. The M-DMA controller transfers five bytes from the source to destination. The source transfer size is a 32-bit word. The data size is an 8-bit byte. The destination transfer size is an 8-bit byte. A 1D transfer descriptor type is used.

Figure 7-15. M-DMA 1D Transfer




Example: The source and memory are regular memory. The M-DMA controller transfers five byte pairs and de-interleaves the pairs as part of the transfer. The source transfer size, data size, and destination transfer size are all 8-bit bytes. A 2D transfer descriptor type is used.

Figure 7-16. M-DMA 2D Transfer



Example: The source and memory are regular memory. The M-DMA controller transfers bytes in the inverse direction of the previous example (note how the source and destination increments are reversed).

Figure 7-17. M-DMA Inverse Direction Transfer



7.2.8 M-DMA Descriptor Structure

The M-DMA descriptor is stored in memory and it consists of eight fields as follows:

Table 7-17. M-DMA Descriptor Structure

Offset	Name	Description
0x00	DESCR_CTL	Descriptor control
0x04	DESCR_SRC	Descriptor source
0x08	DESCR_DST	Descriptor destination
0x0c	DESCR_X_SIZE	Descriptor X loop size
0x10	DESCR_X_INCR	Descriptor X loop increment
0x14	DESCR_Y_SIZE	Descriptor Y loop size
0x18	DESCR_Y_INCR	Descriptor Y loop increment
0x1c	DESCR_NEXT_PTR	Descriptor next pointer

The offset is based on the descriptor pointer position for each channel which is stored in the register (DMAC_CHx_CURR).

The structure and explanation of each field are as follows:

DESCR_CTL - Descriptor control

Table 7-18. M-DMA Descriptor Control

Bit	Name	Description
1:0	WAIT_FOR_DEACT	Specifies whether the controller should wait for the input trigger to be deactivated; that is, the selected system trigger is not active. This field is used to synchronize the controller with the agent that generated the trigger. This field is used only on completion of the transfer as specified by TR_IN. For example, a TX FIFO indicates that it is empty and needs a new data sample. The agent removes the trigger only when the data sample has been written by the controller and received by the agent. Furthermore, the agent's trigger may be delayed by a few cycles before it reaches the controller. This field is used for a level-sensitive trigger, which reflects the state (pulse sensitive triggers should have this field set to '0'). The wait cycles incurred by this field reduce M-DMA controller performance. 0: Do not wait for trigger deactivation (for pulse sensitive triggers). 1: Wait for up to 4 cycles. 2: Wait for up to 16 cycles. 3: Wait indefinitely. This option may result in controller lockup if the trigger is not deactivated.
3:2	INTR_TYPE	Specifies when a completion interrupt is generated: 0: An interrupt is generated after a single transfer. 1: An interrupt is generated after a single 1D transfer or a memory copy transfer ☐ If the descriptor type is "single", the interrupt is generated after a single transfer. ☐ If the descriptor type is 1D or 2D, the interrupt is generated after the execution of a 1D transfer. ☐ If the descriptor type is "memory copy", the interrupt is generated after the execution of a memory copy transfer. ☐ If the descriptor type is "scatter", the interrupt is generated after the execution of a scatter transfer. 2: An interrupt is generated after the execution of the current descriptor. Independent of the value of DESCR_NEXT_PTR.ADDR of the current descriptor. 3: An interrupt is generated after the execution of the current descriptor and the current descriptor's DESCR_NEXT_PTR.ADDR is 0.
5:4	TR_OUT_TYPE	Specifies when an output trigger is generated: 0: An output trigger is generated after a single transfer. 1: An output trigger is generated after a single 1D transfer or a memory copy transfer. ☐ If the descriptor type is "single", the output trigger is generated after a single transfer. ☐ If the descriptor type is 1D or 2D, the output trigger is generated after the execution of a 1D transfer. ☐ If the descriptor type is "memory copy", the output trigger is generated after the execution of a memory copy transfer. ☐ If the descriptor type is "scatter", the output trigger is generated after the execution of a scatter transfer. 2: An output trigger is generated after the execution of the current descriptor. 3: An output trigger is generated after the execution of a descriptor list: after the execution of the current descriptor and the current descriptor's DESCR_NEXT_PTR.ADDR is 0'.
7:6	TR_IN_TYPE	Specifies the input trigger type (not to be confused with the descriptor type): 0: A trigger results in the execution of a single transfer. The descriptor type can be single, 1D or 2D. 1: A trigger results in the execution of a single 1D transfer. ☐ If the descriptor type is "single", the trigger results in the execution of a single transfer. ☐ If the descriptor type is 1D or 2D, the trigger results in the execution of a 1D transfer. ☐ If the descriptor type is "memory copy", the trigger results in the execution of a memory copy transfer. ☐ If the descriptor type is "scatter", the trigger results in the execution of an scatter transfer. 2: A trigger results in the execution of the current descriptor. 3: A trigger results in the execution of the current descriptor and continues (without requiring another input trigger) with the execution of the next descriptor using the next descriptor's information.

Table 7-18. M-DMA Descriptor Control

Bit	Name	Description
8	DATA_PREFETCH	Source data prefetch: 0: No source data prefetch. Source data transfers are only initiated after the input trigger is activated. 1: Source data prefetch. Source data transfers are initiated as soon as the channel is enabled, the current descriptor pointer is not 0 and there is space available in the channel's data FIFO. When the input trigger is activated, the trigger can initiate destination data transfers with data that is already in the channel's data FIFO. This effectively shortens the initial delay of the data transfer. Note: Data prefetch should be used with care, to ensure that data coherency is guaranteed and that prefetches do not cause undesired side effects.
17:16	DATA_SIZE	Specifies the data element size: 0: Byte (8 bits). 1: Halfword (16 bits). 2: Word (32 bits). DATA_SIZE, SRC_TRANSFER_SIZE, and DST_TRANSFER_SIZE together determine how data elements are transferred. The following are the nine legal settings: ☐ DATA is 8 bit, SRC is 8 bit, DST is 8 bit. ☐ DATA is 8 bit, SRC is 32 bit (higher 24 bits are dropped), DST is 8 bit. ☐ DATA is 8 bit, SRC is 8 bit, DST is 32 bit (higher 24 bits are made '0'). ☐ DATA is 8 bit, SRC is 32 bit (higher 24 bits are dropped), DST is 32 bit (higher 24 bits are made '0'). ☐ DATA is 16 bit, SRC is 16 bit, DST is 16 bit. ☐ DATA is 16 bit, SRC is 32 bit (higher 16 bits are dropped), DST is 16 bit. ☐ DATA is 16 bit, SRC is 16 bit, DST is 32 bit (higher 16 bits are made '0'). ☐ DATA is 16 bit, SRC is 32 bit (higher 16 bits are dropped), DST is 32 bit (higher 16 bits are made '0'). ☐ DATA is 32 bit, SRC is 32 bit, DST is 32 bit. Note: This field is not used for a "memory copy" descriptor type. It must be set to '2' for a "initialization" descriptor type.

Table 7-18. M-DMA Descriptor Control

Bit	Name	Description
24	CH_DISABLE	Specifies whether the channel is disabled or not after completion of the current descriptor (independent of the value of the DESC_NEXT_PTR value): 0: Channel is not disabled. 1: Channel is disabled.
26	SRC_TRANSFER_SIZE	Specifies the bus transfer size to the source location: 0: As specified by DATA_SIZE. 1: Word (32 bits). Distinguishing bus transfer size from data element size allows for source components with data elements that are smaller than their 32-bit bus interface width. For example, an ADC source has a 32-bit bus transfer size, but only provides a 16-bit data element. Note: This field is not used for a "memory copy" descriptor type. It must be set to '1' for a "scatter" descriptor type.
27	DST_TRANSFER_SIZE	Specifies the bus transfer size to the destination location: 0: As specified by DATA_SIZE. 1: Word (32 bits). Distinguishing bus transfer size from data element size allows for destination components with data elements that are smaller than their 32-bit bus interface width. For example, a DAC destination has a 32-bit bus transfer size, but only requires a 16-bit data element. Note: This field is not used for a "memory copy" descriptor type. It must be set to '1' for a "scatter" descriptor type.
30:28	DESCR_TYPE	Specifies the descriptor type (not to be confused with the trigger type): 0: Single transfer. The DESCR_X_SIZE, DESCR_X_INCR, DESCR_Y_SIZE, and DESCR_Y_INCR registers are not present. The DESC_NEXT_PTR is at offset 0x0c. 1: 1D transfer. The DESCR_X_SIZE and DESCR_X_INCR registers are present, the DESCR_Y_SIZE and DESCR_Y_INCR are not present. A 1D transfer consists out of DESCR_X_SIZE.X_COUNT+1 single transfers. The DESC_NEXT_PTR is at offset 0x14. 2: 2D transfer. The DESCR_X_SIZE, DESCR_X_INCR, DESCR_Y_SIZE, and DESCR_Y_INCR registers are present. A 2D transfer consists of (DESCR_X_SIZE.X_COUNT+1)*(DESCR_Y_SIZE.Y_COUNT+1) single transfers. The DESC_NEXT_PTR is at offset 0x1c. 3: Memory copy. The DESCR_X_SIZE register is present, the DESCR_X_INCR, DESCR_Y_SIZE, and DESCR_Y_INCR are not present. A memory copy transfer copies DESCR_X_SIZE.X_COUNT+1 Bytes and may use Byte, halfword, and word transfers. The DESC_NEXT_PTR is at offset 0x10. 4: Scatter transfer. The DESCR_X_SIZE register is present, the DESCR_DST, DESCR_X_INCR, DESCR_Y_SIZE, and DESCR_Y_INCR are not present. 5-7: Undefined. After the execution of the current descriptor, the DESC_NEXT_PTR address is copied to the channel's DMAC_CHx_CURR address and DMAC_CHx_IDX.X and DMAC_CHx_IDX.Y are set to '0'.

DESCR_SRC - Descriptor source

Table 7-19. M-DMA Descriptor Source

Bit	Name	Description
31:0	SRC_ADDR	Base address of source location.

DESCR_DST - Descriptor destination

Table 7-20. M-DMA Descriptor Destination

Bit	Name	Description
31:0	DST_ADDR	Base address of destination location.

DEDESCR_X_SIZE - Descriptor X loop size

This register is not present for the single transfer descriptor type.

Table 7-21. M-DMA Descriptor X Loop Size

Bit	Name	Description
15:0	X_COUNT	Number of iterations (minus 1) of the X loop ($X\_COUNT+1$ is the number of single transfers in a 1D transfer). This field is an unsigned number in the range [0, 65535], representing 1 through 65536 iterations. For the memory copy descriptor type, this field specifies the number of transferred Bytes (minus 1). For the scatter descriptor type, this field specifies the number of (address, write data) initialization pairs (times 2, minus 1).

DESCR_X_INCR - Descriptor X loop increment

This register is not present for the single transfer, memory copy, and scatter descriptor types.

Table 7-22. M-DMA Descriptor X Loop Increment

Bit	Name	Description
15:0	SRC_X_INCR	Specifies increment of source address for each X loop iteration (in multiples of SRC_TRANSFER_SIZE). This field is a signed number in the range [−32768, 32767]. If this field is 0, the source address is not incremented. This is useful for reading from RX FIFO structures.
31:16	DST_X_INCR	Specifies increment of destination address for each X loop iteration (in multiples of DST_TRANSFER_SIZE). This field is a signed number in the range [−32768, 32767]. If this field is 0, the destination address is not incremented. This is useful for writing to TX FIFO structures.

DEDESCR_Y_SIZE - Descriptor Y loop size

This register is not present for the single transfer, 1D transfer, memory copy, and scatter descriptor types.

Table 7-23. M-DMA Descriptor Y Loop Size

Bit	Name	Description
15:0	Y_COUNT	Number of iterations (minus 1) of the Y loop ($X\_COUNT+1 \times (Y\_COUNT+1)$ is the number of single transfers in a 2D transfer). This field is an unsigned number in the range [0, 65535], representing 1 through 65536 iterations.

DESCR_Y_INCR - Descriptor Y loop increment

This register is not present for the single transfer, 1D transfer, memory copy, and scatter descriptor types.

Table 7-24. M-DMA Descriptor Y Loop Increment

Bit	Name	Description
15:0	SRC_Y_INCR	Specifies increment of source address for each Y loop iteration (in multiples of SRC_TRANSFER_SIZE). This field is a signed number in the range [−32768, 32767].
31:16	DST_Y_INCR	Specifies increment of destination address for each Y loop iteration (in multiples of DST_TRANSFER_SIZE). This field is a signed number in the range [−32768, 32767].

DESCR_NEXT_PTR - Descriptor next pointer

Note: For a single transfer descriptor type, this register is at offset 0x0c. For a 1D transfer descriptor type, this register is at offset 0x14. For a 2D transfer descriptor type, this register is at offset 0x1c. For a memory copy transfer descriptor type, this register is at offset 0x10. For a scatter transfer descriptor type, this register is at offset 0x0c.

Table 7-25. M-DMA Descriptor Next Pointer

Bit	Name	Description
31:2	ADDR	Address of next descriptor in the descriptor list. When this field is 0, this is the last descriptor in the descriptor list.

7.3 Registers

Table 7-26. P-DMA Registers

Register	Name	Description
DWx_CTL0	Control Register	This register provides P-DMA enable/disable control and ECC checking/injection for SRAM enable/disable control
DWx_STATUS0	Status register	This register provides status of the P-DMA controller
DWx_ACT_DESCR_CTL0	Active descriptor control register	This register provides the copy of DESCR_CTL field of the currently active descriptor
DWx_ACT_DESCR_SRC0	Active descriptor source register	This register provides the copy of DESCR_SRC field of the currently active descriptor.
DWx_ACT_DESCR_DST0	Active descriptor destination register	This register provides the copy of DESCR_DST field of the currently active descriptor.
DWx_ACT_DESCR_X_CTL0	Active descriptor X loop control register	This register provides the copy of DESCR_X_CTL field of the currently active descriptor. If the currently active descriptor does not have X_CTL, this register provides undefined information.
DWx_ACT_DESCR_Y_CTL0	Active descriptor Y loop control register	This register provides the copy of DESCR_Y_CTL field of the currently active descriptor. If the currently active descriptor does not have Y_CTL, this register provides undefined information.
DWx_ACT_DESCR_NEXT_PTR0	Active descriptor next pointer register	This register provides the copy of DESCR_NEXT_PTR field of the currently active descriptor.
DWx_ACT_SRC0	Active source register	This register provides the current address of source location. This location is not a copy of the source address in the descriptor, but provides a real time source address of the active transfer.
DWx_ACT_DST0	Active destination register	This register provides the current address of destination location. This location is not a copy of the destination address in the descriptor, but provides a real time destination address of the active transfer.
DWx_ECC_CTL0	ECC control register	This register provides to specify the word address where an error will be injected and ECC parity to use for ECC error injection.
DWx_CRC_CTL0	CRC control register	This register provides to specify the bit order in which a data byte is processed (reversal is performed after XORing), and to specify whether the remainder is bit reversed (reversal is performed after XORing)
DWx_CRC_DATA_CTL0	CRC data control register	This register provides to specify a byte mask with which each data byte is XOR'd. The XOR is performed before data reversal.
DWx_CRC_POL_CTL0	CRC polynomial control register	This register provides to specify CRC polynomial. The polynomial is represented without the high order bit (this bit is always assumed '1'). The polynomial should be aligned/shifted such that the more significant bits (bit 31 and down) contain the polynomial and the less significant bits (bit 0 and up) contain padding '0's. Some frequently used polynomials: - CRC32: POLYNOMIAL is $0x04c11db7 (x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1)$. - CRC16: POLYNOMIAL is $0x80050000 (x^{16} + x^{15} + x^2 + 1, \text{shifted by } 16 \text{ bit positions})$. - CRC16 CCITT: POLYNOMIAL is $0x10210000 (x^{16} + x^{12} + x^5 + 1, \text{shifted by } 16 \text{ bit positions})$.

Table 7-26. P-DMA Registers

Register	Name	Description
DWx_CRC_LFSR_CTL0	CRC LFSR control register	This register provides the state of a 32-bit Linear Feedback Shift Registers (LFSR) that is used to implement CRC. This register needs to be initialized by software to provide the CRC seed value. The seed value should be aligned such that the more significant bits (bit 31 and down) contain the seed value and the less significant bits (bit 0 and up) contain padding '0's. Note that software can write this field. This functionality can be to used prevent information leakage (through DWx_CRC_LFSR_CTL0 or DWx_CRC_REM_RESULT0).
DWx_CRC_REM_CTL0	CRC remainder control register	This register provides to specifies a mask with which the DWx_CRC_LFSR_CTL0 register is XOR'd to produce a remainder. The XOR is performed before remainder reversal.
DWx_CRC_REM_RESULT0	CRC remainder result register	This register provides the remainder value. The alignment of the remainder depends on DWx_CRC_REM_CTL0.REM_REVERSE. Note that this field is combinatorially derived from DWx_CRC_LFSR_CTL0.LFSR32, DWx_CRC_CTL0.REM_REVERSE and DWx_CRC_REM_CTL0.REM_XOR.
DWx_CH_STRUCTy_CH_CTL	Channel control register	This register provides generic channel control information.
DWx_CH_STRUCTy_CH_STATUS	Channel status register	This register provides channel status which are the sources of interrupt factors and pending state.
DWx_CH_STRUCTy_CH_IDX	Channel current indices register	This register provides the current X and Y indices of the channel into the current descriptor.
DWx_CH_STRUCTy_CH_CURR_PTR	Channel current descriptor pointer register	This register provides the address of the memory location where the current descriptor is located.
DWx_CH_STRUCTy_INTR	Interrupt register	This register provides an interrupt request. Bit 0 is set 1 when interrupt event (completion or error) is detected. Software can clear this by writing to this bit.
DWx_CH_STRUCTy_INTR_SET	Interrupt set register	This register provides interrupt setting. Software can write 1 to this register to set the corresponding DMAC_CHx_INTR register. When read, this register reflects the DWx_CH_STRUCTy_INTR register.
DWx_CH_STRUCTy_INTR_MASK	Interrupt mask register	This register provides interrupt mask setting. The corresponding interrupt is enabled by writing 1 to this register.
DWx_CH_STRUCTy_INTR_MASKED	Interrupt masked register	This register provides interrupt masked. When read, this register reflects a bit-wise AND between the DWx_CH_STRUCTy_INTR and DWx_CH_STRUCTy_INTR_MASK fields.
DWx_CH_STRUCTy_SRAM_DATA0	SRAM data 0 register	DWx_CH_STRUCTy_SRAM_DATA0 and DWx_CH_STRUCTy_SRAM_DATA1 are provided for ECC fault injection functionality.
DWx_CH_STRUCTy_SRAM_DATA1	SRAM data 1 register	DWx_CH_STRUCTy_SRAM_DATA0 and DWx_CH_STRUCTy_SRAM_DATA1 are provided for ECC fault injection functionality.
DWx_CH_STRUCTy_TR_CMD	Software Trigger register	When written with '1', a trigger is generated which sets 'trigger pending' (only if the channel is enabled). A read always returns a '0'.

Note: In DWx_CH_STRUCTy, 'x' signifies the DW/P-DMA instance and 'y' signifies the channel number.

Table 7-27. M-DMA Registers

Register	Name	Description
DMAC_CTL	Control register	This register provides M-DMA enable/disable control
DMAC_ACTIVE	Active channels register	This register provides active channels
DMAC_CHx_CTL	Channel control register	This register provides generic channel control information.
DMAC_CHx_IDX	Channel current indices register	This register provides the current X and Y indices of the channel into the current descriptor.
DMAC_CHx_SRC	Channel current source address register	This register provides the current address of source location.
DMAC_CHx_DST	Channel current destination address register	This register provides the current address of destination location.
DMAC_CHx_CURR	Channel current descriptor pointer register	This register provides the address of the memory location where the current descriptor is located. When this field is 0, there is no valid descriptor.
DMAC_CHx_TR_CMD	Software trigger register	When written with '1', a trigger is generated which sets 'trigger pending' (only if the channel is enabled). A read always returns a '0'.
DMAC_CHx_DESCR_STATUS	Channel descriptor status register	This register provides the validity of other DMAC_CHx_DESCR registers.
DMAC_CHx_DESCR_CTL	Channel descriptor control register	This register provides the copy of DESCR_CTL field of the currently active descriptor.
DMAC_CHx_DESCR_SRC	Channel descriptor source register	This register provides the copy of DESCR_SRC field of the currently active descriptor.
DMAC_CHx_DESCR_DST	Channel descriptor destination register	This register provides the copy of DESCR_DST field of the currently active descriptor.
DMAC_CHx_DESCR_X_SIZE	Channel descriptor X size register	This register provides the copy of DESCR_X_SIZE field of the currently active descriptor.
DMAC_CHx_DESCR_X_INCR	Channel descriptor X increment register	This register provides the copy of DESCR_X_INCR field of the currently active descriptor.
DMAC_CHx_DESCR_Y_SIZE	Channel descriptor Y size register	This register provides the copy of DESCR_Y_SIZE field of the currently active descriptor.
DMAC_CHx_DESCR_Y_INCR	Channel descriptor Y increment register	This register provides the copy of DESCR_Y_INCR field of the currently active descriptor.
DMAC_CHx_DESCR_NEXT	Channel descriptor next pointer register	This register provides the copy of DESCR_NEXT_PTR field of the currently active descriptor.
DMAC_CHx_INTR	Interrupt register	This register provides an interrupt request. There are eight possible causes that can generate an interrupt. These causes are encoded in bits 0 to 7. Software can clear these by writing to these bits.
DMAC_CHx_INTR_SET	Interrupt set register	This register provides interrupt setting. There are eight bits (same as DMAC_CHx_INTR) and software can write 1 to any of these bits to set the corresponding INTR bit. When read, this register reflects the DMAC_CHx_INTR register.
DMAC_CHx_INTR_MASK	Interrupt mask register	This register provides interrupt mask setting for corresponding field in DMAC_CHx_INTR register. There are eight bits (same as DMAC_CHx_INTR) and they can be selectively enabled by writing 1 to the corresponding bits.
DMAC_CHx_INTR_MASKED	Interrupt masked register	When read, this register reflects a bitwise AND between the corresponding DMAC_CHx_INTR and DMAC_CHx_INTR_MASK fields.

The register access size and the initial value are described in the *Traveo II Body Controller High Registers TRM*.

Note: In DMAC_CHx, 'x' signifies the DMAC/M-DMA instance.

8. Code Flash



Code flash is a flash memory used to store programs. Code flash is a part of Cypress' eCT Flash, which is an embedded flash targeted for use in automotive applications. A common usage is as code storage for user application execution and local data storage/update for MCU-based systems in an automotive environment. The eCT Flash also includes work flash, which is the flash memory to store data; for more details, see the [Work Flash chapter on page 116](#).

8.1 Features

This section lists the features of code flash.

- Optional memory size: 4 MB, 6 MB, and 8 MB
- Programming and erasing functions
- ECC function: 64b + 8b
- Erase sector size of 32 KB for large sector and 8 KB for small sector
- Program size: 64b, 256b, and 4096b
- Supports Single Bank and Dual Bank modes
- Supports reading while programming/erasing
- Endurance of 1 k
- Retention of 20 years

Refer to the device datasheet for more information on the erase and program times.

8.2 Configuration


8.2.1 Block Diagram

[Figure 8-1](#) illustrates the position of code flash.

eCT Flash, which contains code flash is a part of the CPU subsystem. The Cortex-M7 cores can access code flash via AXI. and Cortex-M0+ core can access code flash via AHB. The CPU subsystem also has other subsystems connected with the AHB, such as DMA and Crypto.

The SROM APIs are designed for use with Arm Cortex-M0+ (CM0+) on Traveo devices. The SROM library includes APIs for flash programming and testing. The SROM APIs are executed within the Arm CM0+ IRQ0/1 exception generated using the IPC structures.

Figure 8-1. Position of Code Flash



*Refer to the device datasheet to see if this feature is supported

8.2.2 Flash Controller

The flash controller has multiple AHB-Lite interfaces and AXI bus interfaces:

- An AXI bus interface in the fast clock domain for Ethernet MAC 0/1
- An AHB-Lite bus interface in the slow clock domain for the CM0+ CPU.
- An AHB-Lite bus interface in the slow clock domain for Crypto.
- An AHB-Lite bus interface in the slow clock domain for P-DMA0/1.
- An AHB-Lite bus interface in the slow clock domain for DMA controller.
- An AHB-Lite bus interface in the slow clock domain for SDHC.

Note that each master has a dedicated AHB-Lite bus interface. This is unlike the ROM and SRAM controllers, where the slow bus masters are combined in the slow bus infrastructure.

This micro-architecture decision is driven by the difference in data width of the bus infrastructure (32 bit) and the flash memory (32-bit, 64-bit, 128-bit, or 256-bit): an AHB-Lite bus transfer is for a maximum of 32 bits, whereas a flash memory access always provides 32, 64, 128, or 256 bits. As flash memory accesses typically have wait states, it is beneficial to buffer or cache the complete flash memory data, rather than just selecting the requested 32 bits and discarding the rest of the flash memory data. Buffering or caching improves flash controller performance if bus transfers have temporal or spatial locality, as some bus


transfers can be served from the buffer or cache (without wait states), rather than requiring a flash memory access (with wait states). However, there is no temporal or spatial locality between bus transfers from different bus masters. Therefore, a dedicated buffer or cache is required for each bus master. Hence, the dedicated AHB-Lite bus interfaces.

Typically, Crypto, DataWire, and DMA controller transfers have more locality than CPU bus transfers. The former are typically sequential in nature, whereas the latter are less sequential due to jump/branch instructions. In addition, CPU performance is more important than Crypto, DataWire, or DMA controller performance. Therefore, Crypto, DataWire, or DMA controller transfers are supported through buffers only, whereas the CPU transfers are supported through a cache and a buffer. CPU interfaces support a cache for main interface flash memory data and a buffer for work interface memory data.

Other interfaces support a buffer for main interface flash memory data and a buffer for work interface memory data.

Figure 8-2 gives an overview of the flash controller micro-architecture.

Figure 8-2. Flash Controller



8.2.2.1 Bus Error

The flash controller generates an AHB-Lite/AXI bus error under the following conditions:

1. All flash macro write access.
2. A flash macro read access to a logical bank that is currently being programmed/erased.
3. A read access to a memory hole in the logical flash memory region. A memory hole is defined as a flash memory region address to a location that is not occupied by the code region, work region, or supervisory region.
4. Non-correctable ECC error resulting from read access.

The error responses due to 2, 3, and 4 above can be suppressed by setting FLASHC_FLASH_CTL.MAIN_ERR_SILENT.

Table 8-1. Flash Main Error Silent Register

Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	MAIN_ERR_SILENT	Specifies bus transfer behavior for a non-recoverable error on the flash macro main interface. 0: Bus transfer has a bus error. 1: Bus transfer does not have a bus error; that is, the error is silent.

The errors due to 2 and 3 above for read accesses from CPU masters are captured in FLASHC_CM0_STATUS/FLASHC_CM7_0_STATUS/FLASHC_CM7_1_STATUS registers.

Table 8-2. Flash CM0+/7_0/7_1 Main Status Register

Register	Bit Field and Bit Name	Description
FLASHC_CM0_STATUS	MAIN_INTERNAL_ERR	Specifies/registers the occurrence of a flash macro main interface internal error (typically the result of a read access while a program erase operation is ongoing) as a result of a CM0+ access. Software clears this field to '0'. Hardware sets this field to '1' on a flash macro main interface internal error. Typically, software reads this field after a code section to detect the occurrence of an error. Note: This field is independent of FLASHC_FLASH_CTL.MAIN_ERR_SILENT.
FLASHC_CM7_0_STATUS	MAIN_INTERNAL_ERR	See FLASHC_CM0_STATUS.MAIN_INTERNAL_ERROR.
FLASHC_CM7_1_STATUS	MAIN_INTERNAL_ERR	See FLASHC_CM0_STATUS.MAIN_INTERNAL_ERROR.

8.2.2.2 Wait Cycle Count

If CLK_MEM is higher than the maximum operating frequency of the flash memory, it is necessary to insert wait cycles when accessing the flash memory by setting an appropriate value in FLASHC_FLASH_CTL.MAIN_WS register.

Users can set MAIN_WS according to the followings;

- FLASHC_FLASH_CTL.MAIN_WS = 0 for CLK_MEM ≤ 100 MH
- FLASHC_FLASH_CTL.MAIN_WS = 1 for 100 MHz < CLK_MEM ≤ Fmem_max
Fmem_max refers to the maximum frequency of CLK_MEM

A cache miss will cause more than six FLASHC_FLASH_CTL.MAIN_WS wait states inserted for CPU main flash access.

Table 8-3. Flash Main Wait Status Register

Register	Bit Field and Bit Name	Description
FLASHC - FLASH_CTL	MAIN_WS	Flash macro main interface wait states: 0: 0 wait states. ... 15: 15 wait states

8.2.2.3 Power Modes

The flash controller provides Active functionality. In DeepSleep power mode, the following are retained:

- Retention Registers
- Cache tag structure: valid and tags registers
- Cache LRU structure
- Cache data structure: four SRAMs

Note that buffer information (in the AHB-Lite buffer interfaces and synchronization logic) are not retained.

Losing buffer information after deep-sleep transition has limited performance impact.

8.2.2.4 CM0+ CPU Cache

Note: The cache in this flash macro is only for CM0+ core and not for CM7 cores because CM7 cores have their own I and D caches.

The cache has the following features:

- 8 KB read-only capacity. This capacity provides a good hit rate for a range of benchmarks.
- Four-way set associative with an LRU replacement scheme. A four-way associative cache design provides a better hit rate than a direct mapped cache design at the same cache capacity.
- Sequential cache design. The cache tag functionality is performed before the cache data access. A sequential cache design has lower power consumption than a parallel cache design.
- 256 B line/sector, with thirty-two 32 B, sixteen 16 B, or eight 32 B subsectors each. For an 8 KB capacity, this results in a total of 32 lines distributed over eight sets. The subsector design allows for low overhead tag information, as the 16 subsectors in a line/sector share the tag and only have dedicated valid bits.

For each read transfer, the cache tag structure is evaluated before the cache data structure is accessed. The subsector design results in a relatively low number of 32 lines. The 32 associated tags are implemented in flip flops. The cache data structure is implemented using SRAM memory.

Read transfers that "hit" are processed by the cache. Read transfers that "miss" result in a flash controller access.

Each cache set has an associated 6-bit LRU field, which keeps track of the access history (from least recently used to most recently used) of the lines in the set.

Each cache line has an associated cache tag. The cache tag identifies the location of the line in system memory.


- The address bits that identify a byte in a cache line are not part of the cache tag (byte address bits 7 down to 0).
- The address bits that identify a cache set are not part of the cache tag (byte address bits 10 and 8).
- The address bits that are not part of the flash memory address (byte address bits 31 down to 27) are not part of the tag.

The above omissions of address bits result in small tags. As a result, the cache tag structure can be evaluated quickly.

In addition, the cache tag includes 16 valid bits – one valid bit for each subsector in the cache line.

Figure 8-3 gives an overview of the cache design.

Figure 8-3. Cache



A cache “miss” results in a 16 B (subsector) refill. The cache data structure is updated with 16 B of refilled data. Two cases are considered:

- The refilled data is a subsector of a resident cache line. Here, the data is refilled to the cache used by the resident cache line. The subsector's valid field is set to '1' (the valid fields of all other subsectors in the cache line remain unchanged).
- The refilled data is not a subsector of a resident cache line. Here, the data is refilled to the cache identified by the LRU scheme. The cache line address bits are updated, and the subsector's valid field is set to '1' (the valid fields of all other subsectors in the cache line are set to '0'). Note that this case replaces a resident cache line.

The cache has an LRU replacement scheme. Each cache set has an associated 6-bit LRU field:

- LRU[5]: '1' when way 0 is less recently used than way 1, '0' otherwise.
- LRU[4]: '1' when way 0 is less recently used than way 2, '0' otherwise.

- LRU[3]: '1' when way 0 is less recently used than way 3, '0' otherwise.
- LRU[2]: '1' when way 1 is less recently used than way 2, '0' otherwise.
- LRU[1]: '1' when way 1 is less recently used than way 3, '0' otherwise.
- LRU[0]: '1' when way 2 is less recently used than way 3, '0' otherwise.

Although six bits allow for $2^6 = 64$ -bit patterns, only $4 \times 3 \times 2 \times 1 = 24$ -bit patterns are legal LRU representations. The LRU set information is reset to all '1' or 0b111111, representing a set in which way 0 is less recently used than way 1, which is less recently used than way 2, which is less recently used than way 3. In this case, the line in way 0 is replaced when a new line is brought into the set. A line is made the most recently used line of its set, when it is brought into the set, or when its line data is used because of an AHB-Lite data transfer request.

Users can enable/disable the cache through CM0_CA_CTL.CA_EN.

Table 8-4. Flash Cache Enable Registers

Register	Bit Field and Bit Name	Description
FLASHC_CM0_CA_CTL0	CA_EN	Cache enable: 0: Disabled. 1: Enabled.

When the cache is disabled, the cache tag valid bits are reset to '0's and the cache LRU information is set to '1's (making way 0 the LRU way and way 3 the MRU way).

The cache supports prefetching through FLASHC_CM0_CA_CTL0.PREF_EN.

Table 8-5. Flash Cache Prefetch Enable Registers

Register	Bit Field and Bit Name	Description
FLASHC_CM0_CA_CTL0	PREF_EN	Prefetch enable: 0: Disabled. 1: Enabled.

Table 8-6. Flash ECC Enable Registers

Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	MAIN_ECC_EN	Enable ECC checking for flash main interface: 0: Disabled. No correctable or non-correctable faults are reported. 1: Enabled.

If prefetch is enabled, a cache miss results in a 16 B (subsector) refill for the missing data and a 16 B prefetch for the next sequential data (independent of whether this data is already in the cache). The data of the 16 B prefetch is stored in a temporary buffer and only copied into the cache when a future read transfer “misses” in the cache and requires the buffered data.


For debug purposes, the tag and 16 valid bits of a cache line are readable through registers. The LRU information of a cache set is readable through registers.

8.2.2.5 Code Flash ECC

The flash controller supports error correcting code (ECC) for the flash and cache SRAM memories. It can be enabled or disabled using the FLASHC_FLASH_CTL.MAIN_ECC_EN register field.

Figure 8-4 shows an overview of the data path of the flash ECC.

Figure 8-4. Data Path Overview for Flash ECC



ECC protection is added to the flash for functional safety. The ECC implements a Single Error Correction, Dual Error Detection (SECDED) scheme. In the flash code, 64 bits of data are covered by eight ECC bits.

Single-bit error correction is done in-line, without the need to stall the data returning to the flash controller. Adding the delay to correct single-bit errors to the read path will not have a significant effect on performance, due to flash cache. Flash read delay is long enough and ECC time will not be a significant portion of the read delay.

returns a bus error response. The fault handler is responsible for recovering from uncorrectable errors.

Fault Reporting

Both the correctable and non-correctable ECC errors are reported to the central fault structure in the same way.

All data correction and recovery are left to the ISR. There is no hardware support for writing corrected data back to flash.

Programming Code Flash

Code flash is 64 bits wide, and supports the 64-bit, 256-bit, and 4096-bit program.

For a 64-bit or larger program:

1. The first data input is stored in a write buffer and not written to flash.
2. When the last data input arrives, ECC is calculated, including the buffered data; 64-bit data + 8-bit ECC are programmed to flash.

ECC (Single-Bit Errors)

When the ECC logic detects a single-bit error, the error bit can be corrected. The error correction is in-line, so corrected data will be returned with no delay. The fault is reported through the regular fault structure.

ECC Uncorrectable Errors

If the ECC logic detects that the data has more than one bit wrong then the data cannot be corrected. In this case, the flash does not return the data to the CPU, and instead

Error Injection

Error injection is done through FLASHC_FLASH_CTL.MAIN/WORK_ECC_INJ_EN and FLASHC_ECC_CTL.PARITY/WORD_ADDR register fields.

Table 8-7. Flash ECC Error Injection Control Registers

Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	MAIN_ECC_INJ_EN	Enable error injection for flash main interface. 0: Disabled. 1: Enabled. When enabled, the parity bit (FLASHC_ECC_CTL.PARITY [31:24]) is used to load from the FLASHC_ECC_CTL.WORD_ADDR word address.
FLASHC_ECC_CTL	WORD_ADDR	Specifies the word address where an error will be injected. For flash main interface ECC, WORD_ADDR is device address A [26:3]. Device address A is defined as follows. A[31:27] = b'00010 A[26:3] = WORD_ADDER A[2:0] = b'000 On a flash main interface read and when FLASHC_FLASH_CTL.MAIN_ECC_INJ_EN bit is '1', PARITY replaces the flash macro parity.
FLASHC_ECC_CTL	PARITY	Specifies the ECC parity to use for ECC error injection at WORD_ADDR. For flash main interface ECC, the 8-bit parity (PARITY) is for a 64-bit word.

When error injection is enabled, the read address is compared to the device address A. If they are equal, the data read from flash is replaced with the parity register value.

It allows testing of the error recovery routines without continuous interrupts, as every flash read causes an error.

8.2.2.6 Software Generating Code Flash ECC

This section describes an algorithm to generate the correct ECC parity value with software. Note that this algorithm is not implemented in the hardware. Because the actual algorithm is optimized for hardware performance, it is different from the software algorithm described in this section.

“Value” in this algorithm represents the code flash 64-bit data value.

```
CW = 0x0000_0000_0000_0108_0000_0000_0000_0000 | Value
ECC_P0 = 0x01bf_bb75_be3a_72dc_4484_4a88_952a_ad5b
ECC_P1 = 0x02df_76f9_dd99_b971_1108_9311_26b3_366d
ECC_P2 = 0x04ef_cf9f_9ad5_ce97_0611_1c22_38c3_c78e
ECC_P3 = 0x08f7_ecf6_ed67_4e6c_9821_e043_c0fc_07f0
ECC_P4 = 0x10fb_7baf_6ba6_b5a6_e03e_007c_00ff_f800
ECC_P5 = 0x20fd_b7ce_f36c_ab5b_ffc0_007f_ff00_0000
ECC_P6 = 0x40fe_dd7b_74db_55ab_ffff_ff80_0000_0000
ECC_P7 = 0x807f_0000_07ff_ffff_d442_2584_4ba6_5cb7
```

```
parity[0] = ^ (CW & ECC_P0)
parity[1] = ^ (CW & ECC_P1)
...
parity[7] = ^ (CW & ECC_P7)
```

Note: “^” means reduction XOR. For example, $^(4'b0011) = 0^0^1^1$.

8.2.2.7 Cache ECC

The flash controller supports Error Correcting Code (ECC) for the cache SRAM memories. It can be enabled or disabled through the FLASHC_CM0_CA_CTL0.RAM_ECC_EN register fields.

Table 8-8. Flash Cache ECC Enable Registers

Register	Bit Field and Bit Name	Description
FLASHC_CM0_CA_CTL0	RAM_ECC_EN	Enable ECC checking for cache accesses: '0': Disabled. '1': Enabled.

Error Injection

The cache SRAM memory ECC uses 7-bit SECDED parity for each 32-bit data. The cache SRAM ECC supports error injection through FLASHC_CM0_CA_CTL0.RAM_ECC_INJ_EN and FLASHC_ECC_CTL.PARITY/WORD_ADDR register fields: on a fetch of flash memory data to the cache, the parity for a specific 32-bit word can be injected.

Table 8-9. Flash Cache ECC Error Injection Control Registers

Register	Bit Field and Bit Name	Description
FLASHC_CM0_CA_CTL0	RAM_ECC_INJ_EN	Enable error injection for cache. '0': Disabled. '1': Enabled. When enabled, the parity (FLASHC_ECC_CTL.PARITY) is used when a refill is done from the flash macro to the FLASHC_ECC_CTL.WORD_ADDR word address.
FLASHC_ECC_CTL	WORD_ADDR	Specifies the word address where an error will be injected. For cache SRAM ECC, WORD_ADDR is device address A [25:2]. Device address A is defined as follows. A[31:26] = b'000100 A[25:2] = WORD_ADDR A[1:0] = b'00 On a read from the code flash and CM0_CA_CTL.RAM_ECC_INJ_EN bit is '1', the parity (PARITY [6:0]) is injected and stored in the cache.
FLASHC_ECC_CTL	PARITY	Specifies the ECC parity to use for ECC error injection at WORD_ADDR. For cache SRAM ECC, the 7-bit parity is for a 32-bit word. The least significant 7 bits of PARITY will represent the 7-bit parity and the remaining parity bits are ignored.

8.2.2.8 Software Generating Cache ECC

This section describes an algorithm to generate the correct ECC parity value with software. Note that this algorithm is not implemented in the hardware. Because the actual algorithm is optimized for hardware performance, it is different from the software algorithm described in this section.

"Value" in this algorithm represents the code flash 32-bit data value to be fetched to the cache.

```

CW = 0x0000_0000_0000_0000 | Value
ECC_P0 = 0x037f_36db_2254_2aab
ECC_P1 = 0x05bd_eb5a_4499_4d35
ECC_P2 = 0x09dd_dcee_08e2_71c6
ECC_P3 = 0x11ee_bba9_8f03_81f8
ECC_P4 = 0x21f6_d775_f003_fe00
ECC_P5 = 0x41fb_6db4_fffc_0000
ECC_P6 = 0x8103_fff8_112c_965f

parity[0] = ^ (CW & ECC_P0)
  
```

```
parity[1] = ^ (CW & ECC_P1)
...
parity[6] = ^ (CW & ECC_P6)
```

Note: “^” means reduction XOR. For example, $\wedge(4'b0011) = 0^0^1^1$.

8.2.3 Flash Geometry

8.2.3.1 Interface, Regions, and Type of Use


eCT Flash is divided into work flash and code flash.

The top sectors in code flash are assigned as supervisory region and other sectors are assigned as code region. All sectors in work flash are assigned as work region.

The supervisory area is used to store trim parameters, system configuration parameters, protection and security settings, boot scripts, and other Cypress proprietary information. Read access to this region is permitted, but program/erase access is prohibited. Code region is the memory field to store program code flash. Work region is the memory field to store data.

Note that although supervisory region is located in code flash and it is contiguous with code region physically, the memory address of supervisory region is separated from code region. Work region is located between them as shown in [Figure 8-5](#).

Figure 8-5. Regions of eCT Flash



8.2.3.2 Geometries

eCT code sectors are composed of some memory units.

- Program Word: This is the unit of program. It is the smallest unit of code flash, including 64 bits for data and 8 bits for ECC.
- Read Word: This is the unit of read. It is composed of four units of Program Word.

- Page: This is composed of 16 units of Read Word.
- Erase Sector: This is the unit of erase, which has the following types:
 - Large sector – composed of 64 pages.
 - Small sector – composed of 16 pages.

Figure 8-6 shows the geometries for code flash.

Figure 8-6. Code Flash Sector Organization





Figure 8-7 shows the code flash arrays for each memory size. Note that upper two large sectors belong to supervisory region and do not count for code programming.

Note: “LRG” refers to large sector and SMS refers to small sector

Figure 8-7. Code Flash Array Organization



8.2.3.3 Logical Bank

This flash memory controller has the dual bank mode feature. When using dual bank mode, flash memory region is split into two half banks. One is called Logical Bank 0 and the other is called Logical Bank 1. Flash memory always has two logical banks regardless of its size. Figure 8-7 shows an illustration of the Logical Bank. See 8.2.4 Over-The-Air (OTA) Support for details about dual bank mode.

8.2.4 Over-The-Air (OTA) Support

OTA indicates that the flash macro supports a Read While Write operation on the same flash (code or work). OTA is possible on a Logical Bank resolution. This means a write can be done on one Logical Bank and a read can be done from any of the other Logical Banks in the non-write Logical Bank. If the read is done from the same Logical Bank, it will result in an error. In addition, a parallel read from the non-accessed Logical Bank can be performed.

8.2.4.1 Dual Bank Mode and Remap Functionality

The main flash region supports dual bank mode. The user can select the mode through FLASHC_FLASH_CTL.MAIN_BANK_MODE.

Table 8-10. Flash Main Bank Mode Register

Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	MAIN_BANK_MODE	Specifies bank mode of flash macro main array. 0: Single bank mode. 1: Dual bank mode.

This is to support OTA updates of the software image in flash memory. For example, the CPU executes from a current software image in the lower sectors while the higher sectors are programmed with a new software image. When the CPU reboots, the user code changes the MAIN_MAP field, such that the CPU executed from the new image is on the higher sectors.

The hardware remap functionality only affects the read flash region access path; it does not affect the write/program flash access path. The device SROM flash management APIs will perform all necessary address conversions; users do not have to consider this read/write address mismatch.

These address maps are configurable to support bank swapping as follows:

- When configuring Single Bank mode, the entire code and supervisory logical regions are mapped as a single contiguous address region, starting with all large sectors, followed by all small sectors.
- When configuring Dual Bank mode, these logical regions are split into two halves each, and each half is presented

as a separate address region. Furthermore, these halves can be swapped, to support same-location firmware upgrades.

- Choosing Mapping A will present the first half in the lower region and the second half in the upper region.
- Choosing Mapping B will present the first half in the upper region and the second half in the lower region.

Users can select mapping mode through FLASHC_FLASH_CTL.MAIN_MAP.

Table 8-11. Flash Main Remap Register


Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	MAIN_MAP	Specifies remapping of flash macro main region. 0: Mapping A. 1: Mapping B. This field is only used when MAIN_BANK_MODE is '1' (dual bank mode).

Address mappings for each of the six supported code flash densities are shown in the following sections.

8.2.4.2 Address Mapping for 4 MB Memory

The code region has 126 large sectors of 32 KB and 16 small sectors of 8 KB.


Figure 8-8. Code Flash Memory Mapping (4 MB)



8.2.4.3 Address Mapping for 6 MB Memory

The code region has 190 large sectors of 32 KB and 32 small sectors of 8 KB.

Figure 8-9. Code Flash Memory Mapping (6 MB)




6MB

8.2.4.4 Address Mapping for 8 MB Memory

The code region has 254 large sectors of 32 KB and 32 small sectors of 8 KB.

Figure 8-10. Code Flash Memory Mapping (8 MB)



8.3 Operation

Typically, APIs that are preinstalled in the SROM are used to operate the eCT Flash.

8.3.1 SROM APIs

See [37.3 SROM API Library](#) for details. To execute the SROM APIs, it is recommended to use the core M0+ through inter-processor communication. See the [Inter-Processor Communication chapter on page 48](#) for details.

ROM APIs related to code flash operation is listed in [Table 8-12](#).

Table 8-12. SROM APIs For Flash Operation

SROM API	Description
Program Row	Programs the addressed FLASH page
Erase All	Erases all FLASH
Erase Sector	Erases the addressed FLASH sector
Erase Suspend	Suspends ongoing erase operation
Erase Resume	Resumes an erase suspend operation

Notes:

- Reprogramming previously programmed words is not allowed without first erasing the sector. If the data value to be reprogrammed is the same as the value of programmed data words, reprogramming is permitted.
- The flash state will be unknown if reset/power-down occurs during program/erase. Erase the area because it may contain garbage data.

8.4 Registers

The following register map shows the various register definitions and its functionality.

Table 8-13. FLASHC Registers

Offset	Width	Name	Description
0x0000	32	FLASHC_FLASH_CTL	Control
0x0004	32	FLASHC_FLASH_PWR_CTL	Flash power control
0x0008	32	FLASHC_FLASH_CMD	Command
0x02a0	32	FLASHC_ECC_CTL	ECC control
0x0400	32	FLASHC_CM0_CA_CTL0	CM0+ cache control
0x0404	32	FLASHC_CM0_CA_CTL1	CM0+ cache control
0x0408	32	FLASHC_CM0_CA_CTL2	CM0+ cache control
0x0440	32	FLASHC_CM0_CA_STATUS0	CM0+ cache status 0
0x0444	32	FLASHC_CM0_CA_STATUS1	CM0+ cache status 1
0x0448	32	FLASHC_CM0_CA_STATUS2	CM0+ cache status 2
0x0460	32	FLASHC_CM0_STATUS	CM0+ interface status
0x04e0	32	FLASHC_CM7_0_STATUS	CM7 #0 interface status
0x0560	32	FLASHC_CM7_1_STATUS	CM7 #1 interface status
0x0580	32	FLASHC_CRYPT0_BUFF_CTL	Cryptography buffer control
0x0600	32	FLASHC_DW0_BUFF_CTL	Datavire 0 buffer control
0x0680	32	FLASHC_DW1_BUFF_CTL	Datavire 1 buffer control
0x0700	32	FLASHC_DMAC_BUFF_CTL	DMA controller buffer control
0x0780	32	FLASHC_SLOW0_MS_BUFF_CTL	Slow external master 0 buffer control

Table 8-14. FM_CTL_ETC Registers

Offset	Width	Name	Description
0x0400	32	FLASHC_MAIN_FLASH_SAFETY	Main (Code) flash security enable
0x0404	32	FLASHC_STATUS	Status read from flash macro
0x0500	32	FLASHC_WORK_FLASH_SAFETY	Work flash security enable

9. Work Flash



Work flash is a flash memory used to store data. Work flash is a part of Cypress' eCT Flash, which is an embedded flash targeted for use in automotive applications. A common usage is as code storage for user application execution and local data storage/update for MCU-based systems in an automotive environment. The eCT Flash also includes code flash, which is the flash memory to store programs; for more details, see the [Code Flash chapter on page 99](#).

9.1 Features

This section lists the features of work flash.

- Optional memory size: 256 KB
- Programming and erasing functions
- ECC function: 32b + 7b
- Erase sector size is 2 KB for large sector and 128 B for small sector
- Program size: 32b
- Supports Single Bank and Dual Bank modes
- Supports reading while programming/erasing
- Supports differential sensing architecture.
- Endurance of 250 k
- Retention of 10 years

Refer to the device datasheet for more information on the erase and program times.

9.2 Configuration


9.2.1 Block Diagram

[Figure 9-1](#) illustrates the position of work flash.

eCT Flash, which contains work flash is a part of the CPU subsystem. The Cortex-M7 cores can access work flash via AXI, and Cortex-M0+ core can access work flash via AHB. The CPU subsystem also has other subsystems connected with the AHB, such as DMA and Crypto.

The SROM APIs are designed for use with Arm Cortex-M0+ (CM0+) on Traveo devices. The SROM library includes APIs for flash programming and testing. The SROM APIs are executed within the Arm CM0+ IRQ0/1 exception generated using the IPC structures.

Figure 9-1. Position of Work Flash



9.2.2 Flash Controller

Refer to [Flash Controller](#) on page 100.

9.2.2.1 Bus Error

The flash controller generates an AHB-Lite/AXI bus error under the following conditions:

- All flash macro write access.
- A flash macro read access to a logical bank that is currently being programmed/erased.
- A read access to a memory hole in the logical flash memory region. A memory hole is defined as a flash memory region address to a location that is not occupied by the code region, work region, or supervisory region.
- Non-correctable ECC error resulting from read access.

The error responses due to 2, 3, and 4 can be suppressed by setting FLASHC_FLASH_CTL.WORK_ERR_SILENT.

Table 9-1. Flash Work Error Silent Register

Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	WORK_ERR_SILENT	Specifies bus transfer behavior for a non-recoverable error on the flash macro work interface. 0: Bus transfer has a bus error. 1: Bus transfer does not have a bus error; that is, the error is silent.

The errors due to 2 and 3 for read accesses from CPU masters are captured in the FLASHC_CM0_STATUS/FLASHC_CM7_0_STATUS/FLASHC_CM7_1_STATUS registers.

Table 9-2. Flash CM0+/7_0/7_1 Work Status Register

Register	Bit Field and Bit Name	Description
FLASHC_CM0_STATUS	WORK_INTERNAL_ERR	Specifies the occurrence of a flash macro work interface internal error (typically the result of a read access while a program erase operation is ongoing) as a result of a CM0+ access. Software clears this field to "0". Hardware sets this field to "1" on a flash macro work interface internal error. Typically, software reads this field after a work section to detect the occurrence of an error. Note: This field is independent of FLASHC_FLASH_CTL.WORK_ERR_SILENT.
FLASHC_CM7_0_STATUS	WORK_INTERNAL_ERR	See FLASHC_CM0_STATUS.WORK_INTERNAL_ERROR.
FLASHC_CM7_1_STATUS	WORK_INTERNAL_ERR	See FLASHC_CM0_STATUS.WORK_INTERNAL_ERROR.

9.2.2.2 Work Flash ECC

The flash controller supports error correcting code (ECC) for the work flash. It can be enabled or disabled using the FLASHC_FLASH_CTL.WORK_ECC_EN register field.

Table 9-3. Flash ECC Enable Registers

Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	WORK_ECC_EN	Enable ECC checking for flash work interface: 0: Disabled. No correctable or non-correctable faults are reported. 1: Enabled

Refer to [Figure 8-4](#) for an overview of the flash ECC data path.

ECC protection is added to the flash for functional safety. The ECC implements a Single Error Correction, Dual Error Detection (SECCDED) scheme. The flash work area has 32-bit data, covered by seven ECC bits.

ECC (Single-Bit Errors)

Refer to [ECC \(Single-Bit Errors\)](#) on page 105 for details.

ECC Uncorrectable Errors

Refer to [ECC Uncorrectable Errors](#) on page 105 for details.

Fault Reporting

Refer to [Fault Reporting](#) on page 105 for details.

Error Injection

Error injection is done through the WORK_ECC_INJ_EN and FLASHC_ECC_CTL.PARITY/WORD_ADDR register fields.

Table 9-4. Flash ECC Error Injection Control Registers

Register	Bit Field and Bit Name	Description
FLASHC_FLASH_CTL	WORK_EC-C_INJ_EN	Enable error injection for flash work interface. 0: Disabled 1: Enabled When enabled, the parity bit (FLASHC_ECC_CTL.PARITY) is used to load from the FLASHC_EC-C_CTL.WORD_ADDR word address.
FLASHC_EC-C_CTL	WORD_ADDR	Specifies the word address where an error will be injected. For flash work interface ECC, WORD_ADDR is device address A [25:2]. Device address A is defined as follows. A[31:26] = b'000101 A[25:2] = WORD_ADDR A[1:0] = b'00 On a flash work interface read and when FLASHC_FLASH_CTL.WORK_ECC_INJ_EN bit is '1', PARITY replaces the flash macro parity.
FLASHC_EC-C_CTL	PARITY	Specifies the ECC parity to use for ECC error injection at WORD_ADDR. For flash work interface ECC, the 7-bit parity is for a 32-bit word. The least significant 7 bits of PARITY will represent the 7-bit parity and the remaining parity bits are ignored.

When error injection is enabled, the read address is compared to device address A. If they are equal, the data read from flash is replaced with the parity register value.

It allows testing of the error recovery routines without continuous interrupts, as every flash read causes an error.

9.2.2.3 Software Generating Work Flash ECC

This section describes an algorithm to generate the correct ECC parity value with software. Note that this algorithm is not implemented in the hardware. Because the actual algorithm is optimized for hardware performance, it is different from the software algorithm described in this section.

“Value” in the algorithm represents work flash 32-bit data value.

```
CW = 0x0000_0007_0000_0000 | Value
ECC_P0 = 0x037f_36db_2254_2aab
ECC_P1 = 0x05bd_eb5a_4499_4d35
ECC_P2 = 0x09dd_dcee_08e2_71c6
ECC_P3 = 0x11ee_bba9_8f03_81f8
ECC_P4 = 0x21f6_d775_f003_fe00
ECC_P5 = 0x41fb_6db4_fffc_0000
ECC_P6 = 0x8103_fff8_112c_965f
```

```
parity[0] = ^ (CW & ECC_P0)
parity[1] = ^ (CW & ECC_P1)
...
parity[6] = ^ (CW & ECC_P6)
```

Note: “^” means reduction XOR; for example, $^ (4'b0011) = 0^0^1^1$.

9.2.3 Flash Geometry

9.2.3.1 Interface, Regions, and Type of Use


eCT Flash is divided into work flash and code flash.

The top sectors in code flash are assigned as supervisory region and other sectors are assigned as code region. All sectors in work flash are assigned as work region.

The supervisory area is used to store trim parameters, system configuration parameters, protection and security settings, boot scripts, and other Cypress proprietary information. Read access to this region is permitted, but program/erase access is prohibited. Code region is the memory field to store program code. Work region is the memory field to store data.

Note that although supervisory region is located in code flash and it is contiguous with code region physically, the memory address of supervisory region is separated from the code region. Work region is located between them as shown in [Figure 9-2](#).

Figure 9-2. Regions of eCT Flash



9.2.3.2 Geometries

eCT work sectors are composed of some memory units.

- **Word:** This is the unit of data. It is the smallest unit of work flash, including 32 bits for data and 7 bits for ECC.
- **Page:** This is composed of 16 units of Word (64 B, 624 bits).
- **Erase sector:** This is the unit of erase, which has the following types:
 - **Large sector:** composed of 32 pages (2 KB)
 - **Small sector:** composed of two pages (128 B)

Figure 9-3 shows the geometries for work flash.

Figure 9-3. Work Flash Sector Organization





Figure 9-4 shows the work flash arrays for each memory size. “LRG” stands for large sector and “SMS” stands for small sector.

Figure 9-4. Work Flash Array Organization

LRG = Large sector (2KB)
SMS = Small sector (128B)
 : Logical Bank 0
 : Logical Bank 1



9.2.3.3 Logical Bank

This flash memory controller has the dual bank mode feature. When using dual bank mode, flash memory region is split into two half banks. One is called Logical Bank 0 and the other is called Logical Bank 1. Flash memory always has two logical banks regardless of its size. [Figure 9-4](#) shows an illustration of the Logical Bank. See [9.2.4 Over-the-Air \(OTA\) Support](#) for details about dual bank mode.

9.2.4 Over-the-Air (OTA) Support

In OTA, the flash macro supports a read-while-write operation on the same flash (that is, code or work). OTA is possible on a Logical Bank resolution. This means a write can be done on one Logical Bank and a read can be done from any of the other Logical Banks in the non-write Logical Bank. In case the read is done from the same Logical Bank, it will result in an error. In addition, a parallel read from the non-accessed Logical Bank can be performed.

9.2.4.1 Dual Bank Mode and Remap Functionality

The work flash region supports dual bank mode. This mode can be selected using FLASHC_FLASH_CTL.WORK_BANK_MODE.

Table 9-5. Flash Work Bank Mode Register

Register	Bit Field and Bit Name	Description
FLASHC_- FLASH_CTL	WORK_BANK_MODE	Specifies bank mode of flash macro work array. 0: Single bank mode. 1: Dual bank mode.

The hardware remap functionality only affects the read flash region access path; it does not affect the write/program flash access path. The device SROM flash management APIs will perform all necessary address conversions; users do not have to consider this read/write address mismatch.

These address maps are configurable to support bank swapping as follows:

- When configuring Single Bank mode, the entire work region is mapped as a single contiguous address region, starting with all large sectors, followed by all small sectors.
- When configuring Dual Bank mode, this logical region is split into two halves, and each half is presented as a separate address region. Furthermore, these halves can be swapped to support same-location firmware upgrades.
 - Mapping A will present the first half in the lower region and the second half in the upper region.
 - Mapping B will present the first half in the upper region and the second half in the lower region.

Users can select the mapping mode using FLASHC_FLASH_CTL.WORK_MAP.

Table 9-6. Flash Work Remap Register


Register	Bit Field and Bit Name	Description
FLASHC_- FLASH_CTL	WORK_MAP	Specifies remapping of flash macro work region. 0: Mapping A. 1: Mapping B. This field is only used when WORK_BANK_MODE is '1' (dual bank mode).

9.2.5 Address Map of Work Flash

9.2.5.1 Address Mapping for 128 KB Memory

The work region has 48 large sectors of 2 KB and 256 small sectors of 128 B.


Figure 9-5. Work Flash Memory Mapping (128 KB)



9.2.5.2 Address Mapping for 256 KB Memory

The work region has 96 large sectors of 2 KB and 512 small sectors of 128 B.

Figure 9-6. Work Flash Memory Mapping (256 KB)



9.3 Operation

Typically APIs that are preinstalled in the SROM are used to operate the eCT Flash. This section provides a brief summary of SROM APIs.

9.3.1 Read

There are some read restrictions due to the nature of differential flash. Normal usage of work flash is as follows:

1. Erase entire sector
2. Program words
3. Read

Note:

- Reading a word that is still in the erased state will result in random (spurious) data.
 - Caution: This reading will cause an ECC error.
 - Measures: Call the blank check SROM API before reading work flash to check whether the area is in the programmed or erased state.
- Work flash is always read 64-bit wide via AXI.
 - Caution: A read of 32-bit unit data on work flash results in 64-bit access. If the adjacent 32-bit data has not been programmed, it will result in an unexpected ECC error.
 - Measures: Take one of the following measures:
 1. Use DMA (M-DMA or P-DMA) that reads via AHB, which has a 32-bit width.
 2. Program data into the work flash – always aligned to 64 bits and by 64-bit units.
- ECC error can be notified to only one CPU via fault structure.
 - Caution: When multiple cores have to read from work flash, all the cores, except one core, cannot be notified about the ECC error
 - Measures: Take one of the following measures:
 1. Use DMA (M-DMA or P-DMA) to read from work flash. If a non-correctable error occurs during the DMA transmission, it will be detected and informed via one of the following DMA registers:
 - a. M-DMA: Source bus error bit of the interrupt register (DMAC_CHx_INTR.SRC_BUS_ERROR = 1).
 - b. P-DMA: Interrupt cause bit of the status register (DWx_CH_STRUCTy_CH_STATUS.INTR_CAUSE = 2).
 DMA does not detect correctable ECC errors.
 2. Assign one CPU core for non-correctable ECC error handling. This core informs about the error to the core that caused the error,
 3. Set non-correctable ECC error action to reset (may not be acceptable depending to the application).

9.3.2 SROM APIs

Refer to [SROM API Library on page 697](#) for details.

To execute the following SROM APIs, it is recommended to use the core M0+ through inter-processor communication. See the [Inter-Processor Communication chapter on page 48](#) for details.

SROM APIs related to work flash operation are listed in [Table 9-7](#).

Table 9-7. SROM APIs for Flash Operation

SROM API	Description
Program Row	Programs the addressed flash page
Erase All	Erases all flash

Table 9-7. SROM APIs for Flash Operation

SROM API	Description
Erase Sector	Erases the addressed flash sector
Erase Suspend	Suspends ongoing erase operation
Erase Resume	Resumes an erase suspend operation
Blank check	Performs blank check on the addressed work flash

Notes:

- Reprogramming previously programmed words is not allowed without first erasing the sector. If the data value to be reprogrammed is the same as the value of programmed data words, reprogramming is permitted.
- The flash state will be unknown if reset/power-down occurs during program/erase. Because it may contain garbage data, run blank check; if it is not blank, erase that area.

9.4 Registers

The following register map shows various register definitions and its functionality.

Table 9-8. Registers

Offset	Width	Name	Description
0x0000	32	FLASHC_FLASH_CTL	Control
0x0004	32	FLASHC_FLASH_PWR_CTL	Flash power control
0x0008	32	FLASHC_FLASH_CMD	Command
0x02a0	32	FLASHC_ECC_CTL	ECC control
0x0400	32	FLASHC_CM0_CA_CTL0	CM0+ cache control
0x0404	32	FLASHC_CM0_CA_CTL1	CM0+ cache control
0x0408	32	FLASHC_CM0_CA_CTL2	CM0+ cache control
0x0440	32	FLASHC_CM0_CA_STATUS0	CM0+ cache status 0
0x0444	32	FLASHC_CM0_CA_STATUS1	CM0+ cache status 1
0x0448	32	FLASHC_CM0_CA_STATUS2	CM0+ cache status 2
0x0460	32	FLASHC_CM0_STATUS	CM0+ interface status
0x04e0	32	FLASHC_CM7_0_STATUS	CM7 #0 interface status
0x0560	32	FLASHC_CM7_1_STATUS	CM7 #1 interface status
0x0580	32	FLASHC_CRYPT0_BUFF_CTL	Cryptography buffer control
0x0600	32	FLASHC_DW0_BUFF_CTL	Datawire 0 buffer control
0x0680	32	FLASHC_DW1_BUFF_CTL	Datawire 1 buffer control
0x0700	32	FLASHC_DMAC_BUFF_CTL	DMA controller buffer control
0x0780	32	FLASHC_SLOW0_MS_BUFF_CTL	Slow external master 0 buffer control

Table 9-9. FM_CTL_ETC Registers

Offset	Width	Name	Description
0x0400	32	FLASHC_MAIN_FLASH_SAFETY	Main (Code) flash security enable
0x0404	32	FLASHC_STATUS	Status read from flash macro
0x0500	32	FLASHC_WORK_FLASH_SAFETY	Work flash security enable

10. SRAM Interface



SRAM controllers are implemented in the Traveo II family device for the on-chip SRAM memory interface. RAMs are accessible by several masters connected to the fast and slow infrastructures: CPUs, Peripheral-DMA (P-DMA), Memory-DMA (MDMA), Crypto, and external masters. CPUs can also execute code out from these SRAMs.

10.1 Features

This section lists the features of SRAM controller.

- Optional memory size: 1024 KB
- AXI bus interfaces:
 - In the fast clock domain for the CM7 CPUs
- AHB-Lite bus interface:
 - In the slow clock domain for all bus masters (CM0+ CPU, CRYPTO, P-DMA, M-DMA, debug interface, and optional external bus master). The slow bus infrastructure combines the bus masters in the slow clock domain.
- Programmable wait states.
- ECC function
 - Single-bit error correction and double-bit error detection (SECDED)
 - ECC error injection
- RAM retention function
- RAM power up delay control
 - Setting the power stabilization wait after switching on the SRAM power domain.

Note: The first 2 KB of SRAM is reserved and is not available for users. The first 32 KB block of SRAM0 should be in the enabled or retained state in Active, LP Active, Sleep, LP Sleep, and DeepSleep modes.


Note: The SRAM region from (SRAM size minus 6 KB) to (SRAM size minus 2 KB) and the SRAM region of the first word of the last 2 KB are used by Cypress firmware during boot operation. Therefore, this region is available to the user; however, data retention across resets is not guaranteed in this area, because it can be overwritten by Cypress boot firmware.

10.2 Configuration

10.2.1 Block Diagram

The SRAM controller has a 64-bit wide interface to SRAM memory. Figure 10-1 gives an overview of the SRAM controller.

Figure 10-1. SRAM Controller



The SRAM controller has one AXI interface and one AHB-Lite interfaces that connect to the AXI and AHB-Lite infrastructures. The AHB-Lite interface is connected to a synchronization component that translates between the interface clock (CLK_SLOW) and the high-frequency clock (CLK_MEM).

Arbitration is performed on the transfers from the two ports (AHB-Lite interface and AXI interface). Arbitration uses device wide, bus master specific priorities, and round-robin based acceptance within the same priority group. Therefore, although two bus interfaces are provided, one AHB-Lite or AXI transfer is accepted by the port arbitration component.

The SRAM controller supports Error Correcting Code (ECC) for SRAMs. This functionality can be disabled or enabled (CPUSS_RAMx_CTL0.ECC_EN). Initial value of CPUSS_RAMx_CTL0.ECC_EN is '1' (ECC enabled).

- If ECC functionality is disabled (CPUSS_RAMx_CTL0.ECC_EN = '0'), Read and write transfers originating from the AHB-Lite or AXI interface are accessing SRAM directly. Note, that when ECC is disabled, no parity information is written to the RAM. All data sizes can be written to the SRAM by a single access.
- If ECC functionality is enabled (CPUSS_RAMx_CTL0.ECC_EN = '1'), an AHB-Lite or

AXI transfer is translated into one or multiple SRAM accesses. Furthermore, ECC, and write buffer components are used to implement the desired functionality. Read data accesses use parity information stored beside each 64-bit word in the SRAM to correct single bit errors or to detect double bit errors. Write data accesses need to generate parity information for each 64-bit word in the SRAM. Writing full 64-bit size can be done by a single access. Writing smaller data sizes are translated into read accesses from the SRAM and combined (previous read and new partial write data) are stored into a write buffer. Such a partial write operation causes two accesses to the SRAM. Note, that partial write can achieve only half of the possible memory bandwidth. Typically, the pending writes from the write buffer are executed when the SRAM is not accessed by the AHB-Lite or AXI ports.

10.2.2 Wait States

The SRAM controller supports programmable wait states. Dedicated wait states are provided for the fast and slow AHB-Lite bus interfaces. The programmable wait states represent the number of CLK_MEM cycles for a read path through the SRAM memory in either the fast domain (CM7

CPUs and optional AXI masters) or slow domain (CM0+ CPU, P-DMA, M-DMA, etc.).

The SRAM controller supports wait states in the range from 0 to 3. The number of wait states is expressed in CLK_HF clock cycles.

- The wait states for the slow clock domain
 - CPUSS_RAMx_CTL0.SLOW_WS = '0' up to 100 MHz of CLK_HF.
 - CPUSS_RAMx_CTL0.SLOW_WS = '1' from 100 MHz to 200 MHz of CLK_HF.

- The wait states for the fast clock domain

CPUSS_RAMx_CTL0.FAST_WS = '0' up to 200 MHz of CLK_HF.

As the wait states are represented in CLK_MEM cycles, the wait states do not have to be reprogrammed when the fast clock domain frequency (CLK_FAST_0 or CLK_FAST_1) or slow clock domain frequency (CLK_SLOW) is changed. However, it may be necessary to reprogram the wait states when the high frequency clock domain (CLK_HF) or CLK_MEM is changed. To summarize, the required number of wait states is a function of the CLK_MEM frequency.

- The read path through the SRAM memory in the fast domain is faster than the read path through the SRAM memory in the slow domain, because fast clock domain is timing closed at a higher frequency than the slow clock domain. Therefore, the required number of fast wait states (CPUSS_RAMx_CTL0.FAST_WS) should be less or equal than the required number of slow wait states (CPUSS_RAMx_CTL0.SLOW_WS).

10.2.3 Operation

The following describes the SRAM controller with ECC functionality enabled.

SRAM accesses originate from one of the following paths:

- AHB-Lite or AXI transfers.
- Write buffer requests. If ECC functionality is disabled, this path is not used.
- SRAM repair requests. If ECC functionality is disabled, this path is not used.

The AHB-Lite and AXI transfers are the origin for all SRAM accesses; the write buffer and SRAM repair requests result from AHB-Lite and AXI transfers. The SRAM controller differentiates between the following three types of AHB-Lite and AXI transfers:

- AHB-Lite and AXI read transfers.
- 64-bit AXI write transfers.
- 8-bit, 16-bit, and 32-bit AHB-Lite or AXI write transfers (also referred to as partial AHB-Lite or AXI write transfers).

Each type is described in more detail here.

AHB-Lite and AXI read transfers. An AHB-Lite or AXI read transfer is translated into an SRAM read access using the ECC syndrome logic. The ECC syndrome logic corrects recoverable errors. If the read address matches in the write buffer, the SRAM has stale data and the write buffer provides the requested read data.

The ECC syndrome logic reports recoverable and non-recoverable errors to the fault reporting component in the SRAM controller.

A corrected, recoverable error requires an update of the SRAM: the SRAM address needs to be written/repared with the corrected code word.

This automatic repair functionality is enabled when CPUSS_RAMx_CTL0.ECC_AUTO_CORRECT is '1'.

64-bit AXI write transfers. A 64-bit AXI write transfer is translated into an SRAM write access, using the ECC parity logic. If the write address matches in the write buffer, the matching write buffer entries have stale data and these entries are invalidated.

Partial AHB-Lite and AXI write transfers. A partial AHB-Lite or AXI write transfer is translated into an SRAM read access and an SRAM write access. The SRAM read access is the direct result of the partial write transfer and the SRAM write access is the result of a write buffer request. A partial write transfer requires an SRAM read access to retrieve the “missing” data bytes from the SRAM. If the read address matches in the write buffer, the SRAM has stale data and the write buffer provides the requested read data. The requested read data is merged with the partial write data to provide a complete 64-bit data word. The address and the merged write data are written to the write buffer. A future write buffer request results in a SRAM write access with the merged write data.

Only the partial AHB-Lite and AXI write transfers of data size less or equal 32-bit (dependency on data size dependency on data size) use the write buffer.

10.2.4 Write Buffer

The write buffer is a temporary holding station for future SRAM write accesses.

The buffer allows SRAM write accesses to be postponed. This allows for more performance critical AHB-Lite or AXI requests to “overtake” write buffer requests. Memory consistency is guaranteed by matching the SRAM access address with the write buffer entries' addresses: a matching SRAM read access uses the read merge component and a matching SRAM write access invalidates the matching write buffer entries.

When the write buffer is full, an entry needs to be freed to accommodate future partial AHB-Lite or AXI write transfers. Therefore, a full write buffer raises the priority of the write buffer request path.

The state of the write buffer is reflected by `CPUSS_RAMx_STATUS.WB_EMPTY`. The write buffer is not retained in DeepSleep power mode. Therefore, when transitioning to system DeepSleep power mode, the write buffer should be empty.

Note that this requirement is typically met, because a transition to DeepSleep power mode also requires that there are no outstanding AHB-Lite or AXI transfers. If there are no outstanding AHB-Lite or AXI transfers, the write buffer gets SRAM access.

10.3 ECC Details

The SRAM controller supports ECC. Specifically, it supports a hamming code with an additional parity bit. This code supports single error correction, double error detection (SEDED). The ECC is applied to the SRAM data and SRAM address.

- The ECC corrects single-bit errors in an SRAM code word (stored in SRAM memory).
- The ECC detects single-bit and double-bit errors in an SRAM code word and the SRAM address.

The SRAM controller does not generate AHB-Lite or AXI bus errors. In the case of an ECC error, a correctable error is corrected on the fly and a non-correctable error is communicated through the fault reporting structure.

Note that the initial value of SRAM is undefined. Therefore, SRAM should be initialized before reading or partial writing to prevent unintentional ECC faults. For initialization, the `CPUSS_RAMx_CTL0.ECC_CHECK_DIS` bit can be used. When this bit is set, ECC check, notification for fault reporting, and ECC correction are disabled. Disable the `CPUSS_RAMx_CTL0.ECC_CHECK_DIS` bit only for initialization. This bit is ignored when `ECC_EN = 0`.

10.3.1 ECC Parity Generation for SRAM Write Accesses

For 64-bit AXI write bus transfers, only a single SRAM write access is required. For 8-bit, 16-bit, and 32-bit AHB-Lite and AXI write bus transfers, an additional SRAM read access precedes the SRAM write access to retrieve the “missing” data bytes. These missing bytes are required to construct the complete 64-bit data word. The 8-bit parity is calculated over the complete 64-bit data word.

10.3.2 ECC Syndrome Generation for SRAM Read Accesses

For read accesses, the syndrome specifies one of the following:

- No error is detected. The SRAM 64-bit data word can be used as the result for an AHB-Lite or AXI read bus transfer.
- A single error is detected in the data word. This error is recoverable. The syndrome specifies the bit error location. The correction process inverts the bit value at the error location. The corrected data word is used as the result for an AHB-Lite or AXI read bus transfer. An additional SRAM write access is required to update the SRAM code word in case `ECC_AUTO_CORRECT` feature is enabled.
- A single error is detected in the 8-bit parity. This error is recoverable. An additional SRAM write access is required to update the SRAM code word with the correct parity.
- A single error is detected in the word address. This error is non-recoverable.
- A double error is detected. This error is non-recoverable.

For AHB-Lite or AXI read bus transfers, typically only a single SRAM read access is required. However, when a recoverable error is detected, an additional SRAM write access is required for the `ECC_AUTO_CORRECT` feature. Recoverable errors are communicated through the fault reporting structure.

Note that when a non-recoverable error is detected, the data word that is used as the result for an AHB-Lite or AXI bus transfer is incorrect, but no AHB-Lite or AXI bus error is generated. Non-recoverable errors are communicated through the fault reporting structure.

The fault reporting structure supports two types of SRAM controller faults:

- Correctable ECC faults
- Non-correctable ECC faults

For both fault types, the same information is captured by the fault reporting structure:

- SRAM word address
- SRAM syndrome

Note that the SRAM code word (8-bit parity and 64-bit data word) are not captured.

Note, that fault reporting can capture only a certain rate of fault events. It cannot be guaranteed that each of the ECC faults can be captured in case of multiple ECC errors occur short after each other.

10.3.3 ECC Error Injection

The fault reporting structure for ECC faults can be debugged through an SRAM controller ECC parity injection mechanism. This mechanism functions as follows:

- ECC injection is enabled through CPUSS_RAMx_CTL0.ECC_INJ_EN (for SRAM controller x).
- A word address is specified by CPUSS_ECC_CTL.WORD_ADDR[23:0]
(CPUSS_ECC_CTL.WORD_ADDR = (0x00FFFFFF & (RAM_TEST_ADDRESS>>2))).
- A 8-bit parity is specified by CPUSS_ECC_CTL.PARITY[7:0].

When a write transfer to the specified word address is performed, the ECC parity generation uses the specified 8-bit parity, rather than the calculated parity. The data still originates from the bus transfer. Any access size can be used to inject parity. Note that parity injection invalidates the write buffer for this word address. If only a part of 64-bit data is written and consistency should be maintained, CPUSS_RAMx_STATUS.WB_EMPTY=1 should be checked before.

10.3.4 ECC Parity Generation by Software

To inject the ECC error for fault generation, ECC parity must be generated by software. Follow this procedure to generate 8-bit ECC parity.

```
CODEWORD_SW[127:0] = {128 {1'b0}};
CODEWORD_SW[63:0] = ACTUALWORD[63:0];
ADDR_WIDTH = log2(RAM_SIZE)
CODEWORD_SW[ADDR_WIDTH+60:64] = ADDR[ADDR_WIDTH-1:3];
```

Note: RAM_SIZE is size of RAMx, where “x” is the RAM unit number.

```
ECC_P0_SW = 128b00000001_10111111_10111011_01110101_10111110_00111010_01110010_11011100_
01000100_10000100_01001010_10001000_10010101_00101010_10101101_01011011;
ECC_P1_SW = 128b00000010_11011111_01110110_11111001_11011101_10011001_10111001_01110001_
00010001_00001000_10010011_00010001_00100110_10110011_00110110_01101101;
ECC_P2_SW = 128b00000100_11101111_11001111_10011111_10011010_11010101_11001110_10010111_
00000110_00010001_00011100_00100010_00111000_11000011_11000111_10001110;
ECC_P3_SW = 128b00001000_11110111_11101100_11110110_11101101_01100111_01001110_01101100_
10011000_00100001_11100000_01000011_11000000_11111100_00000111_11110000;
ECC_P4_SW = 128b00010000_11111011_01111011_10101111_01101011_10100110_10110101_10100110_
11100000_00111110_00000000_01111100_00000000_11111111_11111000_00000000;
ECC_P5_SW = 128b00100000_11111101_10110111_11001110_11110011_01101100_10101011_01011011_
11111111_11000000_00000000_01111111_11111111_00000000_00000000_00000000;
ECC_P6_SW = 128b01000000_11111110_11011101_01111011_01110100_11011011_01010101_10101011_
11111111_11111111_11111111_10000000_00000000_00000000_00000000_00000000;
ECC_P7_SW = 128b10000000_01111111_00000000_00000000_00000111_11111111_11111111_11111111_
11010100_01000010_00100101_10000100_01001011_10100110_01011100_10110111;
```

As shown here, Reduction XOR of the ANDed result of CODEWORD_SW[127:0] and respective ECC constants will give a single parity bit.

```
parity[0] = ^ (CW_SW[127:0] & ECC_P0_SW)
parity[1] = ^ (CW_SW[127:0] & ECC_P1_SW)
...
parity[7] = ^ (CW_SW[127:0] & ECC_P7_SW)
```

Parity[6:0] gives seven bits parity for 32 bits ACTUALWORD[127:0].

10.4 RAM Retention Configuration

This section covers the steps for transitioning to the RAM retention mode in Traveo II. The registers in [Table 10-1](#) and [Table 10-2](#) are used. The write buffer should be empty (CPUSS_RAMx_STATUS.WB_EMPTY = 1), when the mode is set to RETAINED,

Table 10-1. Power Control Register

Register	Bit Field	Bit Value	Mode	Description
CPUSS_RAM0_PWR_MACRO_CTLy for SRAM#0 ^a CPUSS_RAMx_PWR_CTL for SRAM other than RAM#0 ^a	PWR_MODE	0	OFF	Switch SRAM off
		2	RETAINED	Put SRAM in retained mode
		3 (Default)	ENABLED	Switch SRAM on
	VECTKEYSTAT ^b	0xfa05		Register key (to prevent accidental writes). ■ Should be written with a 0x05fa key value for the write to take effect. ■ Always reads as 0xfa05.

a. SRAM#0 can be fully retained or retained in increments of 32-KB sectors. SRAM unit other than RAM#0 can be retained as a whole unit.

b. VECTKEYSTAT must be written at the same time as PWR_MODE. These registers should be written as the complete 32-bit data.

Table 10-2. RAM Status Register

Register	Bit Field	Description
CPUSS_RAMx_STATUS	WB_EMPTY	Write buffer empty. '0': Write buffer not empty. '1': Write buffer empty.

As mentioned earlier, when transitioning to DeepSleep mode, the write buffer should be empty (CPUSS_RAMx_STATUS.WB_EMPTY = 1).

1. Checking the CPOUSS_RAMx_STATUS.WB_EMPTY register and wait until the WB_EMPTY bit becomes 1.
2. When WB_EMPTY bit becomes 1, set the retained mode to the CPOUSS_RAM0_PWR_MACRO_CTLy or CPOUSS_RAMx_PWR_CTL register.
3. Transfer to DeepSleep mode or issues the software reset.
4. When returning from DeepSleep, it is necessary to set to enable mode before using RAM.

Note: SRAM0_PWR_MACRO_CTL0.PWR_MODE must be set to ENABLE or RETAINED in Active, LP Active, Sleep, LP Sleep, and DeepSleep modes.

10.5 Registers

Table 10-3. List of Registers

Registers Name	Name	Description
CPUSS_RAMx_CTL0	RAMx control register	Specify the operation of the RAMx controller.
CPUSS_RAMx_STATUS	RAMx status register	Indicates RAMx controller status.
CPUSS_RAM0_PWR_MACRO_CTLy	RAM0 power control register	These registers control the system SRAM 0 power states of a single macro. System SRAM 0 consists of up to sixteen 32 kB macros. Each macro is a single power partition and is controlled through a dedicated control field in one of these registers.
CPUSS_RAMx_PWR_CTL	RAMx power control register	This register controls the system SRAMx power states. System SRAMx consists of a single power partition.
CPUSS_RAM_PWR_DELAY_CTL	RAM power up delay control register	Number clock cycles delay needed after power domain power up.
CPUSS_ECC_CTL	ECC control register	Specifies the word address and ECC parity where an error will be injected.

Note: The 'x' in the register name denotes the SRAM memory unit number. The "y" in the register name denotes the SRAM0 memory macro number. Refer to the device datasheet for the specifications.

11. BootROM



System boot is defined as the process of obtaining, validating, and starting the product firmware. Traveo II has embedded ROM, and flash performs its entire boot process in software. The main function of the boot process is to configure the system (apply trims and wounding information, and configure access and protection settings according to the product life-cycle stage), authenticate the application, and transfer control to the application.

11.1 Features

The BootROM of Traveo II supports the following features:

- After any type of reset, the boot code starts execution from ROM on the CM0+.
- The boot process consists of two parts: ROM boot process and flash boot process.
See the [Flash Boot chapter on page 740](#) for more details.
- The ROM boot code applies life-cycle stage and protection state.
- The ROM boot code validates the integrity of the flash boot process before starting it.

11.2 ROM Controller

The Traveo II series has a supervisory ROM that contains BootROM code and SROM APIs. This section gives a brief overview of the ROM controller.

The ROM controller has two AHB-Lite bus interfaces:

- An AHB-Lite bus interface in the slow clock domain for all bus masters in the slow clock domain (CM0+ CPU, P-DMA, M-DMA, and so on). The slow bus infrastructure combines the bus masters in the slow clock domain.

11.2.1 Wait States

The ROM controller supports programmable wait states, which is defined at SLOW_WS[1:0] and FAST_WS[1:0] in the CPUSS_ROM_CTL register. Dedicated wait states are provided for the fast and slow AHB-Lite bus interfaces. The programmable wait states represent the number of CLK_MEM cycles for a read path through the SROM in either the fast or slow domain.

The ROM controller supports wait states in the range 0 to 3. The number of wait states is expressed in CLK_MEM clock cycles. The application that changes CLK_MEM must set the ROM wait states corresponding to the new target of CLK_MEM frequency before CLK_MEM is raised.

- The wait states for the slow clock domain are:
 - CPUSS_ROM_CTL.SLOW_WS = '0' up to 100 MHz of CLK_MEM
 - CPUSS_ROM_CTL.SLOW_WS = '1' from 100 MHz to 200 MHz CLK_MEM
- The wait state for the fast clock domain is:
 - CPUSS_ROM_CTL.FAST_WS = '0' up to 200 MHz of CLK_MEM

11.3 ROM Boot Process

11.3.1 Life-Cycle Stages and Protection States

Life-cycle stages are governed by eFuse and are irreversible. A powered device also has a volatile protection state that reflects its life cycle stage. The protection state is determined on boot and defined by the value of the CPUSS_PROTECTION register. For more details, see the [Device Security chapter on page 158](#).

11.3.2 Multicore Boot

Traveo II starts up with all cores except CM0+ in reset. The ROM boot process will execute on M0+. Its main purpose is to start an M0+ flash boot process.

When other CPUs start, they will enter the ROM boot process. The ROM boot process will recognize this and jump directly to the code pointed by CPUSS.CM7_0/CM7_1_VECTOR_TABLE_BASE for CM7_0/CM7_1 bypassing the full boot. The CPUSS_IDENTITY register is used to determine which CPU is executing the ROM boot process.

11.3.3 Secure Boot

Before CM0+ executes the firmware in supervisory flash and the life-cycle stage is only SECURE and SECURE_WITH_DEBUG, it authenticates flash boot code by comparing the pre-computed SECURE_HASH stored in eFuse with the generated one. Flash boot code will be executed only if it is found to be authentic; otherwise, boot code enters DEAD protection state.

11.3.4 Protection Setting

ROM boot reads the configurations of SMPU, PPU, and SWPU from SFlash and programs the protection units accordingly.

■ DAP Memory Protection Unit (MPU)

This is used to restrict the access rights of DAP as indicated by NORMAL, SECURE, and DEAD access restrictions. The boot uses eight memory regions of MPU to implement the access restrictions.

■ Shared Memory Protection Unit (SMPU)

These are used to implement access restrictions to memory such as ROM, Flash, and RAM. ROM/flash boot reads the SMPU configuration from SFlash and programs the corresponding SMPU registers.

■ Software Protection Unit (SWPU)

These are used to implement access restrictions to flash (program/erase) and eFuse (read/write). There are 32 entries in SWPU. The SWPU is broken into two parts. The first part is stored in SFlash and implements the access restrictions related to PC1 and PCx. Here PC1 means protection context 1 and PCx means one of protection context {2, 3, ..., 15}. See the ["Protection Context" on page 55](#) for details. The second part is stored in SFlash and is used by the application for additional access restrictions specific to the application. ROM/flash boot reads the two parts of SWPU from SFlash and stores them in RAM.

■ Peripheral Protection Unit (PPU)

These are used to implement access restrictions to peripheral registers. Only a subset of the PPU are required to enforce protection for PC1 and PCx and only these are stored in SFlash. Additional PPU will be used by the application (not stored in SFlash) for additional access restrictions specific to the application.

See the [Protection Context chapter on page 55](#) chapter for details on SMPU, PPU, and SWPU.

11.3.4.1 SMPU Configuration in SFlash

In SMPUs, address ranges will be chosen so that the access rights are well defined for both PC1 and PCx. The address ranges are also chosen such that the number of SMPUs required is minimized. One may require SMPUs with overlapping address ranges if the access rights for PC1 are different from the access rights for PCx. In that case, one may have to use the PC Match feature. For SMPUs, both the master and the slave registers are stored in SFlash.

SMPU15 and SMPU14 are configured during boot as follows:

- SMPU15 is configured to protect the first 2KB of SRAM such that only PC0 and PC1 can access it.
 - SMPU15 slave protection attribution ATT0 = 0x8A00037F

- ❑ SMPU15 master protection attribution ATT0 = 0x8700FF49
- SMPU14 is configured to protect system partition of SROM such that it is accessible only by PC0 and PC1. User partition is accessible by all PC.
 - ❑ SMPU14 slave protection attribution ATT0 = 0x8A00037F
 - ❑ SMPU14 master protection attribution ATT0 = 0x8700FF49

11.3.4.2 SWPU Configuration in SFlash

As stated earlier, the SWPU is broken into two parts, which are stored in SFlash. The first part implements the access restrictions related to PC1 and PCx. The second part is used by the application for additional access restrictions specific to the application. ROM/flash boot reads the two parts of SWPU from SFlash and stores them in RAM.

“write” (program/erase/erase suspend/erase resume) access protection for Flash and “read/write” access protection for eFuse are provided using Software Protection Units. These protection units are divided into three groups. The first group of protection units is called FLASH_WRITE_PU, the second group is called FUSE_READ_PU, and the third group is called FUSE_WRITE_PU. Write access protection for Flash is provided by FLASH_WRITE_PU. Read and write accesses to eFuse are provided by FUSE_READ_PU and FUSE_WRITE_PU, respectively.

The maximum number of FLASH_WRITE_PUs, FUSE_READ_PU, and FUSE_WRITE_PU are parameters initialized from SFlash. The value of this parameter is not expected to exceed 32 for FLASH_WRITE_PU and 8 for FUSE_READ_PU and FUSE_WRITE_PU.

A copy of the SWPU structures is stored in SFlash and during boot time the structures are read into RAM. The address range covered by each SWPU entry is fixed when the SWPU is stored in SFlash and cannot be updated in RAM. The integrity of SWPU entries in SFlash is checked by SECURE_HASH during secure boot.

11.3.4.3 PPU Configuration in SFlash

Read and write protection associated with each PPU can be categorized into one of four classes.

The write classes are defined as follows

- Class I – Both PC1 and PCx have write attribute = 0. For example, both PC1 and PCx do not have write access to the CPUSS_PROTECTION register. For the corresponding PPU, both PC1 and PCx must have the attributes “UW=0, PW=0, NS=1” for the master and slave registers.
- Class II – PC1 write attribute = 0 and PCx write attribute = 1. For example, PC1 does not have write access to CPUSS_AP_CTL, but PCx does. For the corresponding PPU, PC1 must have the attributes “UW=0, PW=0, NS=1” and PCx must have the attributes “UW=1, PW=1, NS=1” for the master and slave registers.
- Class III – PCx write attribute = 0. And PC1 write attribute = 1. For example, PCx does not have write access to EFUSE_MXS40.CTL, but PC1 does. For the corresponding PPU, PC1 must have the attributes “UW=1, PW=1, NS=1” and PCx must have the attributes “UW=0, PW=0, NS=1” for the master and slave registers.
- Class IV – PCx and PC1 have write attribute = 1. For example, both PC1 and PCx have write access to IPC_STRUCT1. For the corresponding PPU, both PC1 and PCx must have the attributes “UW=1, PW=1, NS=1” for the master and slave registers.

The read classes are defined as follows

- Class I – Both PC1 and PCx have read attribute = 0. For the corresponding PPU, both PC1 and PCx must have the attributes “UR=0, PR=0, NS=1” for the slave register.
- Class II – PC1 read attribute = 0 and PCx read attribute = 1. For the corresponding PPU, PC1 must have the attributes “UR=0, PR=0, NS=1” and PCx must have the attributes “UR=1, PR=1, NS=1” for the slave register.
- Class III – PCx read attribute = 0. And PC1 read attribute = 1. For the corresponding PPU, PC1 must have the attributes “UR=1, PR=1, NS=1” and PCx must have the attributes “UR=0, PR=0, NS=1” for the slave register.
- Class IV – PCx and PC1 have read attribute = 1. For the corresponding PPU, both PC1 and PCx must have the attributes “UR=1, PR=1, NS=1” for the slave register.

In genera, the read and write classification for each PPU is stored in SFlash only if at least one of them is class I or III. However, other classes may also be stored in SFlash. The storing is done in factory. SFlash has an entry that points to the protection settings. The ROM boot reads this classification and configures PPU accordingly. The following table shows the SFlash representation of the write/read access restrictions for PPUs. Refer to the device datasheet for information about the

protection unit PPU_ID for the corresponding PPU region. For example, neither PC1 or PCx can write, but they can both read the registers in the PERI_MS_PPU_FX_CPUSSE_BOOT region. For those PPUs whose classifications are not stored in SFlash, the ROM boot will configure the PPUs for both read and write to the default class IV.

Table 11-1. SFlash Representation of Write/Read Access Restrictions for Each PPU

Name of Fixed PPU	Access for PC > 0?(slave attributes)	Access for PC > 0?(Master attributes)
PERI_MS_PPU_FX_CRYPT0_BOOT	PC1 - read only PCx - read only	PC1 - read only PCx - read only
PERI_MS_PPU_FX_CPUSSE_BOOT	PC1 - read only PCx - read only	PC1 - read only PCx - read only
PERI_MS_PPU_FX_FLASHC_FlashMgmt	PC1 - full access PCx - No access	PC1 - read only PCx - read only
PERI_MS_PPU_FX_EFUSE_CTL	PC1 - full access PCx - read only	PC1 - read only PCx - read only
PERI_MS_PPU_FX_EFUSE_DATA	PC1 - full access PCx - No access	PC1 - read only PCx - read only
PERI_MS_PPU_FX_SRSS_SECURE	PC1 - read only PCx - read only	PC1 - read only PCx - read only
PERI_MS_PPU_FX_CRYPT0_MAIN	PC1 - full access PCx - full access	PC1 - full access PCx - full access
PERI_MS_PPU_FX_CRYPT0_CRYPT0	PC1 - full access PCx - full access	PC1 - full access PCx - full access
PERI_MS_PPU_FX_IPC_STRUCT0_IPC	PC - full access PCx - full access	PC1 - full access PCx - full access
PERI_MS_PPU_FX_IPC_STRUCT1_IPC	PC1 - full access PCx - full access	PC1 - full access PCx - full access
PERI_MS_PPU_FX_IPC_STRUCT2_IPC	PC1 - full access PCx - full access	PC1 - full access PCx - full access
PERI_MS_PPU_FX_FLASHC_DFT	PC1 - full access PCx - read only	PC1 - read only PCx - read only
PERI_MS_PPU_FX_BIST	PC1 - read only PCx - read only	PC1 - read only PCx - read only
PERI_MS_PPU_FX_CRYPT0_BUF	PC1 - full access PCx - full access	PC1 - full access PCx - full access

The following programmable PPUs are configured during boot. Note that for all programmable PPUs, PC other than PC0 can only modify SL_ATT or MS_ATT; SL_ADDR and SL_SIZE can be modified only in PC0. Therefore, all unused programmable PPUs, that is, PPUs that are not configured during the boot process, are not available to the user. See ["Protection Context" on page 51](#) for details.

- Programmable PPUs 0, 1, and 2 are used to protect the following area of eFuse such that it is not accessible to any PC other than PC0.

PPU ID	SL_ADDR	SL_SIZE
0	0x402c0840	4 bytes
1	0x402c0840	32 bytes
2	0x402c0860	8 bytes

- Programmable PPU 3 is used to protect the CRYPTO register (SL_ADDR = 0x40100000, SL_SIZE = 4 bytes). During boot it is enabled and access is provided to all PCs.

- Programmable PPU 5 is used to protect the CLK_TRIM_ILO0_CTL registers. This PPU is configured to allow write access to the registers for any PC except PC1 (SL_ADDR = 0x40263014, SL_SIZE= 4B).
- Programmable PPU 6 is used to protect the CLK_TRIM_ILO1_CTL registers. This PPU is configured to allow write access to the registers for any PC except PC1 (SL_ADDR = 0x40263220, SL_SIZE= 4B).
- Programmable PPU 7 is used to protect unused part of the CRYPTO register region (SL_ADDR = 0x40108000, SL_SIZE = 32KB). This region will be accessible only to PC0.
- Programmable PPU 8 is used to protect the part of flash controller register region (SL_ADDR = 0x4024f050, SL_SIZE = 16 bytes) in SECURE life cycle, such that they are accessible only to PC0.
- Programmable PPU 9 is used to allow only PC2 access to FLASHC_ECC_CTL registers of the DFT region, from 0x2a0 to 0x2bc offset as DFT region will be protected using a fixed PPU such that only PC0 has access.
- Programmable PPU 10 is used to provide access to EFUSE_SEQ_DEFAULT to all PCx (SL_ADDR = 0x402C0020, SL_SIZE = 4B)

Table 11-2. Programmable PPUs Modifiability Summary

Programmable PPU	Modifiable by PC0	Modifiable by PC1	Modifiable by PCx
PPU0	yes	yes	no
PPU1	yes	no	no
PPU2	yes	no	no
PPU3	yes	yes	yes
PPU5	yes	no	yes
PPU6	yes	no	yes
PPU7	yes	no	no
PPU8	yes	no	no
PPU9	yes	no	no ^a
PPU10	yes	yes	yes

a. Programmable PPU9 can be modified in PC0 and PC2; it cannot be modified in other PCs.

11.3.4.4 Boot Protection Settings in SFlash

Figure 11-1 shows how the protection settings are stored in SFlash.

- Object Size – Size of boot protection object in bytes.
- N_SMPU – Number of SMPU structures (starting from SMPU15) stored in this object.
For example, N_SMPU = 4 indicates SMPU15, SMPU14, SMPU13, and SMPU12 are configured.
- SMPU15 – Contains SMPU region address and SMPU region attributes.
- N_PPU – Number of PPU structures stored in this object.
- PPU_ID, PPU Config defines a PPU – PPU_ID is the PPU number (2 bytes) and the PPU Config is described using 1 byte (4 bits for write class and 4 bits for read class).
- N_FLASH_WRITE_PU – number of FLASH_WRITE_PUs stored in this object. It is followed by the contents of the FLASH_WRITE_PUs.
- FLASH_WRITE_PU – Data structure of FLASH_WRITE_PU.
- N_FUSE_READ_PU – number of FUSE_READ_PUs stored in this object. It is followed by the contents of the FUSE_READ_PUs.
- FUSE_READ_PU – Data structure of FUSE_READ_PU.
- N_FUSE_WRITE_PU – number of FUSE_WRITE_PUs stored in this object. It is followed by the contents of the FUSE_WRITE_PUs.
- FUSE_WRITE_PU – Data structure of FUSE_WRITE_PU.


Figure 11-1. Boot Protection Settings in Flash

...
FUSE_WRITE_PU (16B)
N_FUSE_WRITE_PU (4B)
...
FUSE_READ_PU (16B)
N_FUSE_READ_PU (4B)
...
FLASH_WRITE_PU (16B)
N_FLASH_WRITE_PU (4B)
...
PPU Config. (1B)
PPU_ID (2B)
N_PPU (4B)
...
SMPU15 (16B)
N_SMPU (4B)
Object Size (4B)

11.3.5 Debug and Test Access Restrictions

Depending on the protection state (NORMAL, SECURE, or DEAD), the ROM boot process will enforce access restrictions on the debug access port (SWD/JTAG). See “eFuse Bits” on page 142 for access restriction field bit map. The ROM boot gets access enable bits that allow Debug Access Port (AP) for CM0+, CM7_0/CM7_1, and system from eFuse or SFlash.

Figure 11-2. Debug Access Restrictions Structures



- Three separate structures exist (two for SECURE and DEAD in eFuse and one for NORMAL in supervisory flash). All three structures have the same layout.
- NORMAL access restrictions are stored in SFlash and they can be updated unlike the access restrictions in SECURE and DEAD protection states.

11.3.6 SWD/JTAG Initialization

When access restrictions prohibit use of the SWD/JTAG interface, the boot process does not access or change the SWD/JTAG pins in any way.

The customer firmware can, at any time, change the configuration of the SWD/JTAG pins to another mode, peripheral, configuration, or purpose.

To allow debugging of such applications, a 'listen window' is provided before starting customer firmware. The boot flash process will connect and enable the JTAG/SWD interface and wait for a specified time before starting application firmware. It is expected that application firmware checks the CPUSS_DP_STATUS.SWJ_CONNECTED bit and repurposes the pins only when no SWD/JTAG connection is available.

11.3.7 Waking up from Hibernate

Waking up from Hibernate will result in system boot. The integrity checks on the SFlash trim values and SWD/JTAG connection delay (Listen Window implemented by flash boot) is skipped when waking from hibernate.


11.3.8 Disable Watchdog Timer

The ROM boot code will disable the watchdog timer (WDT) if eFuse.WDT_DISABLE bit is set.

11.3.9 ROM Boot Flow Chart

Figure 11-3 shows the ROM boot flow chart.

Figure 11-3. ROM Boot Flow Chart



11.4 MMIO Registers and eFuse Used by ROM Boot

11.4.1 MMIO Registers

Table 11-3. MMIO Registers used by ROM Boot

Register	Name	Description
CPUSS		
CPUSS_IDENTITY	Identity register	Register that can be used to determine if the ROM boot process is executing on CM0+ or on CM7.
CPUSS_CM0_PC0_HANDLER	CM0+ protection context 0 handler	Register that holds location of NMI vector.
CPUSS_CM0_VECTOR_TABLE_BASE	CM0+ vector table base register	Register that holds location of the vector table (SP and reset exception vector address are provided at offset 0x0 and 0x4) for the M0+ boot image to be used after CPU reset. Typically, this is the location of the Cortex-M vector table in flash. This value must be set before issuing an M0+ CPU reset.
CPUSS_CM7_0/CM7_1_VECTOR_TABLE_BASE	CM7_0/CM7_1 vector table base register	Register that holds location of the vector table (SP and reset exception vector address are provided at offset 0x0 and 0x4) for the CM7_0/CM7_1 boot image. Typically, this is the location of the Cortex-M vector table in flash. This value must be set before releasing the CM7_0/CM7_1 from reset.
CPUSS_PROTECTION	Protection status register	Register that holds the current protection state. Can only be written to a different value according to the state diagram.
CPUSS_WOUNDING	Wounding register	Register that indicates the amount of accessible flash and RAM array in this device.
CPUSS_AP_CTL	Access port control register	Register that disables/enables any usage of CM0+, CM7_0/CM7_1, and system access port interface.
CPUSS_DP_STATUS	Debug port status register	Register that indicates whether a SWD/JTAG connection is established.
SRSS		
TST_MODE	Test mode control register	Register that indicates device is in a test mode. Setting this bit allows programming the flash before the control transfer to application.
PPU/SMPU/MPU/SWPU		Configures read/write access protection and flash program/erase protection. See the Protection Unit chapter on page 49 for details.
IPC		Configures inter-process communication; used to implement system call interface. See the Inter-Processor Communication chapter on page 48 for details.

11.4.2 eFuse Bits

Table 11-4. eFuse Bits Used by BootROM

Name	Bits	Description
SECURE Access Restrictions		
AP_CTL_CM0_DISABLE	1:0	Indicates that this device does not allow access to the CM0+ access port. 00 - Enable M0-AP 01 - Disable M0-AP 1x - Permanently Disable M0-AP
AP_CTL_CM7_DISABLE	3:2	Indicates that this device does not allow access to the CM7 access port. 00 - Enable CM7-AP 01 - Disable CM7-AP 1x - Permanently Disable CM7-AP
AP_CTL_SYS_DISABLE	5:4	Indicates that this device does not allow access to the system access port. 00 - Enable SYS-AP 01 - Disable SYS-AP 1x - Permanently Disable SYS-AP
SYS_AP_MPU_ENABLE	6	Indicates that the MPU on the system access port must be programmed and locked according to the settings in the next six fields.
DIRECT_EXECUTE_DISABLE	7	Disables DirectExecute system call functionality (implemented in software). 0: DirectExecute API execution is allowed 1: DirectExecute API execution is not allowed
FLASH_ALLOWED	10:8	This field indicates what portion of main flash is accessible through the system access port. Only a portion of flash starting at the bottom of the area is exposed. Encoding is as follows: 0: Entire region 1: 7/8th 2: 3/4th 3: 1/2 4: 1/4th 5: 1/8th 6: 1/16th 7: Nothing
SRAM_ALLOWED	13:11	This field indicates what portion of SRAM is accessible through the system access port. Only a portion of SRAM starting at the bottom of the area is exposed. Encoding is the same as FLASH_ALLOWED.
WORK_FLASH_ALLOWED	15:14	This field indicates what portion of work flash is accessible through the system access port. Only a portion of work flash starting at the bottom of the area is exposed. Encoding is as follows: 0: entire region 1: 1/2 2: 1/4th 3: Nothing

Table 11-4. eFuse Bits Used by BootROM

Name	Bits	Description
SFLASH_ALLOWED	17:16	This field indicates what portion of supervisory flash is accessible through the system access port. Only a portion of supervisory flash starting at the bottom of the area is exposed. Encoding is as follows: 0: Entire region 1: 1/2 2: 1/4th 3: Nothing
MMIO_ALLOWED	19:18	This field indicates what portion of the MMIO region is accessible through the system access port. Encoding is as follows: 0: All MMIO registers 1: Only IPC MMIO registers accessible (system calls) 2, 3: No MMIO access
SMIF_XIP_ENABLE	20	This field indicates what portion of SMIF_XIP is accessible through the system access port. Encoding is as follows: 0: Entire region 1: Nothing
DEAD Access Restrictions		
<Same as SECURE Access Restrictions>		The structure is identical to the one above but used when entering DEAD mode. It assumes that this structure is more restrictive than SECURE.
Critical Object Hash		
FACTORY_HASH		SHA-256 (upper 128 bits) that covers objects in TOC Part1. It is checked before transitioning to SECURE_WITH_DEBUG or SECURE.
SECURE_HASH		SHA-256 that covers the flash boot image and other objects in TOC Part1 and Part2. Flash boot code is not started unless this value is correct.
SECURE_HASH_ZEROES		The number of bits that are '0' (fuses that are not blown) in the SHA-256. This guarantees that when a HASH is programmed, it cannot be changed into another valid HASH value.

12. Interrupts



Traveo II supports interrupts and exceptions on both Cortex-M7 and Cortex-M0+ cores. Interrupts refer to events generated by peripherals external to the CPU such as timers, serial communication block, and port pin signals. Exceptions refer to events generated by the CPU such as memory access faults and internal system timer events. Both interrupts and exceptions result in the current program flow being stopped and the exception handler or interrupt service routine (ISR) being executed by the CPU. Both Cortex-M7 and Cortex-M0+ cores provide their own unified exception vector table for both interrupt handlers/ISR and exception handlers.

12.1 Features

Traveo II platform supports the following interrupt features:

- Supports up to 1023¹ system interrupts
 - Eight Cortex-M7 external interrupts and eight Cortex-M7 internal (software only) interrupts. The CPU supports up to 240 interrupts, but only sixteen interrupts are used by the Traveo II interrupt infrastructure. The eight external CPU interrupts support DeepSleep (WIC) functionality.
 - Eight Cortex-M0+ external interrupts and eight Cortex-M0+ internal (software only) interrupts. The CPU supports up to 32 interrupts, but only sixteen interrupts are used by the Traveo II interrupt infrastructure. The eight external CPU interrupts support DeepSleep (WIC) functionality.
 - All the available system interrupt sources are usable in Active power mode and can wake up from Sleep power mode
 - A subset of available system interrupt sources capable of waking the device from DeepSleep power mode
 - Four system interrupts can be mapped to each of the CPU NMI
- Nested vectored interrupt controller (NVIC) integrated with each CPU core, yielding low interrupt latency
- Wakeup interrupt controller (WIC) enabling interrupt detection (CPU wakeup) in DeepSleep power mode
- Vector table may be placed in either flash or SRAM
- Configurable priority levels (eight levels for Cortex-M7 and four levels for Cortex-M0+) for each interrupt
- Level-triggered interrupt signals

1. For the list of system interrupts supported by the device variants, refer to [12.5 Interrupt Sources](#).

12.2 How It Works

Figure 12-1. Traveo II Interrupts Block Diagram




Figure 12-1 shows the interrupt architecture in Traveo II. The 'N' system interrupts of Traveo II are processed by the NVIC of the individual cores. The Traveo II interrupt architecture uses eight CPU interrupts IRQ[7:0] out of the available CPU interrupts for each core. In the CM7 and CM0+ cores, the system interrupt source connection to a particular IRQn of the core is configurable and any of the 'N' system interrupts can be mapped to any of the IRQ[7:0] of each core. This ensures that all the system interrupts can be mapped onto any CPU interrupt simultaneously. The system interrupt to CPU interrupt mapping is independent for both CPUs. Refer to [12.5 Interrupt Sources](#) for more details about the system interrupt to CPU interrupt mapping. The NVIC enables/disables individual interrupt IRQs, priority resolution, and communication with the CPU core. Other exceptions such as NMI and HardFaults are not shown in Figure 12-1 because they are part of CPU core generated events, unlike interrupts, which are generated by peripherals external to the CPU.

In addition to the NVIC, Traveo II supports a wakeup interrupt controller (WIC) and interrupt synchronizer block. The WIC provides detection of DeepSleep interrupts in the DeepSleep CPU power mode. Each CPU can individually

be in DeepSleep; the device is said to be in DeepSleep only when all the CPUs are in DeepSleep. Refer to the [Device Power Modes chapter on page 186](#) for more details about the DeepSleep power mode. Each CPU has independent WIC settings; that is, the interrupts capable of waking up the CPU is configurable independent of the other. However, the device exits DeepSleep mode (System Wakeup signal in Figure 12-1) as soon as one CPU wakes up. For the list of system interrupts capable of waking up the CPU from DeepSleep power mode, refer to [12.5 Interrupt Sources](#). The interrupt synchronizer block synchronizes the interrupts to the CPU clock frequency as the peripheral interrupts can be asynchronous to the CPU clock frequency.


12.3 Interrupts and Exceptions – Operation

12.3.1 Interrupt/Exception Handling

The following sequence of events occurs when an interrupt or exception event is triggered:

1. Assuming that all the interrupt and exception signals are initially low (idle or inactive state) and the processor is executing the main code, a rising edge on any one of the signals is registered by the NVIC, if the interrupt or exception is enabled to be serviced by the CPU. The signal is now in a pending state waiting to be serviced by the CPU.
2. On detecting the signal from the NVIC, the CPU stores its current context by pushing the contents of the CPU registers onto the stack.
3. The CPU also receives the exception number of the triggered interrupt from the NVIC. All interrupts and exceptions have a unique exception number, as given in [Table 12-1](#) and [Table 12-2](#). By using this exception number, the CPU fetches the address of the specific exception handler from the vector table.
4. The CPU then branches to this address and executes the exception handler that follows.
5. Upon completion of the exception handler, the CPU registers are restored to their original state using stack pop operations; the CPU resumes the main code execution.

Figure 12-2. Interrupt Handling When Triggered



When the NVIC receives an interrupt request while another interrupt is being serviced or receives multiple interrupt requests at the same time, it evaluates the priority of all these interrupts, sending the exception number of the highest priority interrupt to the CPU. Thus, a higher priority interrupt can block the execution of a lower priority ISR at any time.

Exceptions are handled in the same way as interrupts. Each exception event has a unique exception number, which is used by the CPU to execute the appropriate exception handler.

Note: Because multiple system interrupts can be mapped on to the eight CPU interrupts (IRQ[7:0]), identification of system interrupts that triggered the CPU interrupt should be done in the CPU interrupt handler. This is described in [12.5 Interrupt Sources](#).

12.3.2 Level Interrupts

The CM0+ and CM7_0/CM7_1 NVICs support only level signals on the interrupt lines (IRQn). Pulse interrupts are not supported by Traveo II.

Figure 12-3. Level Interrupts




Figure 12-3 shows the working of level interrupts. Assuming the interrupt signal is initially inactive (logic low), the following sequence of events explains the handling of level interrupts.

On a rising edge event of the interrupt signal, the NVIC registers the interrupt request. The interrupt is now in the pending state, which means the interrupt requests have not yet been serviced by the CPU.

The NVIC then sends the exception number along with the interrupt request signal to the CPU. When the CPU starts executing the ISR, the pending state of the interrupt is cleared.

If the interrupt signal is still high after completing the ISR execution, it will be pending and the ISR is executed again. Figure 12-3 illustrates this for level triggered interrupts, where the ISR is executed as long as the interrupt signal is high.

12.3.3 Exception Vector Table

The exception vector tables (Table 12-1 and Table 12-2) store the entry point addresses for all exception handlers in Cortex-M0+ and Cortex-M7 cores. The CPU fetches the appropriate address based on the exception number.

Table 12-1. Cortex-M0+ Exception Vector Table

Exception Number	Exception	Exception Priority	Vector Address
–	Initial Stack Pointer Value	Not applicable (NA)	Start_Address = 0x0000 or CM0P_SCS_VTOR ^a
1	Reset	–3, the highest priority	Start_Address + 0x04
2	Non Maskable Interrupt (NMI)	–2	Start_Address + 0x08
3	HardFault	–1	Start_Address + 0x0C
4-10	Reserved	NA	Start_Address + 0x10 to Start_Address + 0x28
11	Supervisory Call (SVCall)	Configurable (0 - 3)	Start_Address + 0x2C
12-13	Reserved	NA	Start_Address + 0x30 to Start_Address + 0x34
14	PendSupervisory (PendSV)	Configurable (0 - 3)	Start_Address + 0x38
15	System Timer (SysTick)	Configurable (0 - 3)	Start_Address + 0x3C
16	External Interrupt (IRQ0)	Configurable (0 - 3)	Start_Address + 0x40
...	...	...	...
23	External Interrupt (IRQ7)	Configurable (0 - 3)	Start_Address + 0x5C
24	Internal (software only) Interrupt (IRQ8)	Configurable (0 - 3)	Start_Address + 0x60
...	...	...	...
31	Internal (software only) Interrupt (IRQ15)	Configurable (0 - 3)	Start_Address + 0x7C

a. Start Address = 0x0000 on reset and is later modified by user code by updating the CM0P_SCS_VTOR register.

Note: Internal interrupts IRQ8–IRQ15 are not connected to any peripheral and can be triggered by software only

Table 12-2. Cortex-M7 Exception Vector Table

Exception Number	Exception	Exception Priority	Vector Address
–	Initial stack pointer value	–	Start_Address = 0x0000 or CM7_0/ CM7_1_SCS_VTOR ^a
1	Reset	–3, highest priority	Start_Address + 0x04
2	Non Maskable Interrupt (NMI)	–2	Start_Address + 0x08
3	HardFault	–1	Start_Address + 0x0C
4	Memory management fault	Configurable (0 – 7)	Start_Address + 0x10
5	Bus fault	Configurable (0 – 7)	Start_Address + 0x14
6	Usage fault	Configurable (0 – 7)	Start_Address + 0x18
7–10	Reserved	–	–
11	Supervisory call (SVCall)	Configurable (0 – 7)	Start_Address + 0x2C
12–13	Reserved	–	–
14	Pend Supervisory (PendSV)	Configurable (0 – 7)	Start_Address + 0x38
15	System Tick timer (SysTick)	Configurable (0 – 7)	Start_Address + 0x3C
16	External interrupt (IRQ0)	Configurable (0 – 7)	Start_Address + 0x40
....			
23	External interrupt (IRQ7)	Configurable (0 – 7)	Start_Address + 0x5C
24	Internal (software only) Interrupt (IRQ8)	Configurable (0 – 7)	Start_Address + 0x60
....			
31	Internal (software only) Interrupt (IRQ15)	Configurable (0 – 7)	Start_Address + 0x7C

a. Start Address = 0x0000 on reset and is later modified by user code by updating the CM7_0/CM7_1_SCS_VTOR register.

Note: Internal interrupts IRQ8–IRQ15 are not connected to any peripheral and can be triggered by software only

In [Table 12-1](#) and [Table 12-2](#), the first word (four bytes) is not marked as exception number zero. This is because the first word in the exception table is used to initialize the main stack pointer (MSP) value on device reset; it is not considered as an exception. In Traveo II, both the vector tables can be configured to be located either in flash memory or SRAM. The vector table offset register (VTOR) present as part of Cortex-M0+ and Cortex-M7 system control space registers configures the vector table offset from the base address (0x00000000). The CM0P_SCS_VTOR register sets the vector offset address for the CM0+ core and CM7_0/CM7_1_SCS_VTOR sets the offset for the CM7_0/CM7_1 core. The VTOR value determines whether the vector table is in flash memory or SRAM. Refer to the device specific datasheet for the address region of flash and SRAM memories. Note that the VTOR registers can be updated only in privilege CPU mode; refer to the [Chip Operational Modes chapter on page 160](#) for details. The advantage of moving the vector table to SRAM is that the exception handler addresses can be dynamically changed by modifying the SRAM vector table contents. However, the nonvolatile flash memory vector table must be modified by a flash memory write.

The exception sources (exception numbers 1 to 15) are explained in [12.4 Exception Sources](#). The exceptions marked as Reserved in [Table 12-1](#) are not used, although they have addresses reserved for them in the vector table. The interrupt sources (exception numbers 16 to 23) are explained in [12.5 Interrupt Sources](#).

12.4 Exception Sources

This section explains the different exception sources listed in [Table 12-1](#) and [Table 12-2](#) (exception numbers 1 to 15).

12.4.1 Reset Exception

Device reset is treated as an exception in Traveo II. Reset exception is always enabled with a fixed priority of -3 , the highest priority exception, in both the cores. When the device boots up, only the Cortex-M0+ core is available. The CM0+ executes the ROM boot code and can enable Cortex-M7 core from the application code. The reset exception of the CM0+ is tied to the device reset or startup. When the CM0+ releases the CM7 reset, the CM7_0/CM7_1 reset exception is executed. A device reset can occur due to multiple reasons, such as POR, external reset signal on XRES_L pin, or watchdog reset. When the device is reset, the initial boot code for configuring the device is executed by the CM0+ from the SROM. The boot code and other data in SROM memory are programmed by Cypress, and are not read/write accessible to external users. After completing the SROM boot sequence, the CM0+ code execution jumps to flash memory. Flash memory address 0x10000004 (Exception#1 in [Table 12-1](#)) stores the location of the startup code in flash memory. The CPU starts executing code out of this address. Note that the reset exception address in the SRAM vector table will never be used because the device comes out of reset with the flash vector table selected. The register configuration to select the SRAM vector table can be done only as part of the startup code in flash after the reset is de-asserted. Note that the reset exception flow for CM7 is the same as CM0+. However, CM7 execution begins only after CM0+ de-asserts the CM7_0/CM7_1 reset. Refer to [Reset System on page 217](#) for details about Reset and start-up.

12.4.2 Non-Maskable Interrupt Exception

Non-maskable interrupt (NMI) is the highest priority exception next to reset. It is always enabled with a fixed priority of -2 . Both the cores have their own NMI exception. There are three ways to trigger an NMI exception in a CPU core:

- **NMI exception from a system interrupt:** Both CM0+ and CM7 provide an option to trigger an NMI exception using four of the available system interrupts for each core. The NMI exception triggered due to the interrupt will execute the NMI handler pointed to by the active vector table. The four CPUSS_CMx_NMI_CTL registers per CPU select the system interrupt sources that can trigger the NMI from hardware.
- **NMI exception by setting NMIPENDSET bit (user NMI exception):** An NMI exception can be triggered in software by setting the NMIPENDSET bit in the interrupt control state registers (CM0P_SCS_ICSR and CM7_0/CM7_1_SCS_ICSR). Setting this bit will execute the

NMI handler pointed to by the active vector table in the respective CPU cores.

12.4.3 HardFault Exception

Both CM0+ and CM7_0/CM7_1 cores support HardFault exception. HardFault is an always-enabled exception that occurs because of an error during normal or exception processing. HardFault has a fixed priority of -1 ; this means, it has higher priority over any exception with configurable priority. HardFault exception is a catch-all exception for different types of fault conditions, which include executing an undefined instruction and accessing an invalid memory addresses. The CPU does not provide fault status information to the HardFault exception handler, but it does permit the handler to perform an exception return and continue execution in cases where software has the ability to recover from the fault situation.

12.4.4 Memory Management Fault Exception

A memory management fault is an exception that occurs because of a memory protection-related fault. The fixed memory protection constraints determine this fault, for both instruction and data memory transactions. This fault is always used to abort instruction accesses to Execute Never (XN) memory regions. The memory management fault is only supported by the CM7_0/CM7_1 core. The priority of the exception is configurable from 0 (highest) to 7 (lowest).

12.4.5 Bus Fault Exception

A bus fault is an exception that occurs because of a memory-related fault for an instruction or data memory transaction. This can be from an error detected on a bus in the memory system. The bus fault is only supported by the CM7_0/CM7_1 core. The priority of the exception is configurable from 0 (highest) to 7 (lowest).

12.4.6 Usage Fault Exception

A usage fault is an exception that occurs because of a fault related to instruction execution. This includes:

- An undefined instruction
- An illegal unaligned access
- Invalid state on instruction execution
- An error on exception return

The following can cause a usage fault when the core is configured to report them:

- An unaligned address on word and halfword memory access
- Division by zero.

The usage fault is only supported by the CM7_0/CM7_1 core. The priority of the exception is configurable from 0 (highest) to 7 (lowest).

12.4.7 Supervisor Call (SVCall) Exception

Both CM0+ and CM7_0/CM7_1 cores support SVCall exception. Supervisor call (SVCall) is an always-enabled exception caused when the CPU executes the SVC instruction as part of the application code. Application software uses the SVC instruction to make a call to an underlying operating system and provide a service. This is known as a supervisor call. The SVC instruction enables the application to issue a supervisor call that requires privileged access to the system.

The priority of a SVCall exception can be configured to a value between 0 and 3 for CM0+ and 0 to 7 for CM7_0/CM7_1 core by writing to the bit fields PRI_11 of the System Handler Priority Register 2 (CM0P_SCS_SHPR2 and CM7_0/CM7_1_SCS_SHPR2). When the SVC instruction is executed, the SVCall exception enters the pending state and waits to be serviced by the CPU. The SVCALLPENDE bit in the System Handler Control and State Register (CM0P_SCS_SHCSR and CM7_0/CM7_1_SCS_SHCSR) can be used to check or modify the pending status of the SVCall exception.

12.4.8 PendSV Exception

Both CM0+ and CM7_0/CM7_1 cores support PendSV exception. PendSV is another supervisor call related exception similar to SVCall, normally being software-generated. PendSV is always enabled and its priority is configurable similar to SVCall. The PendSV exception is triggered by setting the PENDSVSET bit in the Interrupt Control State Register (CM0P_SCS_ICSR and CM7_0/CM7_1_SCS_ICSR). On setting this bit, the PendSV exception enters the pending state, and waits to be serviced by the CPU. The pending state of a PendSV exception can be cleared by setting the PENDSVCLR bit in the Interrupt Control State Register. The priority of a PendSV exception can be configured to a value between 0 and 3 for CM0+ and 0 to 7 for CM7_0/CM7_1 by writing to the bit fields PRI_14 of the System Handler Priority Register 3. See the [Armv6-M Architecture Reference Manual](#) for more details.

12.4.9 SysTick Exception

Both CM0+ and CM7_0/CM7_1 cores in Traveo II support a system timer, referred to as SysTick, as part of their internal architecture. SysTick provides a simple, 24-bit decrementing counter for various timekeeping purposes such as an RTOS tick timer, high-speed alarm timer, or simple counter. The SysTick timer can be configured to generate an interrupt when its count value reaches zero, which is referred to as SysTick exception. The exception is enabled by setting the TICKINT bit in the SysTick Control and Status Register (CM0P_SCS_SYST_CSR and CM7_0/

CM7_1_SCS_SYST_CSR). The priority of a SysTick exception can be configured to a value between 0 and 3 for CM0+ and 0 to 7 for CM7_0/CM7_1 by writing to the bit fields PRI_15 of the System Handler Priority Register 3 (SHPR3). The SysTick exception can always be generated in software by writing a one to the PENDSTSET bit in the Interrupt Control State Register (ICSR). Similarly, the pending state of the SysTick exception can be cleared by writing a one to the PENDSTCLR bit in the ICSR.

Note: The SysTick clock source can be configured through SYSTICK_CTL register in the CPUSS.

12.5 Interrupt Sources

The Traveo II family supports up to 1023 system interrupts from peripherals. However, the available system interrupts depend on the device variant. Check the device datasheet to know the list of system interrupt sources supported by the device variant.

The CM0+ CPU supports a maximum of 32 CPU interrupts (IRQ[31:0]) and the CM7_0/CM7_1 CPU supports a maximum of 240 CPU interrupts (IRQ[239:0]). To allow the support of up to 1023 system interrupts by the Cortex-M7 and M0+ CPUs, an interrupt reduction functionality is used. The interrupt reduction functionality allows each system interrupt to be mapped onto one out of the eight external CPU interrupts (IRQ[7:0]). Multiple system interrupts can be mapped on the same CPU interrupt. Therefore, an active CPU interrupt may indicate one or multiple active system interrupts.

The interrupt controller logic is independent for each CPU and each system interrupt has an associated CM0/CM7_0/7_1_SYSTEM_INT_CTL register:

- CM0/CM7_0/7_1_SYSTEM_INT_CTL.CPU_INT_VALID configures if the system interrupt is enabled for the CPU.
- CM0/CM7_0/7_1_SYSTEM_INT_CTL.CPU_INT_IDX[2:0] configures on which CPU interrupt the system interrupt is mapped.

Typically, the CPU uses different priority levels for the different CPU interrupts and will map system interrupts to CPU interrupts accordingly (all system interrupts that are mapped on the same CPU interrupt have the same priority). In addition to the eight (external) hardware CPU interrupts (IRQ[7:0]), eight (internal) software CPU interrupts are supported (IRQ[15:8]).

As a result of the reduction functionality, multiple system interrupts share a CPU interrupt handler as provided by the CPU's VTOR table. Each CPU interrupt has an associated CM0/CM7_0/7_1_INT_STATUS register:

- CM0+/CM7_0/CM7_1_INT_STATUS.SYSTEM_INT_VALID specifies if any system interrupt is active for the CPU interrupt.

- CM0+/CM7_0/
CM7_1_INT_STATUS.SYSTEM_INT_IDX[9:0] specifies the index (a number in the range [0, 1022]) of the lowest active system interrupt mapped to the corresponding CPU interrupt.

The CPU interrupt handler uses the SYSTEM_INT_IDX field to index a system interrupt lookup table and jumps to the system interrupt handler. The lookup table is typically located in one of the system memories. Note that this scenario introduces a two step approach: a CPU interrupt handler followed by a system interrupt handler. The following code illustrates the approach:

```
void CM7_0/CM7_1_CpuIntr0_Handler (void)
{
    uint32_t      system_int_idx;
    SystemIntr_Handler handler;

    if(CPUSS_CM7_0/CM7_1_INT_STATUS[0].SYSTEM_INT_VALID)
    {
        system_int_idx = CPUSS_CM7_0/CM7_1_INT_STATUS[0].SYSTEM_INT_IDX;
        handler = SystemIntr_Table[system_int_idx];
        handler(); // jump to system interrupt handler
    }
    else
    {
        // Triggered by software or because software cleared a peripheral interrupt flag
        // but did not clear the Pending flag at NVIC
    }
}
...
void CM7_0/CM7_1_CpuIntr7_Handler (void)
{
    uint32_t      system_int_idx;
    SystemIntr_Handler handler;

    if(CPUSS_CM7_0/CM7_1_INT_STATUS[7].SYSTEM_INT_VALID)
    {
        system_int_idx = CPUSS_CM7_0/CM7_1_INT_STATUS[7].SYSTEM_INT_IDX;
        handler = SystemIntr_Table[system_int_idx];
        handler(); // jump to system interrupt handler
    }
    else
    {
        // Triggered by software or because software cleared a peripheral interrupt flag
        // but did not clear the Pending flag at NVIC
    }
}
...
void CM7_0/CM7_1_SystemIntr0_Handler (void)
{
    // Clear the peripheral interrupt request flag by register write
    // Read back the register, to ensure completion of register write access
    // Handle system interrupt 0.
}
...
void CM7_0/CM7_1_SystemIntr1022_Handler (void)
{
    // Clear the peripheral interrupt request flag by register write
    // Read back the register, to ensure completion of register write access
    // Handle system interrupt 1022.
}
```


The system interrupts include standard interrupts from the on-chip peripherals such as TCPWM, serial communication block, CSD block, watchdog, ADC and so on. The interrupt generated is usually the logical OR of the different peripheral states. The peripheral interrupt status register should be read in the ISR to detect which condition generated the interrupt. These interrupts are usually level interrupts, which require that the peripheral interrupt status register be read in the ISR to clear the interrupt. If the interrupt status register is not read in the ISR, the interrupt will remain asserted and the ISR will be executed continuously. See the [I/O System chapter on page 246](#) for details on GPIO interrupts.

12.6 Exception Priority

Exception priority is useful for exception arbitration when there are multiple exceptions that need to be serviced by the CPU. Both CM7_0/CM7_1 and CM0+ cores in Traveo II provide flexibility in choosing priority values for different exceptions. All exceptions other than Reset, NMI, and HardFault can be assigned a configurable priority level. The Reset, NMI, and HardFault exceptions have a fixed priority of -3, -2, and -1 respectively. In Traveo II, lower priority numbers represent higher priorities. This means that the Reset, NMI, and HardFault exceptions have the highest priorities. The other exceptions can be assigned a configurable priority level between 0 and 3 for Cortex-M0+ and 0 to 7 for Cortex-M7.

Both CM0+ and CM7_0/CM7_1 support nested exceptions in which a higher priority exception can obstruct (interrupt) the currently active exception handler. This pre-emption does not happen if the incoming exception priority is the same as or lower than the active exception. The CPU resumes execution of the lower priority exception handler after servicing the higher priority exception. The CM0+ core in Traveo II allows nesting of up to four exceptions; the CM7_0/CM7_1 core allows up to eight exceptions. When the CPU receives two or more exceptions requests of the same priority, the lowest exception number is serviced first.

The registers to configure the priority of exception numbers 1 to 15 are explained in [12.4 Exception Sources](#).

The priority of the eight CM0+ and eight CM7_0/CM7_1 interrupts can be configured by writing to the respective Interrupt Priority registers (CM0P_SCS_IPR and CM7_0/CM7_1_SCS_NVIC_IPR). This is a group of eight (CM0+) and sixty (CM7_0/CM7_1) 32-bit registers with each register

storing the priority values of four interrupts, as given in [Table 12-3](#) and [Table 12-4](#)

Table 12-3. Interrupt Priority Register Bit Definitions for Cortex-M0+ (CM0P_SCS_IPR)

Bits	Name	Description
7:6	PRI_N0	Priority of interrupt number N.
15:14	PRI_N1	Priority of interrupt number N+1.
23:22	PRI_N2	Priority of interrupt number N+2.
31:30	PRI_N3	Priority of interrupt number N+3.

Table 12-4. Interrupt Priority Register Bit Definitions for Cortex-M7 (CM7_0/CM7_1_SCS_NVIC_IPR)

Bits	Name	Description
7:5	PRI_N0	Priority of interrupt number N
15:13	PRI_N1	Priority of interrupt number N+1
23:21	PRI_N2	Priority of interrupt number N+2
31:29	PRI_N3	Priority of interrupt number N+3

12.7 Enabling and Disabling Interrupts

The NVICs of both CM0+ and CM7_0/CM7_1 core provide registers to individually enable and disable the CPU interrupts in software. If an interrupt is not enabled, the NVIC will not process the interrupt requests on that interrupt line. The Interrupt Set-Enable Register (CM0P_SCS_ISER and CM7_0/CM7_1_SCS_NVIC_ISER) and the Interrupt Clear-Enable Register (CM0P_SCS_ICER and CM7_0/CM7_1_SCS_NVIC_ICER) are used to enable and disable the interrupts respectively. These registers are 32-bit wide and each bit corresponds to the same numbered interrupt in CM0+. For CM7_0/CM7_1 core, there are eight ISER/ICER registers. These registers can also be read in software to get the enable status of the interrupts. [Table 12-5](#) shows the register access properties for these two registers. Note that writing zero to these registers has no effect.

Table 12-5. Interrupt Enable/Disable Registers

Register	Operation	Bit Value	Comment
Interrupt Set Enable Register	Write	1	To enable the interrupt
		0	No effect
	Read	1	Interrupt is enabled
		0	Interrupt is disabled
Interrupt Clear Enable Register	Write	1	To disable the interrupt
		0	No effect
	Read	1	Interrupt is enabled
		0	Interrupt is disabled

The Interrupt Set-Enable Register (ISER) and Interrupt Clear-Enable Register (ICER) registers are applicable only for the interrupts. These registers cannot be used to enable or disable the exception numbers 1 to 15. The 15 exceptions have their own support for enabling and disabling, as explained in [12.4 Exception Sources](#).

The Priority Mask (PRIMASK) register in the CPUs (both CM0+ and CM7_0/CM7_1) can be used as a global exception enable register to mask all the configurable priority exceptions irrespective of whether they are enabled. Configurable priority exceptions include all the exceptions except Reset, NMI, and HardFault listed in [Table 12-1](#). When the PRIMASK.PM bit is set, none of the configurable priority exceptions can be serviced by the CPU, though they can be in the pending state waiting to be serviced by the CPU after the PRIMASK.PM bit is cleared.

12.8 Exception States

Each exception can be in one of the following states.

Table 12-6. Exception States

Exception State	Meaning
Inactive	The exception is not active and not pending. Either the exception is disabled or the enabled exception has not been triggered.
Pending	The exception request has been received by the CPU/NVIC and the exception is waiting to be serviced by the CPU.
Active	An exception that is being serviced by the CPU but whose exception handler execution is not yet complete. A high-priority exception can interrupt the execution of lower priority exception. In this case, both the exceptions are in the active state.
Active and Pending	The exception is being serviced by the processor and there is a pending request from the same source during its exception handler execution.

The Interrupt Control and State Register (CM0P_SCS_ICSR and CM7_0/CM7_1_SCS_ICSR) contains status bits describing the various exceptions states.

- The ICSR.VECTACTIVE bits store the exception number for the current executing exception. This value is zero if the CPU does not execute any exception handler (CPU is in thread mode). Note that the value in VECTACTIVE bit fields is the same as the value in bits [8:0] of the Interrupt Program Status Register (IPSR), which is also used to store the active exception number.
- The ICSR.VECTPENDING bits store the exception number of the highest priority pending exception. This value is zero if there are no pending exceptions.
- The ICSR.ISRPENDING bit indicates if a NVIC generated interrupt is in a pending state.

12.8.1 Pending Exceptions

When a peripheral generates an interrupt request signal to the NVIC or an exception event occurs, the corresponding exception enters the pending state. When the CPU starts executing the corresponding exception handler routine, the exception is changed from the pending state to the active state. The NVIC allows software pending of the eight (CM0+/CM7_0/CM7_1) interrupt lines by providing separate register bits to set and clear the pending states of the interrupts. The Interrupt Set-Pending registers (CM0P_SCS_ISPR and CM7_0/CM7_1_SCS_NVIC_ISPR) and the Interrupt Clear-Pending register (CM0P_SCS_ICPR and CM7_0/CM7_1_SCS_NVIC_ICPR) are used to set and clear the pending status of the interrupt lines. These registers are 32 bits wide, and each bit corresponds to the same numbered interrupt. [Table 12-7](#) shows the register access properties for these two registers. Note that writing zero to these registers has no effect.

Table 12-7. Interrupt Set Pending/Clear Pending Registers

Register	Operation	Bit Value	Comment
Interrupt Set-Pending Register (ISPR)	Write	1	To put an interrupt to pending state
		0	No effect
	Read	1	Interrupt is pending
		0	Interrupt is not pending
Interrupt Clear-Pending Register (ICPR)	Write	1	To clear a pending interrupt
		0	No effect
	Read	1	Interrupt is pending
		0	Interrupt is not pending

Setting the pending bit when the same bit is already set results in only one execution of the ISR. The pending bit can be updated regardless of whether the corresponding interrupt is enabled. If the interrupt is not enabled, the interrupt line will not move to the pending state until it is enabled by writing to the ISER.

Note that the ISPR and ICPR are used only for the peripheral interrupts. These registers cannot be used for pending the exception numbers 1 to 15. These 15 exceptions have their own support for pending, as explained in [12.4 Exception Sources](#).

12.9 Stack Usage for Exceptions

When the CPU executes the main code (in thread mode) and an exception request occurs, the CPU stores the state of its general-purpose registers in the stack. It then starts executing the corresponding exception handler (in handler mode). The CPU pushes the contents of the eight 32-bit internal registers into the stack. These registers are the Program and Status Register (PSR), ReturnAddress, Link

Register (LR or R14), R12, R3, R2, R1, and R0. Both Cortex-M7 and Cortex-M0+ has two stack pointers – MSP and PSP. Only one of the stack pointers can be active at a time. When in thread mode, the Active Stack Pointer bit in the Control register is used to define the current active stack pointer. When in handler mode, the MSP is always used as the stack pointer. The stack pointer always grows downwards and points to the address that has the last pushed data.

When the CPU is in thread mode and an exception request comes, the CPU uses the stack pointer defined in the control register to store the general-purpose register contents. After the stack push operations, the CPU enters handler mode to execute the exception handler. When another higher priority exception occurs while executing the current exception, the MSP is used for stack push/pop operations, because the CPU is already in handler mode. See the [CPU Subsystem \(CPUSS\) chapter on page 38](#) for details.

12.10 Interrupts and Low-Power Modes

Traveo II allows device (CPU) wakeup from low-power modes when certain peripheral interrupt requests are generated. The WIC block generates a wakeup signal that causes the CPU to enter Active mode when one or more wakeup sources generate an interrupt signal. After entering Active mode, the ISR of the peripheral interrupt is executed.

The Wait For Interrupt (WFI) or Wait For Event (WFE) instructions executed by the CPU triggers the transition into Sleep and DeepSleep modes. Only the WFI instruction is meant for waking up using interrupts. The WFE instruction puts the CPU to sleep based on the status of an event bit and wakes up from an event signal, typically sent by the other CPU. The sequence to enter the different low-power modes is detailed in the [Device Power Modes chapter on page 186](#). Device low-power modes have two categories of interrupt sources:

- Interrupt sources that are available in the Active, Sleep, and DeepSleep modes (see the [Device Power Modes chapter on page 186](#) for the available sources)
- Interrupt sources that are available only in the Active and Sleep modes

12.11 Exception – Initialization and Configuration

This section discusses the steps to initialize and configure exceptions in Traveo II.

1. Configuring the Exception Vector Table Location:

The first step in using exceptions is to configure the vector table location as required – either in flash memory

or SRAM. This configuration is described in [12.3.3 Exception Vector Table](#).

It is recommended that the vector table be available in SRAM if the application needs to change the vector addresses dynamically. If the table is located in flash, then a flash write operation is required to modify the vector table contents.

2. Configuring Individual Exceptions: The next step is to configure individual exceptions required in an application, as explained in earlier sections.

- a. Configure the exception or interrupt source; this includes setting up the interrupt generation conditions. The register configuration depends on the specific exception required. Refer to the respective peripheral chapter to know more about the interrupt configuration supported by them.
- b. Define the exception handler function and write the address of the function to the exception vector table. [Table 12-1](#) gives the exception vector table format; the exception handler address should be written to the appropriate exception number entry in the table.
- c. For system interrupts, define the system interrupt handler function, specify to which CPU interrupt the system interrupt is to be mapped in CM0/CM7_0/CM7_1_SYSTEM_INT_CTL.CPU_INT_IDX[2:0] and enable the system interrupt by setting CM0/CM7_0/CM7_1_SYSTEM_INT_CTL.CPU_INT_VALID. The CPU interrupt handler function should check the CM0/CM7_0/CM7_1_INT_STATUS.SYSTEM_INT_IDX[9:0] to determine the system interrupt that caused the interrupt and call the corresponding system interrupt handler function.
- d. Set up the exception priority, as explained in [12.6 Exception Priority](#).
- e. Enable the exception, as explained in [12.7 Enabling and Disabling Interrupts](#).

12.12 Registers

Table 12-8. Register List

Register	Name	Description
CPUSS_CM7_0/CM7_1_INTx_STATUS	CM7_0/CM7_1 CPU Interrupt Status Register	The CPUSS_CM7_0/CM7_1_INT0_STATUS – CPUSS_CM7_0/CM7_1_INT7_STATUS registers provide the lowest CM7_0/CM7_1-activated system interrupt index for the eight external CPU interrupts.
CPUSS_CM7_0/CM7_1_VECTOR_TABLE_BASE	CM7_0/CM7_1 Vector Table Base Register	Address of CM7_0/CM7_1 vector table. This register is used for CM7_0/CM7_1 warm and cold boot purposes: the CM0+ CPU initializes the CM7_0/CM7_1_VECTOR_TABLE_BASE register and the CM7_0/CM7_1 boot code uses the register to initialize the CM7_0/CM7_1 internal VTOR register.
CPUSS_CM7_0/CM7_1_NMI_CTLx	CM7_0/CM7_1 NMI Control Register	The CPUSS_CM7_0/CM7_1_NMI_CTL0 – CPUSS_CM7_0/CM7_1_NMI_CTL3 registers allow connecting four system interrupts to the NMI. The four selected system interrupts are logically OR'd into a single CM7_0/CM7_1 NMI input.
CPUSS_CM0_INTx_STATUS	CM0+ CPU Interrupt Status Register	The CPUSS_CM0_INT0_STATUS – CPUSS_CM0_INT7_STATUS registers provide the lowest CM0-activated system interrupt index for the eight external CPU interrupts.
CPUSS_CM0_VECTOR_TABLE_BASE	CM0+ Vector Table Base Register	Address of CM0+ vector table. This register is used for CM0+ warm boot purposes: the CM0+ warm boot code uses the register to initialize the CM0+ internal VTOR register.
CPUSS_CM0_NMI_CTLx	CM0+ NMI Control Register	The CPUSS_CM0_NMI_CTL0 – CPUSS_CM0_NMI_CTL3 registers allow connecting four system interrupts to the CM0+ NMI. The four selected system interrupts are logically OR'd into a single CM0+ NMI input.
CPUSS_CM0_SYSTEM_INT_CTLx	CM0+ System Interrupt Control Register	These registers are used to configure the mapping of system interrupt “x” to one of the eight external CM0+ CPU interrupts.
CPUSS_CM7_0/CM7_1_SYSTEM_INT_CTLx	CM7_0/CM7_1 System Interrupt Control Register	These registers are used to configure the mapping of system interrupt “x” to one of the eight external CM7_0/CM7_1 CPU interrupts.
CM0P_SCS_ISER ^a	Cortex-M0+ Interrupt Set-Enable Register	The CM0P_SCS_ISER enables CM0+ external and internal (software only) interrupts, and shows which interrupts are enabled.
CM0P_SCS_ICER ^a	Cortex-M0+ Interrupt Clear-Enable Register	The CM0P_SCS_ICER disables CM0+ external and internal (software only) interrupts, and shows which interrupts are enabled.
CM0P_SCS_ISPR ^a	Cortex-M0+ Interrupt Set-Pending Register	The CM0P_SCS_ISPR forces CM0+ external and internal (software only) interrupts into the pending state, and shows which interrupts are pending.
CM0P_SCS_ICPR ^a	Cortex-M0+ Interrupt Clear-Pending Register	The CM0P_SCS_ICPR removes the pending state from CM0+ external and internal (software only) interrupts, and shows which interrupts are pending.
CM0P_SCS_IPRx ^a	Cortex-M0+ Interrupt Priority Registers	The CM0P_SCS_IPR registers allow to configure the priority for the CM0+ external and internal (software only) interrupts.
CM0P_SCS_ICSR ^a	Cortex-M0+ Interrupt Control and State Register	The CM0P_SCS_ICSR provides: ■ a set-pending bit for the non-maskable interrupt (NMI) exception ■ set-pending and clear-pending bits for the PendSV and SysTick exceptions The register also indicates: ■ the number of the highest priority pending exception

Table 12-8. Register List

Register	Name	Description
CM0P_SCS_VTOR ^a	Cortex-M0+ Vector Table Offset Register	The CM0P_SCS_VTOR indicates the offset of the CM0+ vector table base address from memory address 0x00000000.
CM0P_SCS_AIRCR ^a	Cortex-M0+ Application Interrupt and Reset Control Register	The CM0P_SCS_AIRCR provides endian status for CM0+ data accesses and reset control of the system.
CM0P_SCS_SHPR2 ^a	Cortex-M0+ System Handler Priority Register 2	The CM0P_SCS_SHPR2 allows to configure the priority for SVCcall exception.
CM0P_SCS_SHPR3 ^a	Cortex-M0+ System Handler Priority Register 3	The CM0P_SCS_SHPR3 allows to configure the priority for SysTick and PendSV exceptions.
CM0P_SCS_SHCSR ^a	Cortex-M0+ System Handler Control and State Register	The CM0P_SCS_SHCSR controls and provides the active and pending status of CM0+ exceptions.
CM7_0/CM7_1_SCS_NVIC_ISERx ^b	CM7_0/CM7_1 Interrupt Set-Enable Registers	The CM7_0/CM7_1_SCS_NVIC_ISER registers enable CM7_0/CM7_1 external and internal (software only) interrupts, and show which interrupts are enabled.
CM7_0/CM7_1_SCS_NVIC_ICERx ^b	CM7_0/CM7_1 Interrupt Clear-Enable Registers	The CM7_0/CM7_1_SCS_NVIC_ICER registers disable CM7_0/CM7_1 external and internal (software only) interrupts, and show which interrupts are enabled.
CM7_0/CM7_1_SCS_NVIC_ISPRx ^b	CM7_0/CM7_1 Interrupt Set-Pending Registers	The CM7_0/CM7_1_SCS_NVIC_ISPR registers force CM7_0/CM7_1 external and internal (software only) interrupts into the pending state, and show which interrupts are pending.
CM7_0/CM7_1_SCS_NVIC_ICPRx ^b	CM7_0/CM7_1 Interrupt Clear-Pending Registers	The CM7_0/CM7_1_SCS_NVIC_ICPR registers remove the pending state from CM7_0/CM7_1 external and internal (software only) interrupts, and shows which interrupts are pending.
CM7_0/CM7_1_SCS_NVIC_IABRx ^b	CM7_0/CM7_1 Interrupt Active Bit Registers	The CM7_0/CM7_1_SCS_NVIC_IABR registers indicate which CM7_0/CM7_1 external and internal (software only) interrupts are active.
CM7_0/CM7_1_SCS_NVIC_IPRx ^b	CM7_0/CM7_1 Interrupt Priority Registers	The CM7_0/CM7_1_SCS_NVIC_IPR registers allow to configure the priority for the CM7_0/CM7_1 external and internal (software only) interrupts.
CM7_0/CM7_1_SCS_STIR ^b	CM7_0/CM7_1 Software Triggered Interrupt Register	The CM7_0/CM7_1_SCS_STIR allows to generate CM7_0/CM7_1 external and internal (software only) interrupts from software. This register has the same function as CM7_0/CM7_1_SCS_NVIC_ISPR except that STIR can be configured to allow access by unprivileged software.
CM7_0/CM7_1_SCS_ICSR ^b	CM7_0/CM7_1 Interrupt Control State Register	The CM7_0/CM7_1_SCS_ICSR provides: ■ a set-pending bit for the NMI exception ■ set-pending and clear-pending bits for the PendSV and SysTick exceptions This register also indicates: ■ the exception number of the exception being processed ■ whether there are preempted active exceptions ■ the exception number of the highest priority pending exception ■ if any interrupts are pending
CM7_0/CM7_1_SCS_VTOR ^b	CM7_0/CM7_1 Vector Table Offset Register	The CM7_0/CM7_1_SCS_VTOR indicates the offset of the CM7_0/CM7_1 vector table base address from memory address 0x00000000.
CM7_0/CM7_1_SCS_AIRCR ^b	CM7_0/CM7_1 Application Interrupt and Reset Control Register	The CM7_0/CM7_1_SCS_AIRCR provides priority grouping control for the exception model, endian status for data accesses of CM7_0/CM7_1, and reset control of the system.
CM7_0/CM7_1_SCS_SHPR1 ^b	CM7_0/CM7_1 System Handler Priority Register 1	The CM7_0/CM7_1_SCS_SHPR1 allows to configure the priority for UsageFault, BusFault, and MemManage exceptions.

Table 12-8. Register List

Register	Name	Description
CM7_0/CM7_1_SCS_SHPR2 ^b	CM7_0/CM7_1 System Handler Priority Register 2	The CM7_0/CM7_1_SCS_SHPR2 allows to configure the priority for SVC call exception.
CM7_0/CM7_1_SCS_SHPR3 ^b	CM7_0/CM7_1 System Handler Priority Register 3	The CM7_0/CM7_1_SCS_SHPR3 allows to configure the priority for SysTick and PendSV exceptions.
CM7_0/CM7_1_SCS_SHCSR ^b	CM7_0/CM7_1 System Handler Control and State Register	The CM7_0/CM7_1_SCS_SHCSR enables the system handlers, and indicates: ■ the pending status of the BusFault, MemManage fault, and SVC exceptions ■ the active status of the system handlers
CPUSS_SYSTICK_CTL	SysTick Timer Control Register	The CPUSS_SYSTICK_CTL register allows to configure the SysTick timer clock source, specify the clock source precision, and the number of clock source cycles that make up 10 ms. Note: If an external clock source is configured using this register, the external clock frequency must be less than the CPU internal clock frequency.

a. Refer to the Arm Cortex-M0+ TRM for details about this register.

b. Refer to the Arm Cortex-M7 TRM for details about this register.

13. Device Security



Traveo II offers several features to protect user designs from unauthorized access or copying. Selecting a secure life-cycle stage, enabling memory and peripheral protection, configuring flash write and eFuse read/write protection, and using hardware-based cryptography can provide a high level of security.

13.1 Features

The Traveo II provides the following device security features:


- Nonvolatile and irreversible life-cycle stages that can limit program and debug access.
- Memory protection units (MPU), shared memory protection units (SMPU), and peripheral protection units (PPU) provide memory and peripheral protection, such as preventing unauthorized reading of sensitive data.
- Software protection units (SWPU) that define flash write (or erase) permissions and eFuse read and write permissions.
- A cryptographic function block that provides hardware-based encryption and decryption of data and code.

13.2 How It Works

13.2.1 Life-Cycle Stages

Traveo II devices have configurable, nonvolatile life-cycle stages. Life-cycle stages follow a strict, irreversible progression governed by invoking system management APIs that will change the one-time programmable (OTP) eFuse settings accordingly.

Figure 13-1. Traveo II Life Cycle Stage Transitions



Traveo II supports the following life-cycle stages:

- **NORMAL_PROVISIONED** – Customers receive parts in this stage.
- **SECURE** – You can secure the device in this stage after the application has been created and tested. A secure device will boot only when there is a successful authentication of its flash boot code and application code. Access restrictions in SECURE mode are controlled by eFuse settings.
- **SECURE_WITH_DEBUG** – This is similar to the SECURE life-cycle stage, except with NORMAL access restrictions applied to enable debugging, even if authentication fails. Devices that are in this stage are only used by developers and testers.
- **RMA** – Devices can be brought into this stage so that Cypress can perform a failure analysis. Sensitive data should be erased before transitioning to this life-cycle stage. The boot process will set access restrictions such that only the “Open for RMA” system management API call can be executed from outside; this requires a part-specific certificate provided by the customer.
- **CORRUPTED** (not shown in state diagram) – This stage is entered in case an error is detected when the boot process tries to determine the current life-cycle stage.

The current nonvolatile life-cycle stage as well as the volatile protection state can be retrieved with the SiliconID system call. The protection state is also available in the CPUSS_PROTECTION register.

The following table shows the mapping of life-cycle stages to the protection states.

Table 13-1. Life-cycle Stage Mapping

Life-cycle Stage	Protection State
VIRGIN ^a	VIRGIN
SORT ^a	
PROVISIONED ^a	
RMA	
NORMAL ^a	NORMAL
NORMAL_PROVISIONED	
SECURE	SECURE
SECURE_WITH_DEBUG	
Any of the above stages (on certain conditions)	DEAD
CORRUPTED	

a. These life-cycle stages are not applicable for final samples.

13.2.2 Memory and Peripheral Protection

The MPU, SMPU, and PPU can be used to restrict access to memory (RAM and flash) or peripheral address space. This can prevent unauthorized code or bus masters from reading/writing sensitive address areas.

For more details, see the [Protection Unit chapter on page 53](#).

13.2.3 Flash Write and eFuse Read/Write Protection

Traveo II devices include software protection units (SWPU), which define permissions for flash writing (or erasing) and eFuse reading and writing. This feature prevents malicious or inadvertent modification of flash or eFuse, or reading of sensitive eFuse data. In addition, unauthorized changes to the application are detected by the secure boot operation.

For more details, see the [Protection Unit chapter on page 53](#).

13.2.4 Hardware-based Cryptography

Traveo II has a cryptographic block (Crypto) that provides hardware implementation and acceleration of cryptographic functions. It implements symmetric key encryption and decryption, hashing, message authentication, random number generation (pseudo and true), cyclic redundancy checking, and hardware acceleration of asymmetric cryptography. See the [Cryptography Block chapter on page 480](#).

14. Chip Operational Modes



Traveo II is capable of executing firmware in four different modes. These modes dictate execution from different locations in flash and ROM, with different levels of hardware privileges. Only three of these modes are used in end-applications; debug mode is used exclusively to debug designs during firmware development. This chapter gives an overview of the Traveo II operational modes. The device power modes are explained in the [Device Power Modes chapter on page 186](#). These modes are independent of privileged and unprivileged access levels of Arm Cortex core.

The operational modes in Traveo II are:

- Boot
- User
- Trusted
- Debug

14.1 Boot

In the Boot mode the device is configured by instructions hard-coded in the device ROM and from supervisory flash. This mode is entered after the end of a reset, provided no debug-acquire sequence is received by the device. Boot mode is a privileged mode; interrupts are disabled so that the boot firmware can set up the device for operation without being interrupted. During boot mode, hardware trim settings are loaded from flash to guarantee proper operation during power-up. After executing ROM boot code, supervisory flash boot code execution begins after flash boot authentication. ROM boot is the root of trust as it is immutable. Flash boot is more flexible and can be modified during the VIRGIN life cycle. However it is treated as an extension of ROM boot because it is authenticated by the ROM boot. After both ROM boot and flash boot, the device enters user mode and code execution from user flash begins. See the [BootROM chapter on page 133](#) for the details of ROM boot and flash boot.

14.2 User

In the User mode normal user firmware from flash is executed. User mode cannot execute code from ROM. The boot process transfers control to this mode after it has completed its tasks. Then the user application starts from the default user application address. Both privileged and unprivileged access levels of Arm Cortex core can be executed in the user mode.

14.3 Trusted

Trusted mode allows execution of special subroutines that are stored in the device ROM. These subroutines cannot be modified by the user and are used to execute proprietary code that is not meant to be interrupted or observed. Debugging is not allowed in the trusted mode.

This mode is entered from user mode by executing the system call (ROM API code). Trusted ROM code can be executed only when the master is in protection context 1. Only the CM0+ (secure CPU) can attain protection context 1 upon a trusted interrupt handler entry. See the [Device Security chapter on page 158](#) for more details on protection contexts. Exit from this mode returns the device to user mode.

14.4 Debug

Debug mode allows observation of the Traveo II operational parameters. This mode is used to debug firmware during development. The debug mode is entered when a debugger connects to the device during the acquire time window, which occurs during device reset. Debug mode allows IDEs to debug the firmware. This mode is available only on devices whose access restriction settings allow debugging. For NORMAL protection state, access restrictions settings are stored in the supervisory flash (SFlash). For DEAD and SECURE states, it is stored in eFuse. For more details on protection states, see the [Device Security chapter on page 158](#). For more details on the debug interface, see the [Program and Debug Interface chapter on page 684](#).

15. Fault Subsystem



The fault subsystem contains information about faults that occur in the system. The subsystem can cause a reset, give a pulse indication, or trigger another peripheral. The Traveo II platform uses a centralized fault report structure. The centralized nature allows for a system-wide, consistent handling of faults, which simplifies software development as follows:

- Only a single fault interrupt handler is required
- The fault report structure provides the fault source and additional fault-specific information from a single set of Memory Mapped Input/Output (MMIO) registers; that is, no iterative search is required for the fault source and fault information
- All pending faults are available from a single set of MMIO registers

The fault subsystem captures faults related to, but not limited to:

- MPU/SMPU/PPU protection violations
- Peripheral-specific errors
- Memory controller specific errors, such as SRAM controller ECC errors, flash controller “read-while-program”, and ECC errors
- Processor tightly-coupled memory (TCM) ECC errors
- Timeout errors

Note that some of the above faults also result in errors on the bus infrastructure. These faults are communicated in two ways:


- As a bus error to the master of the faulting bus transfer
- As a fault in a fault report structure. This fault can be communicated as a fault interrupt to any processor in the system. This allows fault handling on a processor that is not the master of the faulting bus transfer. It is useful for faults that cause the master of the faulting transfer to become unresponsive or unreliable

The fault subsystem only captures faults. It does not take any action to correct it.

15.1 Fault Report Structure

Figure 15-1 gives an overview of the fault report structure.

Figure 15-1. Fault Reporting Structure



Refer to the device datasheet for information about the number of fault report structures (FAULT_NR) supported. Each structure has a dedicated set of control and status registers, and captures a single fault. The captured fault information includes:

- A validity bit field that indicates a fault is captured (FAULT_STRUCTx_STATUS.VALID).
- A fault index that identifies the fault source (FAULT_STRUCTx_STATUS.IDX).
- Additional fault information describing fault specifics (FAULT_STRUCTx_DATAy). This additional information is fault type-specific. Most fault types use only a few of the FAULT_STRUCTx_DATAy registers. For example, an MPU protection violation provides information on the violating bus address, bus master identifier, and bus access control information in only two FAULT_STRUCTx_DATAy registers.

In addition to the captured fault information, each fault report structure supports a signaling interface to notify the rest of the system of the captured fault. This interface supports the following:

- A fault interrupt (interrupts_fault). This interrupt is supported by the platform interrupt registers: FAULT_STRUCTx_INTR, FAULT_STRUCTx_INTR_SET, FAULT_STRUCTx_INTR_MASK, and FAULT_STRUCTx_INTR_MASKED. Only a single interrupt cause is present: FAULT (indicating that a fault is detected). The FAULT_STRUCTx_INTR_MASK register provides a mask/enable for the cause. The interrupt cause is set to '1' when a fault is captured.
- A trigger (FAULT_TR_OUT[FAULT_NR]). An enabled trigger is activated (generating a two-cycle '1' pulse) when FAULT_STRUCTx_STATUS.VALID is set to '1'. The trigger is enabled by


FAULT_STRUCTx_CTL.TR_EN. The trigger can be connected to a DMA controller, for example, which can transfer captured fault information from the fault report structure to memory and can clear the FAULT_STRUCTx_STATUS.VALID field. For failure analysis, a memory location that is retained during warm/soft reset is desirable.

- An output signal (FAULT_OUT_x, x = 0, 1, 2, 3). An enabled output signal is active/'1' when FAULT_STRUCTx_STATUS.VALID is '1'. The output signal is enabled by FAULT_STRUCTx_CTL.OUT_EN. It can be used to communicate non-recoverable faults, for example, to off-chip components (possibly resulting in a device reset).
- A fault reset request (fault_reset_req). An enabled request is active/'1' when FAULT_STRUCTx_STATUS.VALID is '1'. The request is enabled by FAULT_STRUCTx_CTL.RESET_REQ_EN. The reset request feeds into the logic that generates a warm/soft reset.

The four different signaling interfaces provided have their own 'enable' functionality. Each enabled interface is activated when FAULT_STRUCTx_STATUS.VALID is '1'.

As the system resources subsystem (SRSS) has a single fault_reset_req input signal, the individual fault_reset_req[i] signals are combined (logical OR'd) into a single fault_reset_req signal.

Figure 15-2. Fault Reset



A central structure, shared by all fault report structures, keeps track of all pending faults in the system. The `FAULT_STRUCTx_PENDINGy` registers reflect which of the fault sources are pending. These registers provide a dedicated pending bit for up to 96 fault sources. The `FAULT_STRUCTx_PENDINGy` registers are mirrored in each of the fault report structures. The fault source numbering scheme follows the numbering scheme of `FAULT_STRUCTx_STATUS.IDX`.

The fault sources corresponding to a pending bit (which is set) are the ones that are not yet captured by any of the fault structures. When a pending fault is captured by a fault structure, the associated pending bit is cleared to '0'. Each fault report structure is selective in the faults it captures. `FAULT_STRUCTx_MASKy` reflect which pending fault source is captured by a fault structure. These faults are referred to as "enabled" faults. The `FAULT_STRUCTx_MASKy` registers are unique to each fault structure. This allows for the following:

- One fault report structure is used to capture recoverable faults and another is used to capture non-recoverable faults. The former can be used to generate a fault interrupt and the latter can be used to activate a chip output signal or a reset request.
- Two fault report structures are used to capture the same faults. This first fault is captured by the structure with the lower index (for example, fault structure 0) and the second fault is captured by the structure with the higher index (for example, fault structure 1).
Note: `FAULT_STRUCTx_STATUS.VALID` bits are different for each of the fault structures. As an example, consider that the MCWDT lower threshold is linked to Fault Structure#0 and higher threshold is linked to Fault Structure#1.
 Fault Structure#0 occurs first to give a warning; then, Fault Structure#1 occurs to trigger a reset.

A fault structure only captures "enabled" faults when `FAULT_STRUCTx_STATUS.VALID` is '0'.

When a fault is captured, the hardware sets `FAULT_STRUCTx_STATUS.VALID` to '1'. In addition, the hardware clears the associated pending bit to '0'. When a fault structure is processed, the software (if the fault is processed by an interrupt handler) or a DMA transfer (if a triggered DMA transfer copied the captured fault information) should clear `FAULT_STRUCTx_STATUS.VALID` to '0'.

Note that fault capturing does not consider `FAULT_STRUCTx_INTR.FAULT`:

- Fault capturing is only conditioned by `FAULT_STRUCTx_STATUS.VALID` being '0'.
- If an interrupt handler is used to process the fault structure, software should clear `FAULT_STRUCTx_INTR.FAULT` to '0'.

15.2 Fault and Reset

As mentioned, a captured fault may result in a warm/soft reset. This type of reset brings regular MMIO registers to their default/reset state. This is not acceptable for the registers that capture fault information; for failure analysis, fault information should be retained during a warm/soft reset. Therefore, the `FAULT_STRUCTx_STATUS` and `FAULT_STRUCTx_DATAy` registers are connected to a cold reset. This illustrates another benefit of centralized fault report structures: only the centralized structure is connected to a cold reset. The multiple fault sources that are scattered throughout the system can use the regular reset, as a copy of the fault information is captured by the fault structure.

Note: When the fault is configured to trigger reset, then debugging of the configured fault structure is not possible.

15.3 Fault and Power Modes

The fault report structure functionality is available only in Active/Sleep power modes (it is an Active functionality):

- DeepSleep fault sources are not supported. These fault sources require dedicated solutions.
- The interfaces between the active fault sources and the centralized fault report structures is reset in DeepSleep power mode. Note that the fault information is retained.

As the fault report structure is an active functionality, pending faults (in the `FAULT_STRUCTx_PENDINGy` registers) are not retained when transitioning to DeepSleep power mode. This is acceptable, because the fault source itself is an active functionality.

For fault assignments, refer to the device specific datasheet.

15.4 Register List

Table 15-1. Fault Subsystem Register List

Symbol	Name	Description
FAULT_STRUCTx_CTL	Fault Control	This register is used to enable or disable the output trigger, I/O output signal, and reset request when a fault occurs.
FAULT_STRUCTx_STATUS	Fault Status	This register provides the fault source index and validity of data in the fault data registers.
FAULT_STRUCTx_DATAy	Fault Data	The data registers capture fault information.
FAULT_STRUCTx_PENDING0	Fault Pending 0	The FAULT_STRUCTx_PENDINGy registers specify pending (not captured) fault sources. The fault source for which data is captured in FAULT_STRUCTx_DATAy registers and is validated by FAULT_STRUCTx_STATUS.VALID and identified by FAULT_STRUCTx_STATUS.IDX is not included in this list of pending fault sources. When a fault source is captured, its corresponding bit field in FAULT_STRUCTx_PENDINGy is set to 0.
FAULT_STRUCTx_PENDING1	Fault Pending 1	
FAULT_STRUCTx_PENDING2	Fault Pending 2	
FAULT_STRUCTx_MASK0	Fault Mask 0	The FAULT_STRUCTx_MASKy registers specify “enables” for fault sources. Only “enabled” fault sources will be captured by this fault structure (and result in FAULT_STRUCTx_STATUS.VALID and FAULT_STRUCTx_INTR.FAULT being set to 1). When a fault source is captured, its corresponding bit field in FAULT_STRUCTx_PENDINGy is set to 0.
FAULT_STRUCTx_MASK1	Fault Mask 1	
FAULT_STRUCTx_MASK2	Fault Mask 2	
FAULT_STRUCTx_INTR	Interrupt	This register sets the register bit when an enabled pending fault source is captured.
FAULT_STRUCTx_INTR_SET	Interrupt Set	This register sets the corresponding bits in the interrupt request register.
FAULT_STRUCTx_INTR_MASK	Interrupt Mask	Mask for interrupt request register.
FAULT_STRUCTx_INTR_MASKED	Interrupt Masked	Bitwise AND of interrupt request and mask registers.

Note: In FAULT_STRUCTx, 'x' signifies the fault structure instance and 'y' in FAULT_STRUCTx_PENDINGy/MASKy varies from 0 through 2 and FAULT_STRUCTx_DATAy varies 0 through 3.

Section C: System Resources Subsystem (SRSS)




This section encompasses the following chapters:

- [Power Supply and Monitoring chapter on page 167](#)
- [Device Power Modes chapter on page 186](#)
- [Clocking System chapter on page 197](#)
- [Reset System chapter on page 217](#)
- [Watchdog Timer chapter on page 222](#)
- [Real-Time Clock chapter on page 239](#)

Top Level Architecture

System-Wide Resources Block Diagram



16. Power Supply and Monitoring



The Traveo II family supports multiple power supply rails – V_{DD} , V_{DDA} , V_{CCD} , and multiple V_{DDIO} rails. It integrates multiple regulators to power the blocks within the device in various power modes.

Traveo II devices support power-on-reset (POR), brownout detection (BOD), over-voltage detection (OVD), over-current detection (OCD), and low-voltage detection (LVD) circuit for power supply monitoring and failure detection purposes. The low-voltage detection circuit can also be used as a high-voltage detection (HVD) circuit.

- POR provides a reset pulse during the V_{DD} initial power ramp.
- BOD on V_{DD} or V_{CCD} generates a reset if V_{DD} or V_{CCD} voltage dips below the threshold voltage.
- BOD on V_{DDA} can generate a reset or a fault if V_{DDA} voltage dips below the threshold voltage.
- OVD on V_{DD} or V_{CCD} generates a reset if V_{DD} or V_{CCD} voltage goes above the threshold voltage.
- OVD on V_{DDA} can generate a reset or a fault if V_{DDA} voltage goes above the threshold voltage.
- OCD generates a reset if the load current of a regulator is over the regulator limit.
- LVD (HVD) can generate an interrupt or a fault whenever V_{DD} voltage crosses the threshold in the configured direction.

16.1 Features

The features of the Traveo II power supply subsystem are as follows:


- V_{DD} power supply voltage range of 2.7 V to 5.5 V.
- Core supply rail (V_{CCD}).
- Independent multiple power supply rails (V_{DD} , V_{DDA} , V_{CCD} , and multiple V_{DDIO} rails) for Traveo II core peripherals.
- Multiple on-chip regulators.
 - Active regulator to power the MCU in Active/Sleep mode in case of low current consumption
 - DeepSleep regulator to power peripherals operating in DeepSleep mode
 - High-current regulator to support higher current load by using an external pass transistor or by controlling for an external power management integrated circuit (PMIC) or low-dropout (LDO) regulator
- Low-voltage (V_{CCD}) and high-voltage (V_{DD} and V_{DDA}) BOD circuits are available in all power modes except Hibernate and XRES modes.
- Low-voltage (V_{CCD}) and high-voltage (V_{DD} and V_{DDA}) OVD circuits are available in all power modes except Hibernate and XRES modes.
- Two LVD circuits to monitor V_{DD} for falling detection (LVD), rising detection (HVD), or both in all power modes except Hibernate and XRES modes.
- OCD circuit to monitor V_{CCD} current in all power modes except Hibernate and XRES modes. OCD is not monitored for PMIC.

16.2 Power Supply

The regulators and supply pins/rails, shown in [Figure 16-1](#), power various blocks inside the device. The availability of various supply rails/pins for an application will depend on the device package selected. See the device datasheet for details.

All the core regulators draw their input power from the V_{DD} supply pin. V_{CCD} supply is used to power all active domains and DeepSleep domains. From V_{CCD} , there are power domain switches that allow disabling active circuitry while leaving DeepSleep circuitry connected to V_{CCD} . The Hibernate domain does not implement any regulators and the peripherals available in that domain such as ILO operate directly from V_{DD} .

Figure 16-1. Power System Block Diagram



The I/O cells operate from V_{DDA} or multiple V_{DDIO} rails depending on the port they are located. V_{CCD} supply is used to drive logic inside the I/O cells from core peripherals. To know which I/Os operate from which supply see the [I/O System chapter on page 246](#).

16.2.1 Core Regulators

The device includes the following core regulators to power peripherals and blocks in various power modes.

Note that in Hibernate mode, all regulators are OFF and V_{CCD} is not driven. Hibernate related logic operates from V_{DDP} directly. For details, see the [Device Power Modes chapter on page 186](#).

Active Regulator

The device includes a linear LDO regulator to power the Active and Sleep mode peripherals. This regulator generates the core voltage (V_{CCD}) from V_{DDP} during Active and Sleep modes. It is operational during device start up, when servicing interrupts that do not require the M7 CPUs to be fully active, or when the device is active at slow clock frequencies. The high-current regulator powers the MCU in case of high current consumption in Active/Sleep mode, which cannot be supplied by the Active regulator.

DeepSleep Regulator

In addition to the Active regulator, the device includes a DeepSleep regulator, which generates the core voltage (V_{CCD}) during DeepSleep mode. The primary differences from the Active regulator are that it has lower drive capability and consumes much less current.

High-Current Regulator

The High-current regulator controller supports higher load currents than the Active regulator using an external pass transistor, or it can control an external PMIC or LDO.

Traveo II starts with the Active regulator, before switching to the high-current regulator.

16.2.2 Power Pins and Rails

[Table 16-1](#) lists all the power supply pin names available in the device. The PCB must short all identically named pins externally with low-impedance connections (that is, either connect to a plane or use a wide top layer route between pins). And, all ground pins must be at same potential. For details, see the device datasheet.

Table 16-1. Supply Pins

Supply Pin	Ground Pin	Power Supply Voltage Range	Description
V_{DDP}	V_{SSD}	2.7 V to 5.5 V	Digital and I/O supply
V_{CCD}	V_{SSD}	1.1 V to 1.2 V	Core supply
V_{DDA}	V_{SSA}	2.7 V to 5.5 V	Analog supply, $V_{DDA} = V_{DDIO\_2}$
V_{DDIO_1}	V_{SSD}	2.7 V to 5.5 V	I/O supply
V_{DDIO_2}	V_{SSD}	2.7 V to 5.5 V	I/O supply
V_{DDIO_3}	V_{SSIO_3}	2.7 V to 3.6 V	I/O supply
V_{DDIO_4}	V_{SSIO_4}	2.7 V to 3.6 V	I/O supply

16.2.3 Power Sequencing Requirements

V_{DDP} , multiple V_{DDIO} , and V_{DDA} do not have any sequencing limitation and can establish in any order. These supplies except V_{DDA} and V_{DDIO_2} are independent in voltage level. See the device datasheet for details of device operating conditions.

There are operating limits if a supply is not present:

- The part will not boot unless V_{DDP} is present
- V_{DDA} must be equal to V_{DDIO_2}
- A BOD can be configured by software to reset the part when V_{DDA} is not present

16.2.4 Power Supply Sources

Traveo II offers power supply options that support a wide range of application voltages and requirements. The recommended V_{DDP} voltage range is 2.7 V to 5.5 V. If the application voltage is in this range, then Traveo II (V_{DDP}) can be interfaced with any power supply voltage in the range of 2.7 V to 5.5 V. Other supply rails and pins such as V_{DDA} and multiple V_{DDIO} rails exist independent of V_{DDP} and V_{CCD} . See the device datasheet for details of device operating conditions.

16.2.5 Usage of High-Current Regulator Controller

The high-current regulator controller (REGHC) is initially disabled and must be enabled with the REGHC_EN bit [31] of PWR_REGHC_CTL2 register. The Active regulator supports the chip current until the high-current regulator controller is configured, enabled, operating, and ready. The current consumption of the full device must stay within the operation conditions of the Active regulator

The high-current regulator controller has two configurations. One is using an external pass transistor configuration (Figure 16-2) and the other is using an external PMIC or LDO device (Figure 16-4 or Figure 16-5) for driving much higher load currents. Table 16-2 lists the high-current regulator controller pins for each configuration. EXT_PS_CTL0/1/2 pins are shared with the GPIO function. Therefore, the GPIO function must be disabled (high-impedance mode with input buffer disabled same as the default state) by software before REGHC is enabled.

Table 16-2. High-Current Regulator Controller Pins

Pin Name	External Transistor mode		External PMIC mode	
	Direction	Description	Direction	Description
DRV_VOUT	OUT	Gate of the pass transistor	Unused	–
EXT_PS_CTL0	IN	Positive terminal of the current sense resistor	IN	Power good input from PMIC
EXT_PS_CTL1	IN	Negative terminal of the current sense resistor	OUT	Enable output for PMIC
EXT_PS_CTL2	Unused	–	OUT (optional)	Reset threshold adjustment for some PMIC

The high-current regulator controller transitions when:

- Software switches between the Active regulator and the high-current regulator controller.
- The high-current regulator controller is disabled by hardware for transitions to OFF and XRES states. Operation resumes with the Active regulator when the reset condition is removed. Software can change back to the high-current regulator controller after the device reboots.
- A low-voltage reset can leave the device in an unintended state that requires software recovery. See [Transitioning from Active/Sleep to Reset on page 178](#).
- The following options are supported for DeepSleep:
 - Hardware changes to the Active regulator before entering DeepSleep using PWR_REGHC_CTL4.REGHC_PMIC_DPSLP = 0. The device wakes up from DeepSleep using the Active regulator. After waking from DeepSleep, hardware changes back to the high-current regulator controller.
Follow the sequence in [DeepSleep Entry/Exit on page 180](#).
 - PMIC is configured to operate during DeepSleep using PWR_REGHC_CTL4.REGHC_PMIC_DPSLP = 1. Hardware does not change the power system settings when entering or exiting DeepSleep.
- The high-current regulator controller is disabled by hardware for transitions to Hibernate mode. When using an external PMIC, the high-current regulator controller tristates the PMIC enable output and a pull resistor on the PCB or within the PMIC disables the PMIC. The device wakes from Hibernate using the Active regulator. After wakeup, software can reconfigure the high-current regulator controller and change back to it.

Figure 16-2. External Pass Transistor Configuration




Table 16-3. External Transistor Requirement Specification (Type: NPN)

Parameter	Symbol	Min Value	Unit	Conditions
Static forward current transfer ratio	H_{FE}	100	-	$I_C = 1\text{ A}$, $V_{CE} = 1\text{ V}$
Collector-emitter voltage	V_{CEO}	10	V	-
Transition frequency	f_T	100	MHz	-
Collector current	I_C	1	A	-
Collector power dissipation	P_{CD}	2	W	-

Do not exceed the rated temperature; it is necessary to estimate the junction temperature.

Loss of the pass transistor under the operating conditions is calculated using the following equation.

Equation 1

$$P_{Loss} = (V_{DD_{max}} - V_{CCD_{min}}) \times I_{VCCD_{max}}$$

Where:

P_{Loss} : Loss of pass transistor (W)

$V_{DD_{max}}$: Maximum V_{DD} voltage (V)

$V_{CCD_{min}}$: Minimum V_{CCD} voltage (V)

$I_{VCCD_{max}}$: Maximum V_{CCD} load current (V)

The maximum power dissipation of the pass transistor must be greater than or equal to the above P_{Loss} .

The junction temperature of the pass transistor is calculated using the following equation:

Equation 2

$$T_J = P_{Loss} \times \theta_{JA} + T_A$$

Where:

T_J : Pass transistor junction temperature (°C)

T_A : Ambient temperature (°C)

P_{Loss} : Loss by pass transistor (W)

θ_{JA} : Thermal resistance from junction to ambient temperature (°C/W)

T_J must be below the rated temperature.

Table 16-4 shows examples of maximum θ_{JA} .

Table 16-4. Maximum θ_{JA} Examples

$T_{A\ max}$ (°C)	V_{DD}	$P_{Loss\ max}$ (W)	θ_{JA} (°C/W)
+125	5-V power rail	2.64	< 9.5
+105			< 17.0
+85			< 24.6
+125	3.3-V power rail	1.50	< 16.7
+105			< 30.0
+85			< 43.3

Conditions: V_{DD} maximum voltage = 5.5 V/3.6 V, V_{CCD} minimum voltage = 1.10 V, V_{CCD} maximum load current = 0.6 A

16.2.5.1 Transitioning from Active Regulator to High-Current Regulator Controller with External Transistor

To set up the high-current regulator controller with external transistor:

1. Confirm if the circuit board has populated the necessary external components for the high-current regulator controller, including a compatible transistor, capacitor, and current sense resistor.

This sequence is provided to customers as part of the ConfigureRegulator API:

2. Configure the high-current regulator controller. Writes in same registers can be done in same cycle:
 - a. Write `PWR_REGHC_CTL.REGHC_MODE = 0` to configure the external transistor mode.
 - b. Write `PWR_REGHC_CTL.REGHC_TRANS_USE_OCD = 1`, depending on whether a current monitoring resistor is integrated on the PCB. Current monitoring is recommended.
 - c. Write `PWR_REGHC_CTL.REGHC_VADJ` to the required feedback setting. Add the signed offset in `SFLASH_SRSS_PWR_OFFSET.REGHC_TRANS_VADJ_OFFSET` to account for die to die variation.
 - d. Write `PWR_REGHC_CTL.REGHC_CONFIGURED = 1`.

After the high-current regulator controller is set up, it can be enabled without writing the setup again.

This sequence is handled by the SwitchOverRegulators API. Call the API with blocking = 1 to ensure the transition occurs in the proper order:

- e. Execute the system call (`LoadRegulatorTrims`) to change internal regulator trims.
- f. Write `PWR_REGHC_CTL2.REGHC_EN = 1`.
- g. Wait until `PWR_REGHC_STATUS.REGHC_SEQ_BUSY = 0` and `PWR_REGHC_STATUS.REGHC_ENABLED = 1`. This should occur within 15 μ s.
- h. Again, execute the trim change system call.
- i. The device is now operating on the high-current regulator controller with external transistor. Additional current load can be enabled.

16.2.5.2 Transitioning from High-Current Regulator Controller with External Transistor to Active Regulator

To transition from the high-current regulator controller with external transistor to the Active regulator:


- j. Reduce the current consumption to within the Active regulator limit.

This part of the sequence is handled by the SwitchOverRegulators API. Call the API with blocking = 1 to ensure the transition occurs in the proper order.

- k. Execute the system call (LoadRegulatorTrims) to change internal regulator trims.
- l. Write `PWR_REGHC_CTL2.REGHC_EN = 0`.
- m. Wait until `PWR_REGHC_STATUS.REGHC_SEQ_BUSY = 0` and `PWR_REGHC_STATUS.REGHC_ENABLED = 0`. This should occur within 10 μ s.
- n. Again, execute the trim change system call.
- o. The device is now operating on the Active regulator.

Figure 16-3 shows the transitions between active regulator and external transistor.


Figure 16-3. Transitions between Active Regulator and External Transistor



1: Add the signed offset in `SFLASH SRSS_PWR_OFFSET.REGHC_TRANS_VADJ_OFFSET` to account for die to die variation.

Figure 16-4 shows how to connect to a PMIC that does not discharge its output.


Figure 16-4. External PMIC/LDO Configuration



- PMIC EN pin polarity is HIGH for enable. PMIC PG pin polarity is HIGH for power good.
- If EN pin of PMIC does not have the internal pull-down resistor, an external pull-down resistor must be placed to keep the PMIC disabled during power-on reset.
- See the Traveo II device datasheet for C_{S1} .
- Output voltage setting resistors (R1, R2) are needed according to the selected PMIC.

Figure 16-5 shows how a load switch can be used for compatibility with a PMIC that discharges its output when disabled.

Figure 16-5. External PMIC with Load Switch Configuration



- PMIC EN pin polarity is HIGH for enable. PMIC PG pin polarity is HIGH for power good.
- Output voltage setting resistors (R1, R2) and output capacitor (C1) depends on the selected PMIC.
- See the Traveo II device datasheet for C_{S1} .
- Gate signal of load switch is connected to EXT_PS_CTL1 pin of MCU.
A resistor to restrict current of load switch gate charge to less than 1 mA is put between EXT_PS_CTL1 and gate of the load switch.
- EN pin of PMIC is connected to GPIO of MCU. To maintain correct operation, use GPIO in the power domain that keeps the power during DeepSleep.

16.2.5.3 Transitioning from Active Regulator to High-Current Regulator Controller with External PMIC

To setup the high-current regulator controller with external PMIC:

1. Confirm if the circuit board has populated the necessary external components for the high-current regulator controller, including a compatible PMIC. When using an external PMIC, the PMIC is responsible for supervision of V_{CCD} .

This sequence is provided to customers as part of ConfigureRegulator API:

2. Configure the high-current regulator controller for PMIC operation. These fields can be written in same cycle in the PWR_REGHC_CTL register:
 - a. Write PWR_REGHC_CTL.REGHC_MODE = 1 to configure the PMIC mode.
 - b. Write PWR_REGHC_CTL.REGHC_PMIC_STATUS_INEN = 1, to enable the input path for PMIC status. Write PWR_REGHC_CTL.REGHC_PMIC_STATUS_POLARITY to the setting that indicates an error condition (depending on the polarity of the PMIC status output).
 - c. Write PWR_REGHC_CTL.REGHC_PMIC_CTL_POLARITY to the setting that enables the PMIC (depending on polarity of PMIC enable input).
 - d. Write PWR_REGHC_CTL.REGHC_VADJ to the required feedback setting for the chosen PMIC. Add the signed offset in SFLASH_SRSS_PWR_OFFSET.PMIC_VADJ_OFFSET to account for die-to-die variation.
 - e. Customer option: Configure PWR_REGHC_CTL.REGHC_PMIC_USE_LINREG = 1 to keep the internal active regulator enabled for its supply supervision capability. If this feature is not desired, write PWR_REGHC_CTL.REGHC_PMIC_USE_LINREG = 0.

- f. Customer option: Configure `PWR_REGHC_CTL.REGHC_PMIC_USE_RADJ = 1` and `PWR_REGHC_CTL.REGHC_PMIC_RADJ` to generate a reset threshold for the PMIC. If this feature is not needed, write `PWR_REGHC_CTL.REGHC_PMIC_USE_RADJ = 0`. Reset threshold adjustment is offered to reduce the need for external components. If it is used, they must be set such that the PMIC can robustly supply V_{CCD} .
- g. Customer option: Configure `PWR_REGHC_CTL.REGHC_PMIC_STATUS_WAIT` to give additional settling time after PMIC status input is initially correct, until the sequencer continues. This is written if PMIC needs more startup time.
3. PMIC configuration in the `PWR_REGHC_CTL4` register:
 - h. Customer option: Configure `REGHC_PMIC_DPSLP` to specify the PMIC behavior during DeepSleep. If using a PMIC with load switch, and the PMIC is configured to be always enabled on the PCB (not under programmable control), then configure `REGHC_PMIC_DPSLP = 1`.
 - i. Configure `REGHC_PMIC_VADJ_DIS = 1` to reduce power consumption. This is especially useful to reduce DeepSleep current, if `REGHC_PMIC_DPSLP = 1`.
4. Configure this set of fields in a separate write cycle from the previous set. This prevents a possible glitch on PMIC enable output, which can happen if the polarity and output enable are changed at the same time.
 - j. Write `PWR_REGHC_CTL.REGHC_PMIC_CTL_OUTEN = 1`
 - k. Write `PWR_REGHC_CTL.REGHC_CONFIGURED = 1`

After the high-current regulator controller is set up, it can be enabled without writing the setup again. If using a PMIC with load switch, ensure the PMIC is enabled and operating. (Set GPIO to '1' for enabling PMIC. Then, wait until `PWR_REGHC_STATUS.REGHC_PMIC_STATUS_OK = 1`.)

This part of the sequence is handled by the SwitchOverRegulators API. Call the API with `blocking = 1` to ensure the transition occurs in the proper order:

 - l. Execute the system call (`LoadRegulatorTrims`) to change the internal regulator trims.
 - m. Write `PWR_REGHC_CTL2.REGHC_EN = 1`.
 - n. If `blocking = 1`, wait until `PWR_REGHC_STATUS.REGHC_SEQ_BUSY = 0` and `PWR_REGHC_STATUS.REGHC_ENABLED = 1`. This delay depends strongly on the startup time of the PMIC, based on its status output and the value in `PWR_REGHC_CTL.REGHC_PMIC_STATUS_WAIT`.
 - o. Again, execute the trim change system call.
 - p. If `PWR_REGHC_CTL.REGHC_PMIC_USE_LINREG = 0` and `PWR_REGHC_CTL4.REGHC_PMIC_DPSLP = 1`, then it may be possible to disable the DeepSleep regulator by writing `PWR_CTL2.DPSLP_REG_DIS = 1`.
This bit must not be set if it is later intended to switch back to the Active regulator using the sequence in the next section. The API with `blocking = 1` case assumes it is never intended to switch back, and it disables the DeepSleep regulator for this register configuration. If the application wants the future ability to switch back to the Active regulator, it must call the API with `blocking = 0` and not write `PWR_CTL2.DPSLP_REG_DIS` (leave it 0).
 - q. The device is now operating on the high-current regulator controller with external PMIC. Additional current load can be enabled.

16.2.5.4 Transitioning from High-Current Regulator Controller with External PMIC to Active Regulator

Note: This sequence cannot be used if the DeepSleep regulator is disabled (`PWR_CTL2.DPSLP_REG_DIS = 1` at any time). The DeepSleep regulator cannot be re-enabled, and it is needed for the Active regulator to operate. If the application wants to use this sequence, do not disable the DeepSleep regulator in the transition to the PMIC.

To transition from the high-current regulator controller with external PMIC to the Active regulator:

- r. Reduce the current consumption to within the Active regulator limit.

This part of the sequence is handled by the SwitchOverRegulators API. If using a PMIC with load switch, call the API with `blocking = 0`. For all other cases, call the API with `blocking = 1` to ensure the transition completes in the proper order:

- s. Execute the system call (`LoadRegulatorTrims`) to change internal regulator trims.
- t. Write `PWR_REGHC_CTL2.REGHC_EN = 0`.
- u. If `blocking = 1`, wait until `PWR_REGHC_STATUS.REGHC_SEQ_BUSY = 0` and `PWR_REGHC_STATUS.PWR_REGHC_ENABLED = 0`. This delay depends on how long it takes for the external PMIC to deassert its power good signal (that is, until `PWR_REGHC_STATUS.REGHC_PMIC_STATUS_OK = 0`).


- v. Again, execute the trim change system call.
- w. The device is now operating on the Active regulator.

If using a PMIC with load switch and `PWR_REGHC_CTL4.REGHC_PMIC_DPSLP = 0`, wait until the load switch is fully off and then disable the PMIC so it deasserts its power good signal and `PWR_REGHC_STATUS.REGHC_PMIC_STATUS_OK = 0`. This allows the internal state machine to complete.

If the `SwitchOverRegulators` API was called with `blocking = 0`, set GPIO as '0' to disable PMIC after the GPIO that is shared with `EXT_PS_CTL1` is 0 (load switch turned off). Later, wait for the transition to complete (`PWR_REGHC_STATUS.REGHC_SEQ_BUSY = 0` and `PWR_REGHC_STATUS.REGHC_ENABLED = 0`).

Figure 16-6 shows an example of transitions between the active regulator and external PMIC, with `REGHC_PMIC_CTL_POLARITY = 1` and `REGHC_PMIC_STATUS_POLARITY = 0`.

Figure 16-6. Transitions between Active Regulator and External PMIC



- 1: Set `REGHC_PMIC_STATUS_POLARITY` depending on the polarity of the PMIC error status output. This bit affects the polarity of `EXT_PS_CTL0` pin.
- 2: Set `REGHC_PMIC_CTL_POLARITY` depending on the polarity of the PMIC enable input. This bit affects the polarity of `EXT_PS_CTL1` pin.
- 3: Add the signed offset in `SFLASH_SRSS_PWR_OFFSET.PMIC_VADJ_OFFSET` to account for die to die variation.

16.2.5.5 Internal Regulator Configuration when Using PMIC

To configure the high-current regulator controller (REGHC) to PMIC mode, set the Active and DeepSleep regulators to enable or disable.

When PMIC is enabled and the Active regulator is in parallel, the OCD function of the Active regulator is used to supply power from the PMIC. The OCD is the integrated part of the Active regulator; it is enabled only when the Active regulator is enabled. In power save modes such as DeepSleep, the Active regulator is off; therefore, the OCD is also off. V_{CDD} brownout, such as when the PMIC drops out of regulation range or cannot provide a fast load current increase, can be detected by the MCU. This is effective when the PMIC does not have an OCD function. This function is disabled when the Active regulator is disabled.

The DeepSleep regulator supplies the core supply only in DeepSleep power mode. The register setting helps to decide if the PMIC or the DeepSleep regulator supplies power to the MCU in DeepSleep mode.

The following configurations are available depending on the system.

Table 16-5. Internal Regulator Configuration for PMIC case

Use Case	Configuration	Active Regulator (including internal OCD)	DeepSleep Regulator (including internal OCD)
Use PMIC without own OCD feature. PMIC is disabled in DeepSleep power mode.	OCD is enabled when supplying power from PMIC. DeepSleep regulator supplies power in DeepSleep mode.	Enabled	Enabled
Use PMIC with own OCD feature. PMIC is disabled in DeepSleep power mode.	OCD is disabled when supplying power from PMIC. DeepSleep regulator supplies power in DeepSleep mode.	Disabled	Enabled
Use PMIC with own OCD feature. PMIC is enabled in DeepSleep power mode.	OCD is disabled when supplying power from PMIC. PMIC supplies power in DeepSleep mode.	Disabled	Enabled/Disabled

Note: The internal OCD feature is for self-protection of the internal LDOs and the Active and DeepSleep regulators only. The PMIC without its own OCD feature can be used with the internal OCD feature, but this detection works outside the current path of PMIC.

Note: When you set the DeepSleep regulator to “Disabled” (PWR_CTL2.DPSLP_REG_DIS = 1), you cannot switch back from the PMIC to the internal regulator. If your system is required to switch back to the internal regulator again after handover to PMIC, DeepSleep regulator should be enabled.

16.2.5.6 Transitioning from Active/Sleep to Reset


This section describes the reset behavior for SRSS, trimming, and user program parts.

For SRSS part, a reset results in the following behavior:

- For resets that do not reset the power system (low-voltage (LV) resets such as fault, internal system reset, MCWDT, or CSV), the high-current regulator controller (REGHC) settings are not changed. If REGHC is already operating, it continues to operate. The V_{CCD} rail continues uninterrupted.
- For resets that do reset the power system (high-voltage (HV) resets such as POR, BOD, OVD, OCD, WDT, Hibernate wakeup, or XRES_L), it turns off REGHC and the device restarts with the Active regulator.
- For the trimming part, the normal trim download overwrites the regulator targets with the internal settings. When this happens, the current load is within the Active regulator limit so it does not risk false OCD if REGHC is also enabled.
- For the Reset-Recovery user software, it detects if the device is operating on REGHC and restores settings needed to increase the current level beyond the limit of the Active regulator.

Figure 16-7 shows a simplified flow. The Reset-Recovery flow can be run after every reset.

Figure 16-7. Reset Recovery Flow



1. Check for the REGHC configuration. When `PWR_REGHC_CTL.REGHC_CONFIGURED` is set to 1, go to (2). When `PWR_REGHC_CTL.REGHC_CONFIGURED` is set to 0, the power system is ready. `PWR_REGHC_STATUS.REGHC_ENABLED` indicates the REGHC state.
2. If you have the REGHC with load switch configuration, you need to configure GPIO for REGHC enable control. In this case, REGHC enable state before LV reset is reconfigured. The following is an example of GPIO reconfiguration:
 - a. Write $\sim (PWR\_REGHC\_CTL2.REGHC\_EN \wedge PWR\_REGHC\_CTL.REGHC\_PMIC\_CTL\_POLARITY)$ to GPIO data register.
 - b. Configure to GPIO output.
3. Call the `LoadRegulatorTrims` API with Reset Recovery.
When `LoadRegulatorTrims` returns `STATUS_SUCCESS`, go to (4).

LoadRegulatorTrims returns an error if the hardware state machine is still transitioning. This can happen if the reset occurred during the hardware sequence. If this case occurs, software may wait for the transition to complete (for example, a PMIC that is still turning on). If the transition does not complete within the PMIC enable time, it is recommended to reset the entire chip, including the power system, by using WDT or requesting an external system controller trigger XRES_L.

4. Check the REGHC operation and DeepSleep regulator configuration. When PWR_REGHC_STATUS.REGHC_ENABLED is set to 1, PWR_REGHC_CTL.REGHC_PMIC_USE_LINREG is set to 0, and PWR_REGHC_CTL4.REGHC_PMIC_DPSLP is set to 1, go to (5).
5. Set PWR_CTL2.DPSLP_REG_DIS to 1.
 - After the flowchart completes, it is possible for the device to operate from the internal Active regulator or from REGHC. If the device is operating from REGHC, it is indicated in PWR_REGHC_STATUS.REGHC_ENABLED.

16.2.5.7 DeepSleep Entry/Exit

When entering DeepSleep, perform the following steps:

- If PWR_REGHC_CTL4.REGHC_PMIC_DPSLP = 1, DeepSleep can be entered immediately with no other steps.
- Otherwise, reduce current within the Active regulator limit.
- If a fast wake time is required, change from REGHC to Active regulator before entering DeepSleep. The device wakes with the same regulator that was operating before going to DeepSleep, and Active regulator wakeup is usually faster than PMIC startup time.
- Execute the system call (LoadRegulatorTrims) to update regulator targets for DeepSleep entry.
- Enter DeepSleep.

When exiting DeepSleep, perform the following steps:

- If operating from PMIC and PWR_REGHC_CTL4.REGHC_PMIC_DPSLP = 1, there are no special steps to be performed because the PMIC is already operating and the internal regulator settings are already correct.
- If operating from REGHC, because it was enabled by hardware during DeepSleep wakeup:
 - Wait until PWR_REGHC_STATUS.REGHC_SEQ_BUSY = 0 and PWR_REGHC_STATUS.REGHC_ENABLED = 1.
 - Execute the system call (LoadRegulatorTrims) for DEEPSLEEP exit.
 - After the system call is completed, current will be increased.

16.3 Voltage Monitoring

The Traveo II family offers multiple voltage monitoring and supply failure protection options. This includes POR, BOD, OVD, LVD, OCD, and ADC monitoring. [Table 16-6](#) lists the dedicated supply monitors in the device.

Table 16-6. Dedicated Supply Monitors

Monitor	Monitored Supply	Number of Trip Points	Output	Available Power Mode
POR	V _{DDD}	1 (Fixed)	Reset	All power modes
BOD	V _{DDD}	2 (Programmable)	Reset	All power modes except Hibernate and XRES modes
	V _{DDA}	2 (Programmable)	Reset, Fault, or No action	
	V _{CCD}	1 (Fixed)	Reset	
OVD	V _{DDD}	2 (Programmable)	Reset	
	V _{DDA}	2 (Programmable)	Reset, Fault, or No action	
	V _{CCD}	1 (Fixed)	Reset	
OCD ^a	V _{CCD}	1 (Fixed)	Reset	
LVD	V _{DDD}	26 (Programmable)	Interrupt, Fault, or No action	

a. Traveo II does not have OCD for the PMIC.

16.3.1 Power-On-Reset (POR)

POR circuits provide a reset pulse during the initial power ramp. POR circuits monitor only V_{DDD} voltage. See the device datasheet for details on the POR trip-point levels.

16.3.2 Brownout-Detection (BOD)

The BOD circuit detects supply conditions below a threshold and applies reset to the device. Traveo II offers three BOD circuits – BOD on V_{DDD}, BOD on V_{DDA}, and BOD on V_{CCD}. The system will not come out of RESET until V_{DDD} and V_{CCD} supplies are detected to be valid again. BOD on V_{DDA} is initially disabled and is configurable by software. There is no BOD support in Hibernate and XRES modes. Applications that require BOD support should not use Hibernate mode and should disable it. See the [Device Power Modes chapter on page 186](#) for details.

16.3.2.1 BOD on V_{DDD}

The BOD on V_{DDD} supports two voltage levels (thresholds) to monitor < 2.7 V or < 3.0 V. The PWR_SSV_CTL.BODVDDD_VSEL bit selects the threshold levels of the BOD on V_{DDD}. The BOD on V_{DDD} cannot be disabled. For details on supported thresholds, see the device datasheet and the PWR_SSV_CTL register definition in the *Traveo II Body Controller High Registers TRM*. The PWR_SSV_STATUS.BODVDDD_OK bit indicates the status of the BOD on V_{DDD}. This will always read 1 (no brownout voltage detected), because a detected brownout will reset the chip.

16.3.2.2 BOD on V_{DDA}

The BOD on V_{DDA} supports two voltage levels (thresholds) to monitor < 2.7 V or < 3.0 V. The

PWR_SSV_CTL.BODVDDA_VSEL bit selects the threshold levels of the BOD on V_{DDA}. The PWR_SSV_CTL.BODVDDA_ACTION bits can be used to select a reset, a fault or no action (default). The PWR_SSV_CTL.BODVDDA_ENABLE bit can be used to enable or disable (default) the BOD on V_{DDA}. However, it is not available unless V_{DDD} is present and valid. For details on supported thresholds, see the device datasheet and the PWR_SSV_CTL register definition in the *Traveo II Body Controller High Registers TRM*. The PWR_SSV_STATUS.BODVDDA_OK bit indicates the status of the BOD on V_{DDA}.

16.3.2.3 BOD on V_{CCD}

The BOD on V_{CCD} cannot be disabled. The BOD on V_{CCD} is not as robust as the BOD on V_{DDD}/V_{DDA}. The limitation is because of the small voltage detection range available for this circuit on the minimum allowed V_{CCD}. For details on supported thresholds, see the device datasheet. The robust operation is possible with robust BOD on V_{DDD} and robust OCD on V_{CCD}, even without robust BOD on V_{CCD}. The input voltage to the regulator is robustly supervised by BOD on V_{DDD}. The load current of the Active and DeepSleep regulators are monitored for current that exceeds the regulator limit by OCD on V_{CCD}. Therefore, OCD monitors the operating conditions met when using the internal regulators. However, Traveo II does not monitor the current of the PMIC. The PWR_SSV_STATUS.BODVCCD_OK bit indicates the status of the BOD on V_{CCD}. This will always read '1' (no brownout voltage detected), because a detected brownout will reset the chip.

16.3.3 Over-Voltage Detection (OVD)

Traveo II offers three over-voltage detection circuits that monitor V_{CCD} , V_{DDD} , and V_{DDA} supply. Similar to the BOD circuit, the OVD circuit detects supply conditions above a threshold and applies a reset. As the name suggests, the OVD circuit maintains a device reset, if V_{CCD} or V_{DDD} supply stays higher than thresholds. The OVD circuit can generate a reset in all device power modes except the Hibernate and XRES modes, provided the V_{DDD} and V_{DDA} supply ramp satisfies the datasheet maximum supply ramp limits in that mode. Applications that require OVD support should not use Hibernate mode and should disable it. See the [Device Power Modes chapter on page 186](#) for details.

16.3.3.1 OVD on V_{DDD}

The OVD on V_{DDD} supports two voltage levels (thresholds) to monitor > 5.5 V or > 5.0 V. The `PWR_SSV_CTL.OVDVDDD_VSEL` bit selects the threshold levels of the OVD on V_{DDD} . The OVD on V_{DDD} cannot be disabled. For details on supported thresholds, see the device datasheet and the `PWR_SSV_CTL` register definition in the *Traveo II Body Controller High Registers TRM*. The `PWR_SSV_STATUS.OVDVDDD_OK` bit indicates the status of the OVD on V_{DDD} . This will always read 1 (no overvoltage detected), because a detected overvoltage will reset the chip.

16.3.3.2 OVD on V_{DDA}

The OVD on V_{DDA} supports two voltage levels (thresholds) to monitor > 5.5 V or > 5.0 V. The `PWR_SSV_CTL.OVDVDDA_VSEL` bit selects the threshold levels of the OVD on V_{DDA} . The `PWR_SSV_CTL.OVDVDDA_ACTION` bits can be used to select a reset, a fault, or no action (default). The `PWR_SSV_CTL.OVDVDDA_ENABLE` bit can be used to enable or disable (default) the OVD on V_{DDA} . However, it is not available unless V_{DDD} is present and valid. For details on supported thresholds, see the device datasheet and the `PWR_SSV_CTL` register definition in the *Traveo II Body Controller High Registers TRM*. The `PWR_SSV_STATUS.OVDVDDA_OK` bit indicates the status of the OVD on V_{DDA} .

16.3.3.3 OVD on V_{CCD}

The OVD on V_{CCD} cannot be disabled. For details on supported thresholds, see the device datasheet. The `PWR_SSV_STATUS.OVDVCCD_OK` bit indicates the status of the OVD on V_{CCD} . This will always read 1 (no overvoltage detected), because a detected overvoltage will reset the chip.

16.3.4 Low-Voltage-Detection (LVD)


Two LVD circuits monitor external supply voltage (V_{DDD}) and detects depletion of the energy source. The LVD detectors generate an interrupt or a fault that causes the system to take preventive measures. The `PWR_LVD_CTL/2.HVLVD1/2_ACTION` bit can be used to select an interrupt or a fault.

These low-voltage detection circuits can also be used for high-voltage detection (HVD). They can each be configured as LVD (falling detection), HVD (rising detection), or both by the `PWR_LVD_CTL/2.HVLVD1/2_EDGE_SEL` bits. Each LVD supports up to 26 voltage levels (thresholds) to monitor between 2.8 V and 5.3 V. The `PWR_LVD_CTL/2.HVLVD1/2_TRIPSEL_HT` bits select the threshold levels of the HVLVD1/2. The LVD should be disabled before selecting the threshold. The `PWR_LVD_CTL/2.HVLVD1/2_EN_HT` bit can be used to enable or disable the HVLVD1/2. The LVD operates in Active, Sleep, and DeepSleep modes. It does not operate in Hibernate and XRES modes. To use HVLVD1/2 in DeepSleep mode, the `PWR_LVD_CTL/2.HVLVD1/2_DPSLP_EN_HT` bit should be enabled.

Whenever the voltage level of the supply being monitored crosses the threshold, the LVD generates an interrupt or a fault. This interrupt status is available in the `SRSS_INTR` register. And the real-time status is available in the `PWR_LVD_STATUS/2` register.

The `SRSS_INTR` register indicates a pending LVD interrupt. The `SRSS_INTR_MASK` register decides whether LVD interrupts are forwarded to the CPU.

For details on supported LVD thresholds, see the device datasheet and the `PWR_LVD_CTL/2` register definition in the *Traveo II Body Controller High Registers TRM*.



16.3.5 Over-Current Detection

The OCD circuit monitors V_{CCD} current and detects if the load current of a regulator is higher than expected. If the current is over the regulator limit, the OCD circuit generates a reset to protect the device. OCD for the high-current regulator controller is not active unless the high-current regulator controller is enabled. For details on detection range, see the device datasheet. OCD operates in Active, Sleep, and DeepSleep modes. Because the regulators are disabled in Hibernate mode, the OCD circuit also does not operate. The PWR_SSV_STATUS register indicates the OCD status. Traveo II does not have OCD for the PMIC. However, when $PWR\_REGHC\_CTL\_REGHC\_PMIC\_USE\_LINREG = 1$, MCU keeps the internal Active regulator and its OCD enabled to improve supply supervision of V_{CCD} for the external PMIC mode. When using this feature, if the PMIC fails to keep V_{CCD} above the internal regulator target, then the internal regulator will attempt to recover V_{CCD} . If the regulator current is too high, the regulator triggers an OCD reset.

16.3.6 Voltage Monitoring by ADC

In addition to the dedicated monitors, analog connections are provided to allow the ADC to monitor all high-voltage supplies and grounds. This is the only monitor capability provided for supplies without dedicated monitors (such as multiple V_{DDIO} rails).

To facilitate voltage monitoring by the ADC, a monitor switch in the power pad creates a connection between the power or ground pad and the AMUXBUS. The HSIOM_MONITOR_CTL_0 register controls the connectivity of power/ground pads to either AMUXBUS_A or AMUXBUS_B respectively. For details, see [Table 16-7](#). The power monitor cell can connect the power pad to AMUXBUS_A as shown in [Figure 16-9](#). The ground monitor cell can connect ground pad to AMUXBUS_B. It is shown in [Figure 16-10](#). For details on HSIOM_MONITOR_CTL_0 register, see the *Traveo II Body Controller High Registers TRM*.

The series resistor is intended to allow voltage division using a matching resistor in the ADC. This enables measuring supplies outside of the V_{DDA} (V_{REFH})/ V_{SSA} (V_{REFL}) limits. For details on the ADC, see the [Reference Buffer on page 679](#).

Table 16-7. Relation between HSIOM_MONITOR_CTL_0 Register and Power/Ground Pads

HSIOM_MONITOR_CTL_0	Power/Ground Pads	AMUXBUS	BGA-320	BGA-272 for CYT4BF	BGA-272 for CYT3BB/4BB	TEQFP-176	TEQFP-144	TEQFP-100
Bit 0	V_{DDD}	A	F8, F9, H15, J15, K15, L15, M15, N15, R12, R13	F8,H13,J13,K13,L13,N11	F8,H13,J13,K13,L13,N11	176	144	100
Bit 2						22	18	12
Bit 4						43	35	24
Bit 13						110	90	62
Bit 15						132	108	75
Bit 17						153	124	86
Bit 1	V_{SSD}	B	A1, A20, C3, C10, C18, H9, H10, H11, H12, H13, J9, J10, J11, J12, J13, J18, K9, K10, K11, K12, K13, K18, L9, L10, L11, L12, L13, M9, M10, M11, M12, M13, N12, V3, V4, V15, Y1, Y20	A1, A18, D9, G7, G12, H9, H10, H11, J9, J10, J11, J15, K9, K10, K11, M7, M12, R5, R14, V1, V18	A1,A18,D9,G7,G12,H9,H10,H11,J9,J10,J11,J15,K9,K10,K11,M7,M12,R5,R14,V1,V18,L9,L10	1	1	1
Bit 3						23	19	13
Bit 6						45	37	26
Bit 12						89	73	51
Bit 14						114	94	66
Bit 16						133	109	76
Bit 19						155	126	88
Bit 5	V_{DDIO_1}	A	F10, F11, F12, F13	F9, F10, F11	F9,F10,F11	44	36	25
Bit 18	V_{SSD_1} (BGA)	B	N13	L11	L11	154	125	87
	V_{SSD} (TEQFP)							
Bit 7	V_{REFL}	B	M8	K8	K8	76	62	41
Bit 8	V_{SSA}	B	N8	L8	L8	77	63	42
Bit 9	V_{DDA}	A	N6	L6	L6	78	64	43
Bit 10	V_{REFH}	A	M6	K6	K6	79	65	44
Bit 11	V_{DDIO_2}	A	R8	N8	N8, N9, N10	88	72	50

Table 16-7. Relation between HSIOM_MONITOR_CTL_0 Register and Power/Ground Pads

HSIOM_MONITOR_CTL_0	Power/Ground Pads	AMUXBUS	BGA-320	BGA-272 for CYT4BF	BGA-272 for CYT3BB/4BB	TEQFP-176	TEQFP-144	TEQFP-100
Bit 20	V _{DDIO_3}	A	H6, J6, K6, L6	H6, J6	H6, J6	—	—	—
Bit 21	V _{SSIO_3}	B	H8, J8, K8, L8	H8, J8	H8, J8	—	—	—
Bit 22	V _{DDIO_4}	A	R9, R10, R11	N9, N10	—	—	—	—
Bit 24						—	—	—
Bit 23	V _{SSIO_4}	B	N9, N10, N11	L9, L10	—	—	—	—
Bit 25						—	—	—

Figure 16-9. Power Monitor Cell





Figure 16-10. Ground Monitor Cell



16.4 Register List

Register	Name	Description
PWR_LVD_CTL	High-Voltage/Low-Voltage Detector (HVLVD) Configuration Register	This register shows the configuration bits for HVLVD1
PWR_LVD_CTL2	High-Voltage/Low-Voltage Detector (HVLVD) Configuration Register 2	This register shows the configuration bits for HVLVD2
PWR_LVD_STATUS	High-Voltage/Low-Voltage Detector (HVLVD) Status Register	This register shows the real time status for HVLVD1
PWR_LVD_STATUS2	High-Voltage/Low-Voltage Detector (HVLVD) Status Register 2	This register shows the real time status for HVLVD2
PWR_SSV_CTL	Supply Supervision Control Register	This register shows the controls for BOD and OVD
PWR_SSV_STATUS	Supply Supervision Status Register	This register shows the status for BOD and OVD
SRSS_INTR	SRSS Interrupt Register	This register shows interrupt requests from the SRSS peripheral.
SRSS_INTR_SET	SRSS Interrupt Set Register	This register is used for firmware testing.
SRSS_INTR_MASK	SRSS Interrupt Mask Register	This register controls forwarding of the interrupt to CPU.
SRSS_INTR_MASKED	SRSS Interrupt Masked Register	This register shows the logical AND of the corresponding SRSS interrupt request (SRSS Interrupt register) and mask bits (SRSS Interrupt Mask register)
HSIOM_MONITOR_CTL_0	Power/Ground Monitor Cell Control 0 Register	This register controls the connectivity of Power/Ground monitor cells to either AMUXBUS A or B respectively.
PWR_REGHC_CTL	High-current Regulator Control Register	This register shows the control for the high-current regulator controller.
PWR_REGHC_CTL2	High-current Regulator Control Register 2	This register shows the control for the high-current regulator controller.
PWR_REGHC_CTL4	High-current Regulator Control Register 4	This register shows the control for the high-current regulator controller.
PWR_REGHC_STATUS	High-current Regulator Status Register	This register shows the status register for the high-current regulator controller.

17. Device Power Modes



The TVII-B-H device can operate in different power modes that are intended to minimize the average power consumption in an application. The power modes supported by TVII-B-H in the order of decreasing power consumption are:

- Active – all peripherals are available
- Low-Power Active (LPACTIVE) profile – Low-power profile of Active mode where all peripherals including the CPU are available, but with limited capability
- Sleep – all peripherals except the CPU are available
- Low-Power Sleep (LPSLEEP) profile – Low-power profile of Sleep mode where all peripherals except the CPU are available, but with limited capability
- DeepSleep – only low-frequency peripherals are available
- Hibernate – the device and I/O states are frozen and the device resets on wakeup
- XRES – the device enters this state when the XRES_L pin is asserted

Active, Sleep, and DeepSleep are standard Arm-defined power modes supported by the Arm CPUs and Instruction Set Architecture (ISA). Hibernate mode is an additional low-power mode supported in TVII-B-H. LPACTIVE and LPSLEEP are similar to Active and Sleep modes, respectively; however, the high-current components are either frequency or current limited or turned off. Hibernate mode and XRES state are the lowest power mode/state that the TVII-B-H device can be in. On wakeup from XRES or Hibernate mode, the CPU and most peripherals go through a reset. Peripherals such as RTC or watchdog can be used during any of these power modes and also to trigger a transition to other active power modes.

17.1 Features

TVII-B-H power modes have the following features:

- Software can use power modes to optimize power consumption in an application
- Low-power DeepSleep mode with support for multiple wakeup sources and configurable amount of SRAM retention
- Ultra-low-power Hibernate mode with wakeup from I/O and timer alarms

The power consumption in different power modes is controlled by using the following methods:

- Enabling and disabling clocks to peripherals
- Powering on/off clock sources
- Powering on/off peripherals and parts inside the device

17.2 Device Power Modes

Table 17-1 summarizes the power modes available in TVII-B-H, their description, and details on entry and exit conditions.

Table 17-1. TVII-B-H Power Modes

Power Mode	Description	Entry Condition	Wakeup Source	Wakeup Action
Active	Primary mode of operation; all peripherals are available (programmable).	Wakeup from Sleep/DeepSleep modes, Hibernate reset, or any other reset.	Not applicable	Not applicable
Low-Power Active Profile	A low-power profile of Active mode; most peripherals are available with limited capabilities	Register write from Active mode and wakeup from LPSLEEP/DeepSleep modes.	Not applicable	Not applicable
Sleep	CPU is in Sleep mode; all other peripherals are available.	Register write from Active mode or wakeup from DeepSleep through debugger	Any interrupt to CPU	Interrupt
Low-Power Sleep Profile	A low-power profile of Sleep mode; CPU is in Sleep mode; most peripherals are available with limited capabilities.	Register write from LPACTIVE mode.	Any interrupt to CPU	Interrupt
DeepSleep	All high-frequency clocks and peripherals are turned off. Low-frequency clock (ILO) and low-power analog and digital peripherals are available for operation and as wakeup sources. SRAM can be retained (configurable).	Register write from Active or LPACTIVE modes.	GPIO interrupt, event generators, SCB ^a , watchdog timer, and RTC alarms ^b and debugger	Interrupt or debug
Hibernate	GPIO states are frozen; all high-frequency clocks and peripherals are switched off. Low-frequency clocks (32 kHz), WCO, or LPECO can function. Device resets on wakeup event.	Register write from Active or LPACTIVE modes.	WAKEUP pins, RTC alarm, and watchdog timer	Hibernate Reset

a. See the device-specific datasheet for the SCB-instance capable of waking up the device from DeepSleep mode.

b. RTC (along with optional WCO/LPECO) is supplied with V_{DD} and is available irrespective of the device power mode. RTC alarms are capable of waking up the device from any power mode.

17.2.1 Active and Sleep Modes

The Active and Sleep modes are the standard Arm-defined power modes supported by both Cortex-M7 and Cortex-M0+ cores.

The device enters Active mode upon any reset. In this mode, the CPU executes code along with all logic and memory powered. The firmware may decide to enable or disable specific peripherals and power domains depending on the application and power requirement. All the peripherals are available for use in Active mode.

In Sleep mode, the CPU clock is turned off and the CPU enters sleep. Note that in TVII-B-H, both Cortex-M7 and Cortex-M0+ support their own CPU sleep modes and each CPU can be in sleep, independent of the state of the other CPU. But the device is said to be in Sleep mode, when both the cores are in sleep. All peripherals available in Active mode are available in Sleep mode. Any unmasked interrupt can wake up the CPU to Active mode.

17.2.1.1 Low-Power Profiles - LPACTIVE and LPSLEEP

Low-power profiles are intended to reduce power consumption during Active or Sleep mode. They are software-controlled configurations to fine-tune current consumption. Power consumption can be reduced by controlling the following parameters:

- Reducing frequency either by selecting a slower source (such as IMO) or selecting a different output frequency (such as FLL and PLL), or dividing the clock using the pre-divider or PERI dividers.
- Disabling unnecessary clocks either at the source or by disabling the clock root muxes.
- Disabling unused circuitry such as low-voltage detection (LVD), which are used periodically to monitor external power supply source (such as battery).
- Disabling internal clock sources that are not generating a system clock. All clock sources are initially disabled, except the IMO. Note that some clock sources, such as the crystal oscillators (WCO and ECO) have relatively long startup times. Switching these circuits off and on may result in more overall current if the system must idle while they start up.

- Either Cortex-M7 or Cortex-M0+, or both can be put to the sleep state by controlling the clock provided to them. Cores can be put to either sleep or deep-sleep states depending on the configurations.
- Firmware may allow disabling the flash macro. An unused macro can be disabled to reduce static current consumption; this can be done dynamically based on the application need to access a macro. Further, some code can be copied to SRAM and run from there, because reading from SRAM takes less current than reading from flash. In such a case, it may be possible to disable the flash macro. The current savings needs to be compared with the cost of re-enabling the macro and copying the code.

Examples of low-power profiles are as follows.

- **LPACTIVE:** Low-speed source clock (IMO), PLL/FLL off, Cortex-M7 in Sleep mode, and Cortex-M0+ in Active mode.
- **LPSLEEP:** Low-speed source clock (IMO), PLL/FLL off, Cortex-M7 in DeepSleep mode, and Cortex-M0+ in Sleep mode.

Using such configurations in combination with cyclic wakeup from DeepSleep can help achieve low-power operation.

17.2.1.2 CM7 Power Domains

When the system is in Active power mode, each CM7 CPU can be put in one of the following CM7 power modes shown in [Table 17-2](#) using the CPUSS_CM7_x_PWR_CTL MMIO register. See the Registers TRM for the register description.

Table 17-2. CM7 Power Domains

Power Mode	Description
ENABLED	Switch CM7 on. Power on, clock on.
RESET	Reset CM7. Clock off, no retain and reset. Note: CM7 CPU has the AIRCR.SYSRESETREQ register field, which allows the CM7 to reset the complete device.
RETAINED	Put CM7 in Retained mode. This can only become effective if CM7 is in DeepSleep CPU power mode. Check the PWR_DONE flag to see if CM7 RETAINED state is reached. Power off, clock off, retain and no reset.
OFF	Switch CM7 off. Power off, clock off, no retain and reset.

Note: Accessing the CM7 ITCM/DTCM by other masters, at addresses CM7_0_ITCM_ADDR, CM7_0_DTCM_ADDR, CM7_1_ITCM_ADDR, and CM7_1_DTCM_ADDR is possible only when the power mode of the corresponding CM7 is ENABLED.

Before changing the CM7 power mode from ENABLED to another mode, perform the following steps (accessing the TCM without these steps may cause the device to hang):

1. Disable access to the CM7 TCM by setting the CM7_0/1_CTL.TCMC_EN field to '0'.
2. Confirm that there are no outstanding accesses to the CM7 TCM from other bus masters by checking if the CM7_0/1_STATUS.TCMC_* fields are '0'. Repeat step 2 if necessary.
3. Now it is safe to change the CM7 power mode.

When the debugger is connected, OFF and RESET modes behave similarly. If the mode is changed from ENABLED to RESET/OFF of any CM7 core, the respective CM7 comes to the reset handler and starts execution from the Vector table base.

When the debugger is not connected, CM7 can be transitioned to ENABLED mode from either RESET or OFF, by configuring CPUSS_CM7_x_PWR_CTL.PWR_MODE to ENABLED.

It is recommended to use RESET when CM7 is intended to go through reset, and OFF when the intention is to save power by switching off CM7.

17.2.2 DeepSleep Mode

In DeepSleep mode, all the high-speed clock sources are off and high-speed peripherals are unusable. Low-speed clock sources and peripherals continue to operate, if configured and enabled by the firmware. In addition, peripherals that do not need a clock or receive clock from their external interface continue to operate, if configured for DeepSleep operation. TVII-B-H provides an option to configure the amount of SRAM, in blocks of 32 KB, to be retained during DeepSleep.

Note that both Cortex-M0+ and Cortex-M7 can enter their local DeepSleep mode independently. However, the entire device enters DeepSleep mode only when both the CPUs are in the deep-sleep state. The device can enter DeepSleep mode after the following conditions are met.

- PWR_CTL.LPM_READY should read '1'. This ensures the device is ready to enter low-power modes. If the PWR_CTL.LPM_READY reads '0', then the device will enter normal CPU sleep instead of DeepSleep until the bit is set, at which instant the device will automatically enter DeepSleep mode, if requested.
- Both Cortex-M0+ and Cortex-M7 are in DeepSleep. This is achieved by setting SCR.SLEEPDEEP of both Cortex-M0+ (CM0P_SCS_SCR) and Cortex-M7 (CM7_0/CM7_1_SCS_SCR).
- Debugger is not connected.

Refer to [Debugger Effect on Device Power Modes on page 194](#) for more information about how the debug session affects power mode transitions. Cortex-M0+ must make sure

that there are no pending flash memory transactions (write/erase operation) before going to DeepSleep mode.

In this mode, the Active mode regulator is turned off and a low-power DeepSleep regulator supplies peripherals in DeepSleep mode. Table 17-6 provides the list of resources available in DeepSleep mode.

Interrupts from low-speed asynchronous or low-power analog peripherals can cause a CPU wakeup from DeepSleep mode. A debug wakeup from DeepSleep returns to Sleep mode.

17.2.3 Hibernate Mode

Hibernate mode is the lowest power mode of the device when external supplies are still present and XRES_L is deasserted. It is intended for applications in a dormant state. In this mode, both the Active and DeepSleep regulators are turned off and GPIO states must be frozen.

Hibernate mode is entered by performing three identical writes to the PWR_HIBERNATE register. Each of these writes should have the UNLOCK code, set FREEZE, set Hibernate commands, and load the other fields (TOKEN, POLARITY_HIBPIN, MASK_HIBPIN, MASK_HIBALARM, MASK_HIBWDT) as desired. The first write unlocks Hibernate; the second freezes the I/Os; and the third enters Hibernate mode. Unlike entry to DeepSleep mode, active debug session cannot prevent transition to Hibernate mode. Instead, after the device enters Hibernate mode, debugger host will be disconnected.

Hibernate mode is exited by either generating a wakeup event or asserting XRES_L. A wakeup event can come from dedicated wakeup pins (up to four pins) with configurable polarity or through alarms from RTC or WDT wakeup event. All wakeup signals from GPIO pins or XRES_L are level-sensitive and must be held long enough for the Hibernate bit to be cleared. Set the respective mask bits for alarm, WDT, or for external pins to wake up the device from Hibernate mode. See the device datasheet for the supported number of pins that can wake up the device from Hibernate mode.

Note: See the device datasheet for information about the number of wakeup pins supported. For unsupported pins, the respective MASK_HIBPIN bits must not be set to high.

The device goes through a reset (except RTC, Backup registers, and Hibernate registers) on wakeup and I/O pins must be unfrozen by firmware upon entering Active mode. The PWR_HIBERNATE (except the Hibernate bit [31]) register along with the PWR_HIB_DATA register are retained through the Hibernate wakeup sequence and can be used by the application to retain some content through the Hibernate wakeup sequence. Note that these registers are reset by other reset events. On a Hibernate wakeup event, PWR_HIBERNATE.HIBERNATE bit is cleared.

Asserting XRES_L in Hibernate mode will lead to device reset; however, it is not the wakeup event. In this case, GPIOs will lose their frozen state and will be tristated.

Consider these restrictions while using Hibernate mode:

- Supplies must be stable through Hibernate mode
- Supplies must remain stable from 250 μ s before entering Hibernate mode until Hibernate is fully entered. This allows the key writer to write all requested keys completely. Failure to observe this requirement can result in undefined behavior.
- The brownout detect (BOD) or overvoltage detect (OVD) blocks are not available in Hibernate mode. As a result, the device will not recover from a brownout or overvoltage event in Hibernate mode. If detection is needed, an external supervisor can be used to assert XRES_L in a brownout or overvoltage condition. Otherwise, it is recommended not to enter Hibernate mode in applications that require brownout or overvoltage detection.

If these restrictions are unacceptable, accidental entry into Hibernate mode can be prevented using the disable option – set PWR_HIBERNATE.HIBERNATE_DISABLE. Note that this bit is a write-once bit during execution and will be cleared on reset.

Notes:

- SRAM cannot be retained in Hibernate mode.
- The device has a separate clock domain, which can be ON, irrespective of the power modes mentioned earlier. This domain contains an RTC and WCO. The RTC provides an option to wake up the device from any of the low-power modes. It can be clocked by an external clock source such as WCO or LPECO¹, or by the internal low-speed oscillator (ILO0). This always-on domain is powered internally by V_{DD}. It also offers a set of 32-bit backup registers (BACKUP_BREGx), which will retain the data through DeepSleep and Hibernate modes.

17.2.4 Other Operational States

In addition to the power modes discussed in the previous sections, there are two other states the device can be in – XRES and OFF state. You do not need a firmware action to enter these states or an interrupt or wakeup event to exit them. The device may be in these states if it is not in any of the modes described earlier.

1. See the device-specific datasheet to check whether LPECO is supported.

XRES is the device state when an external reset (XRES_L pin) is applied. XRES is not a power mode. During the XRES state, all the components in the device are powered down and I/Os are tristated keeping the power consumption to a minimum. The OFF state simply represents the device state with no power or insufficient power applied. The XRES and OFF states are discussed for completeness of all possible states the device can be in.

Reset is an intermediate state while the device starts.

Figure 17-1 shows various states the device can be in along with possible power mode transition paths. The transitions are described in detail in later sections.

The flowchart details the power management logic, including transitions between states like **RESET**, **ACTIVE** (with **LOW POWER PROFILE #0** and **MAIN PROFILE**), **SLEEP** (with **LOW POWER PROFILE #0** and **MAIN PROFILE**), **DEEPSLEEP**, **HIBERNATE**, **XRES**, and **OFF**. Transitions are triggered by events such as **Internal Resets**, **External Reset Event**, **Power Down or Brown Out Event**, **Firmware or Debugger/ATE Action**, **Debugger/ATE Action**, and **Other Event**. Specific actions like **Register Configuration**, **Interrupt**, **WFI/WFE**, **DeepSleep Interrupt**, **Debugger Session Starts/Ends**, **WAKEUP Asserts**, **XRES, L Deasserts**, **External Power Ramps**, and **PowerDown or BrownOut** are also shown.

Traveo II Body High Architecture TRM, Document No. 002-24401 Rev. *D

17.3.1 Power-up Transitions

Table 17-3 summarizes various power-up transitions, their type, triggers, and actions.

Table 17-3. Power Mode Transition

Initial State/Mode	Final State/Mode	Type	Trigger	Actions
OFF	XRES	External	Power rail (V_{DD}) ramps up above POR voltage level with XRES_L pin asserted.	1. All high-voltage logic such as SRSS, WDT, BOD/POR, Hibernate control, and IMO are reset. Low-voltage logic is powered off.
OFF	Reset	External	Power rail (V_{DD}) ramps up above POR and BOD voltage level with XRES_L pin deasserted.	1. All high-voltage logic is reset 2. Low-voltage (internal Active and DeepSleep mode) regulators and references are ramped up 3. All low-voltage logic (operating from internal regulators) such as CPUs, high-speed peripherals, MCWDT, and low-speed peripherals are reset 4. IMO clock is started
XRES	Reset	External	XRES_L pin is deasserted with V_{DD} present and above POR and BOD level.	1. Low-voltage regulators and references are ramped up 2. All low-voltage logic is reset 3. IMO clock is started
Reset	Active	Internal	Reset sequence completes. This transition can also be caused by internal resets or Hibernate wakeup event.	1. Clock is released to the system 2. System reset is deasserted 3. CPU starts execution

17.3.2 Low-Power Mode Transition

Table 17-4 discusses various low-power mode transitions.

Table 17-4. Low-Power Mode Transitions

Initial State/Mode	Final State/Mode	Type	Trigger	Actions
Active	Sleep	Internal	Firmware action 1. Clear SCR.SLEEPDEEP for both Cortex-M0+ (CM0P_SCS_SCR) and Cortex-M7 (CM7_0/CM7_1_SCS_SCR). 2. Optionally, set SCR.SLEEPONEXIT if the CPU runs only on interrupts. When this bit is set, the CPU will not return to application code after the WFI/WFE instruction is executed. The CPU will wake up on any enabled interrupt or event and will enter Sleep/DeepSleep mode as soon as it exits the interrupt or services the event. 3. Optionally, set SCR.SEVONPEND if the application needs to wake up the CPU from any pending interrupt. If this bit is set, any interrupt that enters a pending state will wake up the CPU. 4. Execute WFI/WFE instruction on both CPUs.	1. CPU clocks are gated off 2. CPU waits for an interrupt or event to wake it up.

Table 17-4. Low-Power Mode Transitions

Initial State/Mode	Final State/Mode	Type	Trigger	Actions
Active	DeepSleep	Internal	Firmware action Perform these steps to enter DeepSleep mode (PWR_CTL.LPM_READY should read '1' before performing these steps): 1. Set SCR.SLEEPDEEP for both Cortex-M0+ (CM0P_SCS_SCR) and Cortex-M7 (CM7_0/CM7_1_SCS_SCR). 2. Optionally, set SCR.SLEEPONEXIT if the CPU runs only on interrupts. When this bit is set, the CPU will not return to application code after the WFI/WFE instruction is executed. The CPU will wake up on any enabled interrupt or event and will enter Sleep/DeepSleep mode as soon as it exits the interrupt or services the event. 3. Optionally, set SCR.SEVONPEND if the application needs to wake up the CPU from any pending interrupt. If this bit is set, any interrupt that enters a pending state will wake up the CPU. 4. Execute WFI/WFE instruction on both CPUs. Note: Executing this sequence before the low-power mode is ready ((PWR_CTL.LPM_READY==1)) will make the transition first to Sleep mode. The device state will automatically move to DeepSleep when PWR_CTL.LPM_READY is set. Note: Make sure that any write transfer made before executing the WFI instruction is followed by the read access to the same memory location. This ensures that the write operation is successful.	1. CPU enters low-power mode. 2. High-frequency clocks are shut down. 3. I/O cells associated with DeepSleep-enabled blocks will be functional; the remaining I/Os and their configurations will be frozen automatically. 4. Retention is enabled and non-retention logic is reset. 5. Active regulator is disabled and DeepSleep regulator takes over.

Table 17-4. Low-Power Mode Transitions

Initial State/Mode	Final State/Mode	Type	Trigger	Actions
Active	Hibernate	Internal	Firmware action 1. Set PWR_HIBERNATE.TOKEN (optional) and PWR_HIB_DATA register to some application-specific branching data that can be used on a wakeup event from Hibernate mode. 2. Set PWR_HIBERNATE.UNLOCK to 0x3A, this ungates writes to PWR_HIBERNATE.FREEZE and PWR_HIBERNATE.HIBERNATE bits. 3. Configure wakeup pins polarity (PWR_HIBERNATE.POLARITY_HIBPIN), wakeup pins mask (PWR_HIBERNATE.MASK_HIBPIN), wakeup alarm mask (PWR_HIBERNATE.HIBALARM), and watchdog interrupt mask (PWR_HIBERNATE.MASK_HIBWDT) based on the application requirement. 4. Set PWR_HIBERNATE.FREEZE to freeze the I/O pins. 5. Set PWR_HIBERNATE.HIBERNATE to enter Hibernate mode. 6. Read the PWR_HIBERNATE register to make sure that the write has taken effect. 7. Execute WFI instruction on both CPUs. Note: Transition to HIBERNATE mode can be canceled before setting PWR_HIBERNATE.HIBERNATE. To do so, clear PWR_HIBERNATE.FREEZE and UNLOCK to return to ACTIVE mode. Note: It is recommended to trigger Hibernate mode atomically. This means, when entering the Hibernate mode, disable all the interrupts and do a write operation on the PWR_Hibernate register. Note: Make sure that any write transfer made before executing the WFI instruction is followed by the read access to the same memory location. This ensures that the write operation is successful.	1. CPU enters low-power mode. 2. Both high-frequency and low-frequency clocks except RTC are shut down. 3. Pin output states and configurations are frozen. 4. Both Active and DeepSleep regulators are powered down. The peripherals that are active in the Hibernate domain operate directly out of V_{DD}.
Sleep	DeepSleep	Internal	When the debugger is not connected and DeepSleep mode is triggered, but PWR_CTL.LPM_READY==0, the device internally enters the Sleep mode. The device will automatically transit to DeepSleep when PWR_CTL.LPM_READY==1. If the debugger is connected and DeepSleep mode is triggered by the firmware, the device will enter DeepSleep only when the following conditions are met. 1. PWR_CTL.LPM_READY==1 2. Debugger is disconnected	1. High-frequency clocks are shut down. 2. I/O cells associated with DeepSleep-enabled blocks will be functional; the remaining I/Os and their configurations will be frozen automatically. 3. Retention is enabled and non-retention logic is reset. 4. Active regulator is disabled and DeepSleep regulator takes over.

17.3.3 Wakeup

Table 17-5 shows the sequence from low-power mode to Active mode.

Table 17-5. Wakeup Sequence

Initial State/Mode	Final State/Mode	Trigger Source	Action
Sleep	Active	Any enabled interrupt	CPU exits Sleep mode and executes the interrupt
DeepSleep	Active	Low-speed peripherals or interrupt from DeepSleep peripheral	Device returns to the configuration it had while entering DeepSleep mode. 1. IMO/clocks enabled 2. Non-retained state is reset 3. GPIOs are unfrozen 4. CPU exits low-power mode and executes interrupt
DeepSleep	Sleep	Debug wakeup	1. Non-retained state is reset 2. GPIOs are unfrozen 3. High-frequency and low-frequency clocks are ON 4. CPU remains in Sleep
Hibernate	Active	RTC, WDT, wakeup from up to four pins	Hibernate wakeup is implemented as transition to Active mode through Reset. 1. Low-voltage (internal Active and DeepSleep mode) regulators and references are ramped up 2. All low-voltage logic is reset 3. IMO clock starts 4. CPU starts execution 5. Software must unfreeze the I/Os by unlocking and clearing the PWR_HIBERNATE.FREEZE bit.

17.3.4 Internal Reset Transitions

When an internal reset occurs:

- I/O cells are disabled (excluding the PMIC control interface).
- Most low-voltage logic is reset. Exceptions include reference settings, regulator settings, reset cause registers, and fault logging system.
- Most high-voltage logic is not reset, including hibernate peripherals, WDT, and RTC and BREG registers.

When the device is in Active/Sleep mode and internal reset occurs, the reference and regulator settings are not changed; SRSS enables the IMO (if disabled) and makes the Reset to Active transition.

While the device is in DeepSleep mode and internal reset occurs, then I/O cells are disabled (Hi-Z), most low-voltage logic is reset, retention is disabled, regulators are enabled, IMO starts, and the device enters Active state.

17.3.5 Powering Down/Brownout/Overvoltage

This transition occurs when power is partly or partially lost, and as a result one of the brownout or overvoltage detectors execute reset.

Note that the detectors are disabled in Hibernate mode. If V_{DD} is removed or becomes invalid in Hibernate mode,

then the system must restart with XRES_L applied. This is because the logic dependent on the V_{DD} will slowly discharge and may become invalid.

17.3.6 Debugger Effect on Device Power Modes

The debugger uses non-AHB registers through SWD and JTAG port to transition to/from debug mode. After the debugger connection is established with the device, CPUSS_DP_STATUS.SWJ_CONNECTED bit is set. Debugger connection is possible when the device is in Active/Sleep/DeepSleep power modes but not when the device is in Hibernate power mode.

Device will behave differently in certain cases when the debug connection is active. Some instances are as follows:

- Attempt to enter DeepSleep mode results in a transition to Sleep mode instead, with power and clocks unchanged
 - System power consumption is the same as Sleep mode, which is higher than DeepSleep mode.
 - Wakeup time will be the Sleep wakeup time, which is shorter than a DeepSleep wakeup time.
 - Non-retention registers will not be reset upon wakeup and may lead to behavior that is different from its actual operation during a debug session.

If the debugger is disconnected during active DeepSleep request, SRSS will transition to DeepSleep mode.

- A debugger connection request, when device is in DeepSleep mode, will take the device to Sleep mode. Mode transition from DeepSleep to Sleep is described in [Wakeup on page 194](#).

An active debug session cannot prevent Hibernate or OFF mode entry. When the device enters these modes, the debugger will be disconnected because the debug port is no longer powered.

17.4 Summary

Table 17-6 captures various device components and their availability during the device power modes/states.

Table 17-6. Resource Available in Different Power Modes/States

Component	Power Modes/States			
	Active/Sleep	DeepSleep	Hibernate	XRES
Wakeup				
Wakeup Event	Any interrupt	DeepSleep peripherals/GPIOs	Dedicated wakeup pins ^a	XRES
Wakeup Action	Interrupt	Interrupt	Reset	Reset
Core Function				
CPU	On/Sleep	Retention	Off	Off
SRAM	On	Retention (opt) ^b	Off	Off
Flash	On	Off	Off	Off
High Speed Clock (IMO, ECO, PLL)	On	Off	Off	Off
LVD	On (opt)	Slow (opt) ^c	Off	Off
ILO	On	On	On	Off
CSV	On (opt)	On (opt)	Off	Off
Peripherals				
M_TTCAN	On (opt) ^d	Retention	Off	Off
LIN	On (opt)	Retention	Off	Off
WDT/MCWD	On (opt)	On (opt)	WDT (opt)/MCWD off	Off
ADC	On (opt)	Retention	Off	Off
TCPWM	On (opt)	Off	Off	Off
SCB	On (opt)	On (opt)	Off	Off
GPIO	On	On/Freeze	Freeze	Hi-Z
Supplies and Reset				
XRES_L	Deassert	Deassert	Deassert	Assert
BOD	On	Slow ^e	Off	Off
POR	On	On	On	On
V _{DD} /V _{DDD} /V _{DDA}	On	On	On	On
Backup Domain				
RTC	On (opt)	On (opt)	On (opt)	On (opt)
WCO	On (opt)	On (opt)	On (opt)	On (opt)
LPECO ^f	On (opt)	On (opt)	On (opt)	On (opt)
Backup Registers	On	On (opt)	On (opt)	On (opt)

- a. See the device-specific document to find the supported number of pins to wake up the device.
b. Write buffers are not retained in DeepSleep mode.
c. See the device-specific document for LVD Slow (DeepSleep) specification.
d. When all M_TTCAN channels in a group are powered down, Message RAM will be powered off to save power.
e. See the device-specific document for the BOD Slow (DeepSleep) specification.
f. See the device-specific document to check if LPECO is supported.

17.5 Register List

Table 17-7. Register List

Register	Name	Description
PWR_CTL	Power Control	Power mode status register shows the current state and device ready status
PWR_HIBERNATE	Hibernate mode register	Controls various Hibernate mode entry/exit related options
PWR_HIB_DATA	Hibernate mode data register	Data register that is retained through a hibernate wakeup sequence
CM7_0/ CM7_1_SCS_SCR	Cortex-M7 system control register	Controls the CPU level Sleep/DeepSleep decisions on WFI/WFE instruction execution
CM0P_SCS_SCR	Cortex-M0+ system control register	Controls the CPU level Sleep/DeepSleep decisions on WFI/WFE instruction execution

18. Clocking System



The Traveo II family clocking system includes these resources:


- Three internal clock sources:
 - 8 MHz internal main oscillator (IMO)
 - Internal low-speed oscillators (ILO0/ILO1)
- Four external clock sources
 - External clock connected to one of two EXT_CLK inputs
 - External crystal oscillator (ECO)
 - External watch crystal oscillator (WCO)
 - Low-power external crystal oscillator (LPECO)¹
- One frequency-locked loop (FLL)
- Several phase-locked loops (PLL)

18.1 Block Diagram

Figure 18-1 gives a generic view of the clocking system in Traveo II family devices.

1. Refer to the device-specific datasheet to confirm whether LPECO is present.

Figure 18-1. Clocking System Block Diagram



Active Domain: Active Domain is the region for operating only during active power mode.


DeepSleep Domain: DeepSleep Domain is the region for operating only during Active mode and DeepSleep mode.

Hibernate Domain: Hibernate Domain is the region for operating in the All Power mode.

CSV: CSV mean the Clock Supervision. The Clock Supervision has the function of monitoring the clock.

MUX: MUX mean the multiplexer. The multiplexer select the output from among the input signals.

*1: LPECO and LPECO Prescaler has devices that are not implemented. Refer to device Specific Datasheet for checking if these functions is present or not.

LEGEND1 :


LEGEND2 :
D = the number of direct select paths
P = the number of PLLs
R = the number of clock roots (CLK_HF)
Note: Refer to device Specific Datasheet for checking about the number of D, P and R.

18.2 Clock Sources

18.2.1 Internal Main Oscillator (IMO)

The IMO is an accurate, high-speed internal (crystal-less) oscillator that produces a fixed frequency. Refer to the datasheet for the IMO frequency. The IMO output can be used by the PLL or FLL to generate a wide range of higher frequency clocks, or it can be used directly by the high-frequency root clocks. The IMO is enabled and disabled with CLK_IMO_CONFIG.ENABLE.

The IMO should not be disabled if it is the source of the clock path feeding high-frequency clock zero (CLK_HF0). CLK_HF0 is the source clock for the CPU. Therefore, if the IMO is in the source path of CLK_HF0, disabling the IMO disables the CPU.

The IMO is only available in Active and Sleep modes.

18.2.2 External Crystal Oscillator (ECO)

The Traveo II family contains an oscillator to drive an external crystal. Refer to the datasheet for the frequency of ECO. This clock source is built using an oscillator circuit. The circuit employs an external crystal that needs to be populated on the external crystal pins of the Traveo II device.

The ECO can be enabled using the CLK_ECO_CONFIG register fields.

18.2.2.1 ECO Trimming

The ECO supports a wide variety of crystals and ceramic resonators with the nominal frequency range specification described in the datasheet. The crystal manufacturer typically provides numerical values for parameters, namely the maximum drive level (D_L), the equivalent series resistance (ESR), the ideal shunt capacitance (C_0) and the parallel load capacitance (C_L). These parameters can be used to calculate the transconductance (g_m) and the maximum peak oscillation voltage across the crystal (V_P).

The formula of V_P is as follows. ECO does not support V_P less than 0.3 V.

$$\text{Max peak value: } V_P = \frac{\sqrt{\frac{D_L}{2\text{ESR}}}}{\pi f(C_0 + C_L)}$$

The formula of Transconductance (g_m) is as follows. The ECO block can deliver a maximum Transconductance (g_m) of 17.6 mA/V.

$$\text{Transconductance: } g_m > 20 \times \text{ESR} \times (2\pi \times f)^2 \times (C_0 + C_L)^2$$

The formula of Negative resistance (R_{neg}) is as follows. To guarantee crystal start up, negative resistance needs to be at least five times larger than ESR.

$$\text{Negative Resistance: } |R_{\text{neg}}| = \frac{g_m \times 4 \times C_L^2}{(2\pi \times f)^2 \times (4 \times C_L^2 + 4 \times C_L \times C_0)^2}$$

The ATRIM, WDTRIM, and FTRIM fields can be found in the CLK_ECO_CONFIG2 register. The ATRIM and WDTRIM settings control the trim for amplitude of the oscillator output. The FTRIM setting controls the filter used to prevent the third harmonic oscillation.

Amplitude trim (ATRIM) sets the crystal drive level when automatic gain control (AGC) is enabled (CLK_ECO_CONFIG.AGC_EN = 1). AGC must be enabled for $V_P < 1.1$ V and disabled for all other cases.

Watch dog trim (WDTRIM) sets the threshold on XO magnitude where the ECO block releases the clock to the system.

Filter trim (FTRIM) tunes the low-pass filter between the ECO_IN pin and the amplifier, which is used to prevent amplification of harmonics of the intended crystal frequency. The FTRIM value at 0x03 can be used; no other value is required.

WARNING: The V_P setting is critical for reliable system performance. If the V_P settings are too large (or AGC is disabled), the crystal can be damaged or suffer premature aging due to excessive power dissipation. If the V_P settings are too small, the oscillation will be more susceptible to system noise.

Based on the V_P value, the ATRIM, WDTRIM, and FTRIM values are set as shown in [Table 18-1](#).

Table 18-1. ATRIM, WDTRIM, and FTRIM Settings based on V_P

V _P [V]	AGC_EN	ATRIM	WDTRIM	FTRIM
0.5 ≤ V _P < 0.55	0x1	0x4	0x2	0x3
0.55 ≤ V _P < 0.60	0x1	0x5		
0.60 ≤ V _P < 0.65	0x1	0x6	0x3	
0.65 ≤ V _P < 0.70	0x1	0x7		
0.70 ≤ V _P < 0.75	0x1	0x8	0x4	
0.75 ≤ V _P < 0.80	0x1	0x9		
0.80 ≤ V _P < 0.85	0x1	0xA	0x5	
0.85 ≤ V _P < 0.90	0x1	0xB		
0.90 ≤ V _P < 0.95	0x1	0xC	0x6	
0.95 ≤ V _P < 1.00	0x1	0xD		
1.00 ≤ V _P < 1.05	0x1	0xE	0x7	
1.05 ≤ V _P < 1.10	0x1	0xF		
1.10 ≤ V _P	0x0	0x0-0xE ^a		

a. It is acceptable to select any value from 0x0 to 0xE.

The GTRIM sets up the trim for amplifier gain based on the calculated g_m , as shown in [Table 18-2](#).

Table 18-2. GTRIM Settings

g_m [mA/V]	GTRIM
$0 \leq g_m < 2.2$	0x00
$2.2 \leq g_m < 4.4$	0x01
$4.4 \leq g_m < 6.6$	0x02
$6.6 \leq g_m < 8.8$	0x03
$8.8 \leq g_m < 11$	0x04
$11 \leq g_m < 13.2$	0x05
$13.2 \leq g_m < 15.4$	0x06
$15.4 \leq g_m \leq 17.6$	0x07

RTRIM should be oscillator feedback resistor, as shown in [Table 18-3](#).

Table 18-3. RTRIM Settings

Nominal Frequency f [MHz]	RTRIM
$28.6 < f$	0x00
$23.33 < f \leq 28.6$	0x01
$16.5 < f \leq 23.33$	0x02
$f \leq 16.5$	0x03

First, set up the trim values based on [Table 18-1](#) through [Table 18-3](#) and then enable the ECO. After the ECO is enabled, the CLK_ECO_STATUS register can be checked to ensure it is ready.

18.2.3 External Clock (EXT_CLK)

The external clock can be sourced from a signal on a designated I/O pin. This clock can be used as the source clock for either the PLL or FLL, or can be used directly by the high-frequency clocks.

When manually configuring a pin as the input to the EXT_CLK, the drive mode of the pin must be set to high impedance digital to enable the digital input buffer. See the [I/O System chapter on page 246](#) for more details. Consult the device datasheet to determine the specific pin used for EXT_CLK.

The EXT_CLK function is bi-directional. See [18.3 Clock Generation](#) for more details.

18.2.4 Internal Low-speed Oscillator (ILO)

The two ILO blocks operate with no external components and output a stable clock. Refer to the datasheet for the frequency of the two ILOs. The ILO block is relatively low power and low accuracy. It is available in all power modes. If the ILO is to remain active in Hibernate mode, and across power-on-reset (POR) or brownout detect (BOD), CLK_ILO0_CONFIG.ILO0_BACKUP must be set.

The ILO blocks can be used as the clock source for:

- CLK_LF: CLK_LF in turn can be used as a source for the backup domain (CLK_BAK). CLK_BAK runs the Real Time Clock (RTC). This can be useful if you do not wish to populate a WCO. Although the ILO is not suitable as an RTC due to its poor accuracy, it can be used as a HIBERNATE wakeup source using the wakeup alarm facility in the RTC. In this case, CLK_ILO0_CONFIG.ILO0_BACKUP must be set.
- DSI_Mux: While the ILOs are routable through the DSI_MUX, there are no supported use cases for doing so.
- ILO0 is the clock for the watchdog timer (WDT) and DeepSleep CSV (clock supervision).
- ILO1 is used only if all of the following are true.
 - CLK_LF must be available in DeepSleep mode (otherwise use ECO output).
 - Clock supervision of CLK_LF is necessary (otherwise, use ILO0).
 - WCO is not available (otherwise, use WCO).

The ILO0 and the ILO1 are enabled/disabled with CLK_ILO0_CONFIG.ENABLE and CLK_ILO1_CONFIG.ENABLE respectively. It is recommended to always leave ILO0 enabled as it is the source of the WDT.

If the WDT is enabled, the only way to disable the ILO0 is to first clear WDT_CTL.WDT_LOCK and then clear CLK_ILO0_CONFIG.ENABLE. If the WDT_CTL.WDT_LOCK is set, any register write to disable the ILO0 will be ignored. Enabling the WDT will automatically enable the ILO0.

The MCU provides an opportunity to calibrate the ILOx by using the calibration counter, described in [Clock Calibration Counters chapter on page 210](#). For instance, the ECO can be used as a reference clock. This result can be used to determine how the ILOx needs to be adjusted. The ILO0 and ILO1 can be trimmed using the CLK_TRIM_ILO0_CTL and CLK_TRIM_ILO1_CTL registers.

18.2.5 Watch Crystal Oscillator (WCO)

The WCO is a highly accurate clock source. Refer to the datasheet for the frequency of WCO. It is the primary clock source for the RTC. The WCO can also be used as a source for CLK_LF.

The WCO can be enabled and disabled by setting BACKUP_CTL.WCO_EN for the backup domain. The WCO can also be bypassed and an external 32.768-kHz clock can be routed on a WCO output pin. This is done by setting BACKUP_CTL.WCO_BYPASS for the backup domain.

It is possible to improve the accuracy of the RTC by calibrating the WCO. WCO is routable through the DSI_MUX and can then be routed as the source of the FLL.

18.2.6 ECO Prescaler

ECO Prescaler divides the ECO and creates a clock that can be used with CLF_LF clock. This feature is only available during Active and Sleep mode. It cannot be used during DeepSleep or Hibernate modes.

The division function has a 10-bit integer divider and 8-bit fractional divider. This function is configured using the CLK_ECO_PRESCALE register.

18.2.7 LPECO

The Traveo II family has a low-power external crystal oscillator (LPECO). Refer to the datasheet for the frequency of LPECO.

LPECO can be thought of as an ECO that operates during low-power modes. LPECO replaces the function of WCO for the real-time clock (RTC). This means LPECO must continue to operate during XRES.

The LPECO can be controlled using the BACKUP_LPECO_CTL and BACKUP_LPECO_STATUS register fields.

18.2.8 LPECO Prescaler

The LPECO prescaler divides the LPECO, and creates a clock that can be used with the RTC. This feature is available during Active, DeepSleep, and Hibernate modes, and XRES. The LPECO prescaler is a fractional clock divider. See *Traveo II Body Controller High Registers TRM* for more details.

The LPECO prescaler can be controlled using the BACKUP_LPECO_PRESCALE register fields.

18.3 Clock Generation


This section explains phase locked loop (PLL) and frequency locked loop (FLL) implemented in the Traveo II family.

Traveo II family has two types of PLLs; PLL without SSCG and fractional operation, and PLL with SSCG and fractional operation. Traveo II family has one type of FLL.

18.3.1 PLL without SSCG and Fractional Operation

Refer to the datasheet to identify where this PLL type is used. The datasheet also specifies the frequency range that can be input to the PLL and the frequency range that the PLL can output. This makes it possible to use the IMO or other clock to generate much higher clock frequencies for the rest of the system. [Figure 18-2](#) shows the block diagram of a PLL without SSCG and fractional operation.

Figure 18-2. PLL without SSCG and Fractional Operation



The PLL is configured following these steps:

Note: Fref is the input frequency of the PLL; that is, the frequency of the input clock such as 8 MHz for the IMO.


1. Determine the desired reference clock frequency (Fref) and desired output frequency (PLL_OUT). Calculate the reference (Q), feedback (P), and output (OUTPUT_DIV) dividers subject to the following constraints:
 - a. PFD frequency (phase detector frequency). $F_{pfd} = F_{ref} / Q$. There may be multiple reference divider values that meet this constraint.
 - b. VCO frequency. $VCO = F_{pfd} \times P$. There may be multiple feedback divider values that meet this constraint with different REFERENCE_DIV choices.
 - c. Output frequency. $PLL\_OUT = VCO / OUTPUT\_DIV$. It may not be possible to get the exact desired frequency due to granularity; therefore, consider the frequency error of the two closest choices.
 - d. Choose the best combination of divider parameters depending on the application.
2. Program the divider settings in the appropriate CLK_PLL_CONFIGx register. Do not enable the PLL on the same cycle as configuring the dividers. Do not change the divider settings while the PLL is enabled.

3. Enable the PLL (CLK_PLL_CONFIGx.ENABLE = 1). Wait at least 1 μ s for PLL circuits to start.
4. Wait until the PLL is locked before using the output. By default, the PLL output is bypassed to its reference clock and will automatically switch to the PLL output when it is locked. This behavior can be changed using CLK_PLL_CONFIGx.BYPASS_SEL. The status of the PLL can be checked by reading CLK_PLL_STATUSx. This register contains a bit indicating the PLL has locked. It also contains a bit indicating if the PLL lost the lock status.

18.3.2 PLL with SSCG and Fractional Operation (400-MHz PLL)

Refer to the datasheet to identify where this PLL type is used. The datasheet also specifies the frequency range that can be input to the PLL and the frequency range that the PLL can output. This makes it possible to use the IMO or other clock to generate much higher clock frequencies for the rest of the system. Figure 18-3 shows the block diagram of a PLL with SSCG and fractional operation. This type of PLL is configured in the CLK_PLL400Mx_CONFIG register and the status is confirmed in the CLK_PLL400Mx_STATUS register.

Figure 18-3. PLL with SSCG and Fractional Operation




18.3.2.1 Spread Spectrum Clock Generation (SSCG)

Spread spectrum clock generation (SSCG) is a method by which the energy contained in the narrow band of a clock source is spread over a wider band in a controlled manner, thus reducing the peak spectral amplitude of the fundamental and the harmonics to lower the radiated emission from the clock source. This is achieved by modulating the clock frequency with a waveform. The configuration of the SSCG uses the CLK_PLL400Mx_CONFIG3 register.

Table 18-4 describes the three SSCG parameters:

Table 18-4. SSCG Parameters

Parameters	Description
Modulation Rate	Modulation rate (MR) is the rate (in Hz) at which the energy of the clock source is distributed over the band of frequencies around the output clock frequency. Modulation rate must be much lower than the source clock frequency, and must be above the audio frequency range.
Modulation Depth	Modulation depth (also known as deviation) is the frequency range over which the clock changes while varying at the modulation rate. It is specified as a percentage (%), which is the ratio of the bandwidth of frequency excursion (ΔF) to the source clock frequency. This determines the amount of peak EMI reduction achievable. Generally, the larger the modulation depth, the greater the EMI reduction.
Modulation Type	Modulation type (or spreading mode) specifies the relationship of the frequency deviation of the modulated clock relative to the non-modulated clock. For general SSCG, center, up, and down spread are available, but this PLL only supports down spread. Down spreading is where the maximum frequency of the spread spectrum clock is the same as that of the non-modulated clock. In a down-spread system, the output clock varies between $(F_c - \Delta F)$ and F_c at the modulation rate and following the modulation profile.  It can be represented as "Fout = Fc - ΔF".

Fout = the modulated output clock frequency

Fc = the source or carrier frequency

ΔF = total frequency deviation (min to max)

18.3.2.2 Fractional Operation

Fractional operation is a function that provides an output frequency that is a fractional multiple of the input frequency. When using fractional operation, the following formula holds.


$$PLL\_OUT = (F_{ref} / Q) \times (P + \text{Frac\_div}) / \text{OUTPUT\_DIV}$$

The configuration of fractional operation uses the PLL400_CONFIG2 register. Frac_div is the value set by CLK_PLL400Mx_CONFIG2.FRAC_DIV divided by 2^{24} . While the fractional divider has 24 bits, accuracy is only guaranteed for the upper 21 bits.

18.3.3 Frequency Locked Loop (FLL)

The Traveo II family device contains one FLL, which resides on CLK_PATH0. Refer to the datasheet for the frequency range that can be input to the FLL and the frequency range that the FLL can output. This makes it possible to use the IMO to generate much higher clock frequencies for the rest of the system. Figure 18-4 shows the block diagram of FLL.

Figure 18-4. FLL Block Diagram



The FLL is similar in purpose to a PLL but is not equivalent; there are some differences:

- FLL can start up (lock) much faster than the PLL.
- It consumes less current than the PLL.
- FLL does not lock the phase. At the heart of the FLL is a current-controlled oscillator (CCO). The output frequency of this CCO is controlled by adjusting the trim of the CCO; this is done in hardware.
- FLL can produce a clock with good duty cycle through its divided clock output.
- FLL reference clock can be the WCO, IMO (8 MHz), or any other periodic clock source.

The CCO can output a stable frequency in the 48 MHz to 160 MHz range. This range is divided into five sub-ranges as shown by [Table 18-5](#).

Table 18-5. CCO Frequency Range

CCO Range	0	1	2	3	4
Fmin	48 MHz	64 MHz	85 MHz	113 MHz	150 MHz
Fmax	64 MHz	85 MHz	113 MHz	150 MHz	200 MHz

Note: The output of the CCO has an option to enable a divide by two or not. For this device, the divide by two must always be enabled. The output range of the FLL is shown in [Table 18-5](#).

Within each range, the CCO output is controlled via a 9-bit trim field. This trim field is updated via hardware based on the control algorithm described here.

A reference clock must be provided to the FLL. This reference clock is typically the IMO, but can be many different clock sources. The FLL compares the reference clock against the CCO clock to determine how to adjust the


CCO trim. Specifically, the FLL will count the number of CCO clock cycles inside a specified window of reference clock cycles. The number of reference clock cycles to count is set by `CLK_FLL_CONFIG2.FLL_REF_DIV`.

After the CCO clocks are counted, they are compared against an ideal value and an error is calculated. The ideal value is programmed into `CLK_FLL_CONFIG.FLL_MULT`.

As an example, the reference clock is the IMO (8 MHz), the desired CCO frequency is 100 MHz, the value for `CLK_FLL_CONFIG2.FLL_REF_DIV` is set to 146. This means that the FLL will count the number of CCO clocks within 146 clock periods of the reference clock. In one clock cycle of the reference clock (IMO), there should be $100 / 8 = 12.5$ clock cycles of the CCO. Multiply this number by 146 and the value of `CLK_FLL_CONFIG.FLL_MULT` should be 1825.

If the FLL counts a value different from 1825, it attempts to adjust the CCO such that it achieves 1825 the next time it counts. This is done by scaling the error term with `CLK_FLL_CONFIG3.FLL_LF_IGAIN` and `CLK_FLL_CONFIG3.FLL_LF_PGAIN`. [Figure 18-5](#) shows how the error (err) term is multiplied by `FLL_LF_IGAIN` and `FLL_LF_PGAIN` and then summed with the current trim to produce a new trim value for the CCO. `CLK_FLL_CONFIG4.CCO_LIMIT` can be used to put an upper limit on the trim adjustment; this is not needed for most situations.

Figure 18-5. FLL Error Correction Diagram



The FLL determines if it is “locked” by comparing the error term with CLK_FLL_CONFIG2.LOCK_TOL.

When the error is less than LOCK_TOL the FLL is considered locked.

After each adjustment to the trim the FLL can be programmed to wait a certain number of reference clocks before doing a new measurement. The number of reference clocks to wait is set in the CLK_FLL_CONFIG3.SETTLING_COUNT. It is recommended to set this such that the FLL waits ~1 μ s before a new count. Therefore, if the 8 MHz IMO is used as the reference this field should be programmed to ‘8’.

When configuring the FLL there are two important factors that must be considered: lock time and accuracy. Accuracy is the closeness to the intended output frequency. These two numbers are inversely related to each other via the value of CLK_FLL_CONFIG2.FLL_REF_DIV.

Higher CLK_FLL_CONFIG2.FLL_REF_DIV values lead to higher accuracy, whereas lower CLK_FLL_CONFIG2.FLL_REF_DIV values lead to faster lock times.

In the example used previously the 8 MHz IMO was used as the reference, and the desired FLL output was 100 MHz. For that example, there are 12.5 CCO clocks in one reference clock. If the value for CLK_FLL_CONFIG2.FLL_REF_DIV is set to ‘8’ then CLK_FLL_CONFIG.FLL_MULT must be set to either ‘13’ or ‘12’. This will result in a CCO output of either 96 MHz or 104 MHz, and an error of 4 percent from the desired 100 MHz. Therefore, the best way to improve this is to increase CLK_FLL_CONFIG2.FLL_REF_DIV. However, the larger CLK_FLL_CONFIG2.FLL_REF_DIV is, the longer each measurement cycle takes, thus increasing the lock time. In this example, CLK_FLL_CONFIG2.FLL_REF_DIV was set to 146. This means each measurement cycle takes $146 \times (1/8 \text{ MHz}) = 18.25 \mu\text{s}$, whereas when CLK_FLL_CONFIG2.FLL_REF_DIV is set to 1, each measurement cycle takes $1 \times (1/8 \text{ MHz}) 0.125 \mu\text{s}$.

Another issue with lower CLK_FLL_CONFIG2.FLL_REF_DIV values is that the minimum CLK_FLL_CONFIG2.LOCK_TOL is 1, so the

output of the CCO can have an error of ± 1 . In the example where CLK_FLL_CONFIG2.FLL_REF_DIV = 1 and CLK_FLL_CONFIG.FLL_MULT = 13, the MULT value can really be 12, 13, or 14 and still be locked.

This means the output of the FLL may vary between 96 and 112 MHz, which may not be desirable.

A choice must be made between faster lock times and more accurate FLL outputs. The biggest change is the value of REF_DIV. The CLK_FLL_STATUS register checks the status of FLL.

18.4 Clock Trees

Traveo II family clocks are distributed throughout the device, as shown in Figure 18-1. The clock trees are described in this section:

- Path Clocks (CLK_PATH)
- High-Frequency Root Clocks (CLK_HF)
- Low-Frequency Clock (CLK_LF)
- Timer Clock (CLK_TIMER)

18.4.1 Path Clocks

The Traveo II family device has several clock paths: CLK_PATH0 contains the FLL, CLK_PATH1 to CLK_PATH<P+1> contains the PLLs, and CLK_PATH<P+D> is a direct connection to the high-frequency root clocks. Note that the FLL and PLL can be bypassed if they are not needed. These paths are the input sources for the high-frequency clock roots.

Each clock path has a mux to determine the source clock for that path. This configuration is done in the CLK_PATH_SELECTx register. Table 18-6 shows the clock path source selections.

Table 18-6. Clock Path Source Selections

Name	Description
PATH_MUX	Selects the source for CLK_PATH
	0: IMO
	1: EXT_CLK
	2: ECO
	3: Reserved. Do not use
	4: DSI_MUX
	5: LPECO Prescaler
	6-7: Reserved. Do not use.

The DSI mux is configured through the CLK_DSI_SELECTx register. Table 18-7 shows the DSI mux source selections.

Table 18-7. DSI Mux Source Selections

Name	Description
DSI_MUX	Selects the source for DSI_MUX
	0-15: Reserved. Do not use
	16: ILO0
	17: WCO
	18: Reserved. Do not use
	19: Reserved. Do not use
	20: ILO1
	21-31: Reserved. Do not use

18.4.2 High-Frequency Root Clocks

Traveo II family has several high-frequency root clocks (CLK_HF). Each CLK_HF has a particular destination on the device; refer to the datasheet for details.

Each high-frequency root clock has a mux to determine its source. This configuration is done in the CLK_ROOT_SELECTx register. The number of CLK_PATH depends on the device. Refer to the datasheet for details.

Each CLK_HF has a predivider, which is set in the CLK_ROOT_SELECTx register.

CLK_HF0 is always enabled because it is the source of the CPU clock. Other CLK_HF can be enabled or disabled using CLK_ROOT_SELECTx.ENABLE.

Table 18-8. CLK_HF Divider Selection

Name	Description
ROOT_DVI	Selects Predivider value for the clock root
	0: No Divider
	1: Divide clock by 2
	2: Divide clock by 4
	3: Divide clock by 8

18.4.3 Low-Frequency Root Clocks

The low-frequency clock (CLK_LF) in Traveo II family has four input options: ILO0, WCO, ILO1, and ECO_Prescaler. CLK_LF is the source for the multi-counter watchdog timers

(MCWDT), and can also be a source for the RTC. The source of CLK_LF is set in CLK_SELECT.LFCLK_SEL.

Table 18-9. LFCLK Input Selection Bits LFCLK_SEL

Name	Description
LFCLK_SEL	LFCLK input clock selection
	0: ILO0
	1: WCO
	2: Reserved. Do not use
	3: Reserved. Do not use
	4: ILO1
	5: ECO Prescaler
	6: LPECO Prescaler
	7: Reserved. Do not use

18.4.4 Timer Clock

The timer clock (CLK_TIMER) can be used as a clock source for the CPU SYSTICK timer. The source for CLK_TIMER can only be the IMO.

The CPU SYSTICK timer operation is configured by the CPUSS_SYSTICK_CTL register.

Figure 18-6. Connection to CLK_TIMER




Table 18-10. Timer Clock Selection

Name	Description
TIMER_SEL	Timer clock selection
	0: CLK_IMO
	1: CLK_HF0

18.4.5 Clock Output Function

The EXT_CLK terminal is bi-directional. The EXT_CLK terminal can input clock to Traveo II; it is also possible to output the Traveo II internal clock. To use EXT_CLK as an input, configure the HSIOM to select EXT_CLK on a supported pin, and set that GPIO to the High Impedance Digital drive mode. To use EXT_CLK as an output, configure the HSIOM to select EXT_CLK on a supported pin, and set that GPIO to a mode setting capable of driving the selected clock frequency.

The clock source available on EXT_CLK is CLK_HF3, which can choose any internal clock source including ECO. Note that CLK_HF3 is a shared resource; changing the Predivider setting impacts all the connections.

18.5 CLK_HF Distribution

Traveo II has several CLK_HFs, which connect CLK_FAST_0, CLK_FAST_1, CLK_MEM, CLK_SLOW, CLK_PERI, CLK_GR and PCLK. CLK_FAST_0, CLK_FAST_1, CLK_MEM, CLK_SLOW, CLK_PERI and CLK_GR are source clocks of the CPU subsystem, SRSS, and some peripheral functions. Also, a part of CLK_HF connects to CSV. Refer to the datasheet for the connection relationship between CLK_HF and the source clocks.

18.5.1 CLK_MEM

CLK_MEM clocks the CPUSS fast infrastructure. This clock is a divided version of CLK_HF. The divider for this clock is set in the CPUSS_MEM_CLOCK_CTL register.

18.5.2 CLK_SLOW

CLK_SLOW is the source clock for the Cortex-M0+. This clock is a divided version of CLK_MEM. The divider for this clock is set in CPUSS_SLOW_CLOCK_CTL.INT_DIV.

18.5.3 CLK_FAST_0 and CLK_FAST_1

CLK_FAST_x clocks the CM7. This clock is a divided version of CLK_HF1; the divider for this clock is set in the CPUSS_FAST_0_CLOCK_CTL and CPUSS_FAST_1_CLOCK_CTL registers.

18.5.4 CLK_PERI

CLK_PERI is the source clock for all programmable peripheral clock dividers. It is a divided version of CLK_HF. The divider for this clock is set in CPUSS_PERI_CLOCK_CTL.INT_DIV.

18.5.5 PCLK

PCLK is the source of peripheral functions: CAN FD, FLEX-RAY, LIN, TCPWM, SCB, and SAR ADC. For details, [Peripheral Clock Dividers on page 207](#).

PCLK are of different kinds; it connects to various peripheral functions. Refer to the Peripheral Clocks section in the datasheet for more information.

18.5.6 CLK_GR

CLK_GR is a clock input to peripheral functions. It is grouped by the clock gater. Each GR_CLK is divided by the source clock and generated. The divider for the clock is configured in PERI_GRx_CLOCK_CTL.INT_DIV. CLK_GR can be enabled or disabled. The configuration is configured using the PERI_GRx_SL_CTL register.

18.5.7 CLK_TRC_DBG

CLK_TRC_DBG is a clock used for debugging. This clock is generated by dividing a clock source. See the device datasheet for the clock source. Clock division uses the CPUSS_TRC_DBG_CLOCK_CTL register.


18.6 Peripheral Clock Dividers

Traveo II family peripherals such as SCBs and TCPWMs require a clock. These peripherals can be clocked only by a peripheral clock divider.

The Traveo II family has several peripheral clocks (PCLK) and peripheral clock dividers. Peripheral clock dividers can be 8-bit, 16-bit, 16.5-bit (16 integer bits and five fractional bits) and 24.5-bit (24 integer bits and five fractional bits). The output of any of these dividers can be routed to any peripheral. The divider also outputs a signal to divide (enable) the clock signal.

Refer to the datasheet for the number of dividers and assignment of peripheral clocks. [Figure 18-7](#) shows the peripheral clock divider diagram.

Figure 18-7. Peripheral Clock Divider



18.6.1 Fractional Clock Dividers

Fractional clock dividers allow the clock divisor to include a fraction of 0..31/32. For example, a 16.5-bit divider with an integer divide value of 3 generates a 16-MHz clock from a 48-MHz CLK_PERI. A 24.5-bit divider with an integer divide value of 4 generates a 12-MHz clock from a 48-MHz CLK_PERI. A 24.5-bit divider with an integer divide value of 3 and a fractional divider of 16 generates a $48 / (3 + 16/32) = 48 / 3.5 = 13.7$ MHz clock from a 48-MHz CLK_PERI. Not all 13.7-MHz clock periods are equal in size; half of them will be three CLK_PERI cycles and half of them will be two CLK_PERI cycles.

Fractional dividers are useful when a high-precision clock is required (such as a UART/SPI serial interface). Fractional dividers are not used when a low jitter clock is required, because the clock periods have a jitter of one CLK_PERI cycle.

18.6.2 Peripheral Clock Divider Configuration

The peripheral clock dividers are configured using registers from the Peripheral block; specifically PERI_PCLK_GRx_CLOCK_CTLy, PERI_PCLK_GRx_DIV_CMD, PERI_PCLK_GRx_DIV_8_CTLy, PERI_PCLK_GRx_DIV_16_CTLy, PERI_PCLK_GRx_DIV_16_5_CTLy, and PERI_PCLK_GRx_DIV_24_5_CTLy registers.

First the clock divider needs to be configured. This is done via the PERI_PCLK_GRx_DIV_8_CTLy, PERI_PCLK_GRx_DIV_16_CTLy, PERI_PCLK_GRx_DIV_16_5_CTLy, and PERI_PCLK_GRx_DIV_24_5_CTLy registers. The number of each divider can be found in the Peripheral Clock Dividers section of the datasheet. Dividers not listed in “Divider Type”

are not implemented. The divider selection is determined by the group number and the instance number. For example, use the PERI_PCLK_GR1_DIV_16_CTL3 register to configure the third 16-bit divider with Group 1.


After the divider is configured, use the PERI_PCLK_GRx_DIV_CMD register to enable the divider. This is done by setting the PERI_PCLK_GRx_DIV_CMD.DIV_SEL to the divider number you want to enable, and setting the PERI_PCLK_GRx_DIV_CMD.TYPE_SEL to the divider type. For example, if you wanted to enable the 0th 24.5-bit divider, write '0' to PERI_PCLK_GRx_DIV_CMD.DIV_SEL and '3' to PERI_PCLK_GRx_DIV_CMD.TYPE_SEL. If you wanted to enable the tenth 16-bit divider, write '10' to PERI_PCLK_GRx_DIV_CMD.DIV_SEL and '1' to PERI_PCLK_GRx_DIV_CMD.TYPE_SEL. See the *Traveo II Body Controller High Registers TRM* for more details.

To connect a peripheral to a specific divider, the PERI_PCLK_GRx_CLOCK_CTLy register is used. There is a PERI_PCLK_GRx_CLOCK_CTLy register corresponding to each PCLK mentioned in the Peripheral Clock section of the datasheet. For example, to select the twelfth 16-bit divider for PCLK with Group 1 and Output 29, configure DIV_SEL to '12' and TYPE_SEL to '1' in the PERI_PCLK_GR1_CLOCK_CTL29 register. Also, the 'Output' order matches the order of y of the PERI_PCLK_GRx_CLOCK_CTLy register.

18.6.2.1 Phase Aligning Dividers


For specific use cases, it is required to generate clocks that are phase aligned. For example, consider the generation of two gated clocks at 24 and 12 MHz, both of which are derived from a 48-MHz CLK_PERI. If phase alignment is not considered, the generated gated clocks can appear as follows.

Figure 18-8. Non Phase-Aligned Clock Dividers




These clock signals may or may not be acceptable, depending on the logic functionality implemented on these two clocks. If the two clock domains communicate with each other, and the slower clock domain (12 MHz) assumes that each high/'1' pulse on its clock coincides with a high/'1' phase pulse in the higher clock domain (24 MHz), the phase misalignment is not acceptable. To address this, it is possible to have PCLK dividers produce clock signals that are phase-aligned with any of the other (enabled) clock dividers. Therefore, if (enabled) divider x is used to generate the 24-MHz clock, divider y can be phase-aligned to divider x and used to generate the 12-MHz clock. The generated gated clocks can appear as follows.

Figure 18-9. Phase-Aligned Clock Dividers



Phase alignment also works for fractional divider values. If (enabled) divider x is used to generate the 38.4 MHz clock (divide by $1 \frac{8}{32}$), divider y can be phase-aligned to divider x and used to generate the 19.2-MHz clock (divide by $2 \frac{16}{32}$). The generated gated clocks can appear as follows.

Figure 18-10. Phase-Aligned Fractional Dividers



Divider phase alignment requires that the divider to which it is phase aligned is already enabled. This requires the dividers to be enabled in a specific order. This order can be represented by a divider dependency graph.

Phase alignment is implemented by controlling the start moment of the divider counters in hardware. When a divider is enabled, the divider counters are set to '0'. The divider counters will only start incrementing from '0' to the programmed integer and fractional divider values when the divider to which it is phase aligned has an integer counter value of '0'.

Note that the divider and clock multiplexer control register fields are all retained during DeepSleep power mode. However, the divider counters that are used to implement the integer and fractional clock dividers are not. These counters are set to '0' during DeepSleep power mode. Therefore, when transitioning from DeepSleep to Active power mode, all dividers (and clock signals) are enabled and phase-aligned by design.

Phase alignment is accomplished by setting the `PERI_PCLK_GRx_DIV_CMD.PA_DIV_SEL` and `PERI_PCLK_GRx_DIV_CMD.PA_DIV_TYPE` before enabling the clock. For example, to align the fourth 8-bit divider to the third 16-bit divider, set `PERI_PCLK_GRx_DIV_CMD.DIV_SEL` to '4', `PERI_PCLK_GRx_DIV_CMD.TYPE_SEL` to '0', `PERI_PCLK_GRx_DIV_CMD.PA_DIV_SEL` to '3', and `PERI_PCLK_GRx_DIV_CMD.PA_TYPE_SEL` to '1'.

18.7 Clock Calibration Counters

A feature of the clocking system in Traveo II family is built-in hardware calibration counters. These counters can be used to compare the frequency of two clock sources against one another. The primary use case is to take a higher accuracy clock such as the ECO and use it to measure a lower accuracy clock such as the ILOx. The result of this measurement can then be used to trim the ILOx.

There are two counters: Calibration Counter 1 is clocked off of Calibration Clock 1 (generally the high-accuracy clock) and it counts down; Calibration Counter 2 is clocked off of Calibration Clock 2 and it counts up. When Calibration Counter 1 reaches 0, Calibration Counter 2 stops counting up and its value can be read. From that value the frequency of Calibration Clock 2 can be determined with the following equation.

$$\text{CalibrationClock2frequency} = \frac{\text{Counter2FinalValue}}{\text{Counter1InitialValue}} \times \text{CalibrationClock1Frequency}$$

For example, if Calibration Clock 1 = 8 MHz, Counter 1 = 1000, and Counter 2 = 5

Calibration Clock 2 Frequency = $(5/1000) \times 8 \text{ MHz} = 40 \text{ kHz}$.

Calibration Clock 1 and Calibration Clock 2 are selected with the CLK_OUTPUT_FAST register. All clock sources are available as a source for these two clocks. CLK_OUTPUT_SLOW is also used to select the clock source.

Calibration Counter 1 is programmed in CLK_CAL_CNT1. Calibration Counter 2 can be read in CLK_CAL_CNT2.

When Calibration Counter 1 reaches 0, CLK_CAL_CNT1.CAL_COUNTER_DONE is set.

18.8 Clock Supervision (CSV)

18.8.1 Overview

This section provides an overview of the clock supervision features.

- The CSV circuit checks whether the frequency of the monitored clock is within the allowed frequency window. If the monitored clock stops, or fails to start, it is detected as a low frequency.
- All CLK_HFs have the CSV. All CLK_HF CSVs use the same reference clock (CLK_REF_HF). The reference clock is a selection of one of the Active clock sources in the CSV_REF_SEL register. Typically, the IMO is selected (default).
- Note that ILO0 is supervised both by CLK_LF and CLK_REF_HF, CLK_LF is needed for DeepSleep supervision and CLK_REF_HF is needed for accuracy (while Active).
- CSV_REF monitors CLK_REF_HF with CLK_ILO0.
- All CSV_HFs and CSV_REF are in the Active domain.
- There are two CSVs (CSV_LF and CSV_ILO) in the DeepSleep domain. CSV_LF is used to monitor the selected CLK_LF clock with ILO0. CSV_ILO is used to monitor ILO0 with CLK_LF. ILO1 is provided to enable clock supervision of ILO0 during DeepSleep when the WCO is not being used.
- The Hibernate domain, where the RTC is located, does not have CSVs.

Figure 18-11 gives an overview of the location of CSVs for the Traveo II.

Refer to the datasheet for details about the relationship between the monitored clock and reference clock for each CSV component.





18.8.2 CSV Operation

The basic operation principle of the CSV circuit is as follows:

- The monitored clock generates a Monitor event (Period), and reference clock generates a Lower and Upper limit.
- The Monitor event is compared against a Lower or Upper limit
- An error is reported if Monitor event $\leq$ Lower limit, or Monitor event $>$ Upper limit


Figure 18-12 shows an example of the CSV operation signal.

Figure 18-11. CSV Diagram




*1: LPECO and LPECO Prescaler has devices that are not implemented. Refer to device Specific Datasheet for checking if these functions is present or not.


Figure 18-12. An Example of CSV Operation Signal



(1) The error is not reported: $\text{Monitor Event} > \text{Lower limit}$, and $\text{Monitor Event} \leq \text{Upper limit}$



(2) The error is reported: $\text{Monitor Event} \leq \text{Lower limit}$



(3) The error is reported: $\text{Monitor Event} > \text{Upper limit}$

Key points for the CSV operation:

- “Period” has the following formula.

$$\text{Period} = \text{Target} / (\text{Reference frequency} / \text{Monitor frequency})$$
- “Lower_limit” and “Upper_limit” can be expressed as follows.

$$\text{Lower\_limit} = \text{Target} - \text{Required accuracy}/2$$

$$\text{Upper\_limit} = \text{Target} + \text{Required accuracy}/2$$
- Increasing Target increases accuracy and latency. For example, for an accuracy of 1% the Target must be at least 200.
- If the two clocks are asynchronous (typical) then there will be a one-cycle variation of Monitor Event periods.
- The frequency window needs to account for the maximum clock tolerance on both clocks.
- Lower_limit must be at least one less than Upper_limit.
- All CSVs are initially off and require configuration before enabling.
- The Active domain CSVs are automatically stopped during DeepSleep.
 - After wakeup they will automatically restart
 - Each CSV has a software programmable startup time.
- All Active domain CSVs can either generate a Reset or Fault report.
 - The CLK_HF0 CSV must use reset because the fault structure runs on CLK_HF0.
 - All other CSVs should use a fault report to allow software to shut down.
 - A fault report will result in an interrupt.
- The DeepSleep domain CSVs operate during Active and DeepSleep.
 DeepSleep domain CSVs can only report faults (no reset option).
 - A CSV error detection will wake up the system (if needed), which enables fault reporting
 - The fault report will result in an interrupt (no direct interrupt from CSV)
- All CSVs are enabled independent of the monitored clock; therefore:
 - Software should disable the CSV before stopping or reconfiguring the monitored clock (to avoid a false error detection)
 - The CSV needs to be reconfigured accordingly and restarted after the monitored clock is re-started
- The CSV_REF_SEL register elects a source to be used as the reference clock for CSV in the Active domain. The registers to configure the CSV function are as follows. These registers can enable CSV, and can configure an action when CSV is activated.
 - CSV_HF_CSVx_REF_CTL
 - CSV_REF_CSVx_REF_CTL
 - CSV_LF_CSVx_REF_CTL
 - CSV_ILO_CSVx_REF_CTL
 The following registers can configure upper limit and lower limit.
 - CSV_HF_CSVx_REF_LIMIT
 - CSV_REF_CSVx_REF_LIMIT
 - CSV_LF_CSVx_REF_LIMIT
 - CSV_ILO_CSVx_REF_LIMIT
 The following registers can configure PERIOD (time period).
 - CSV_HF_CSVx_MON_CTL
 - CSV_REF_CSVx_MON_CTL
 - CSV_LF_CSVx_MON_CTL
 - CSV_ILO_CSVx_MON_CTL
- When CSV violation occurs, SRSS_CSV (0x18) is indicated in the IDX bit of the FAULT_STRUCT_STATUS register. The CSV fault report is captured in FAULT_STRUCT_DATA.

18.9 Registers

Registers related to the clock system are shown in [Table 18-11](#).

Table 18-11. Clock System Registers

Register	Name	Description
CLK_DSI_SELECTx	Clock DSI Select Register	Configures DSI mux in the clock generation path. Each path has its own copy of this register. See 18.4.1 Path Clocks for DSI signal connectivity list.
CLK_OUTPUT_FAST	Fast Clock Output Select Register	Two signals can be selected to enable comparison of clocks. Fast clocks output is only observable in (LP)Active/(LP)Sleep and are clamped low during DeepSleep.
CLK_OUTPUT_SLOW	Slow Clock Output Select Register	Two signals can be selected to enable comparison of clocks. Slow clock outputs are observable in (LP)Active/(LP)Sleep/DeepSleep.
CLK_CAL_CNT1	Clock Calibration Counter 1	This register is a calibration counter that counts down by CLK_FAST selected by the CLK_OUTPUT_FAST register.
CLK_CAL_CNT2	Clock Calibration Counter 2	This register is a calibration counter that is counted up by CLK_FAST selected by the CLK_OUTPUT_FAST register.
CLK_PATH_SELECTx	Clock Path Select Register	Selects a source for clock path. The output of this mux can be used as the root of a clock tree. If there is a PLL on the path, this mux output is the PLL reference clock. The related PLL register contains a mux to select whether the clock path uses the PLL output or is bypassed to the PLL reference clock.
CLK_ROOT_SELECTx	Clock Root Select Register	Selects a root for a high-frequency clock tree and DSI input. Each clock root has a copy of this register.
CLK_SELECT	Clock selection Register	Clock source selection register.
CLK_ILO0_CONFIG	ILO0 Configuration	Configuration register for ILO0.
CLK_ILO1_CONFIG	ILO1 Configuration	Configuration register for ILO1.
CLK_IMO_CONFIG	IMO Configuration	Internal high-speed R/C oscillator configuration register. Note that this oscillator comes up active on power up. The oscillator provides the primary system clock (HFCLK) on power up until firmware configures differently. This oscillator is also used before system start to count out power up delays.
CLK_ECO_CONFIG	ECO Configuration Register	Internal high-speed oscillator configuration register for the external-crystal.
CLK_ECO_PRESCALE	ECO Prescaler Configuration Register	Fractional prescaler value to bring down the ECO frequency to 32768 Hz if used as CLK_LF. Do not divider settings while ECO Prescaler is enabled or enabling.
CLK_ECO_STATUS	ECO Status Register	Error and status indications.
CLK_FLL_CONFIG	FLL Configuration Register	This register contains frequency lock loop (FLL) configuration. FLL circuit settings should not be changed while it is a selected clock (connected to logic). This prevents clock glitches that can crash the logic.
CLK_FLL_CONFIG2	FLL Configuration Register2	
CLK_FLL_CONFIG3	FLL Configuration Register3	
CLK_FLL_CONFIG4	FLL Configuration Register4	
CLK_FLL_STATUS	FLL Status Register	This register indicates status for the FLL. This register is synchronized during an AHB read transaction. This causes a number of wait-states to be inserted in the transaction depending on the frequency ration between system and FLL frequency.

Table 18-11. Clock System Registers

Register	Name	Description
CLK_ECO_CONFIG2	ECO Configuration Register 2	Internal high-speed oscillator configuration register for the external-crystal.
CLK_PLL_CONFIGx	PLL Configuration Register	This register contains PLL configuration. Each PLL has a copy of this register. PLL circuit settings should not be changed while it is a selected clock (connected to logic).
CLK_PLL_STATUSx	PLL Status Register	This register indicates status for the PLL. Each PLL has a copy of this register.
CLK_ILO0_TRIM_CTL	ILO0 TRIM Register	This register configures trim of ILO0 frequency.
CLK_ILO1_TRIM_CTL	ILO1 TRIM Register	This register configures trim of ILO1 frequency.
PERI_GRx_CLOCK_CTL	Divider Control Register	This register configures division of CLK_GR.
PERI_GRx_SL_CTL	Slave Control	This register controls whether CLK_GR is enabled or disabled.
PERI_PCLK_GRx_CLOCK_CTLy	Divider Clock Control Register	This register configures DIV_SEL and TYPE_SEL of PCLKs.
PERI_PCLK_GRx_DIV_8_CTLy	Divider Control Register (for 8.0 divider)	This register controls the 8-bit divider.
PERI_PCLK_GRx_DIV_16_CTLy	Divider Control Register (for 16.0 divider)	This register controls the 16-bit divider.
PERI_PCLK_GRx_DIV_16_5_CTLy	Divider Control Register (for 16.5 divider)	This register controls the 16.5-bit divider.
PERI_PCLK_GRx_DIV_24_5_CTLy	Divider Control Register (for 24.5 divider)	This register controls the 24.5-bit divider.
PERI_PCLK_GRx_DIV_CMD	Divider Command Register	This register controls whether each divider is enabled or disabled.
CSV_REF_SEL	Select CSV Reference Clock Register for Active domain	Selects a source to be used as the reference clock for CSV in the Active domain.
CSV_HF_CSVx_REF_CTL	Clock Supervision Reference Control Register for CLK_HF	This register sets the control of CSV function.
CSV_REF_CSVx_REF_CTL	Clock Supervision Reference Control Register for CLK_REF_LF	
CSV_LF_CSVx_REF_CTL	Clock Supervision Reference Control Register for CLK_LF	
CSV_ILO_CSVx_REF_CTL	Clock Supervision Reference Control Register for clk_ilo0	
CSV_HF_CSVx_REF_LIMIT	Clock Supervision Reference Limits Register for CLK_HF	This register sets LOWER and UPEER to be used for the CSV function.
CSV_REF_CSVx_REF_LIMIT	Clock Supervision Reference Limits Register for CLK_REF_LF	
CSV_LF_CSVx_REF_LIMIT	Clock Supervision Reference Limits Register for CLK_LF	
CSV_ILO_CSVx_REF_LIMIT	Clock Supervision Reference Limits Register for clk_ilo0	
CSV_HF_CSVx_MON_CTL	Clock Supervision Monitor Control Register for CLK_HF	This register sets PERIOD to be used for the CSV function.
CSV_REF_CSVx_MON_CTL	Clock Supervision Monitor Control Register for CLK_REF_LF	
CSV_LF_CSVx_MON_CTL	Clock Supervision Monitor Control Register for CLK_LF	
CSV_ILO_CSVx_MON_CTL	Clock Supervision Monitor Control Register for clk_ilo0	

Table 18-11. Clock System Registers

Register	Name	Description
CLK_PLL400Mx_CONFIG	400-MHz PLL Configuration Register	This register contains PLL configuration for 400-MHz PLL. Each PLL 400 has a copy of this register. Configuration settings should not be changed while it is a selected clock (connected to logic). This prevents clock glitches that can crash the logic.
CLK_PLL400Mx_CONFIG2	400-MHz PLL Configuration Register2	This register configures the fractional divider.
CLK_PLL400Mx_CONFIG3	400-MHz PLL Configuration Register3	This register configures the SSCG.
CLK_PLL400Mx_STATUS	400-MHz PLL Status Register	This register indicates status for a PLL400. Each PLL400 has a copy of this register.
CPUSS_FAST_0_CLOCK_CTL	Fast 0 Clock Control Register	This register configures CLK_FAST_0.
CPUSS_FAST_1_CLOCK_CTL	Fast 1 Clock Control Register	This register configures CLK_FAST_0.
CPUSS_SLOW_CLOCK_CTL	Slow Clock Control Register	This register configures CLK_SLOW.
CPUSS_PERI_CLOCK_CTL	Peripheral Interconnect Clock Control Register	This register configures CLK_PERI.
CPUSS_MEM_CLOCK_CTL	Memory Clock Control Register	This register configures CLK_MEM.
CPUSS_TRC_DBG_CLOCK_CTL	Trace and Debug Clock Control Register	This register configures CLK_TRC_DBG.
BACKUP_LPECO_CTL	LPECO Control Register	This register configures LPECO.
BACKUP_LPECO_STATUS	LPECO Status Register	This register indicates status for LPECO.
BACKUP_LPECO_PRESCALE	LPECO Prescaler Register	This register configures LPECO prescaler.

19. Reset System



Traveo II supports several types of resets that guarantee error-free operation during power up and allow the device to reset based on user-supplied external hardware or internal software reset signals. Resets have a broad scope and are generally aligned with power domains and global power modes. The resets described in this chapter cause a reboot that ends in Active mode. Some blocks may have local resets that are described in their respective chapters. Reset assertion is asynchronously propagated and reset deassertion is synchronized to each clock domain where it is used. Traveo II also contains hardware to record which reset occurs.

The fault manager and processors can work together to reset parts of the device. The fault manager converts a fault into a high-priority interrupt (such as NMI) to give the processor an opportunity to return to a safe state, such as halting memory writes and releasing peripherals. The processor can then trigger its own local reset or a system reset. This allows recovery from faults generated by safety circuits (clock supervision, supply supervision, and multi-counter watchdog timer). These circuits can also generate their own direct reset for cases when the fault manager itself is unresponsive or possibly corrupted.

Traveo II has these reset sources:

- Power-on reset (POR) to hold the device in reset while the power supply is below the level required for initialization of startup circuits.
- Brownout detection reset (BOD) to reset the device if the power supply falls below the device specifications during normal operation.
- Over-voltage detection reset
- Over-current detection reset of the Active or DeepSleep regulator
- External reset (XRES_L) to reset the device using an external input
- Watchdog resets of the basic watchdog timer (WDT) and the multi-counter watchdog timers (MCWDT) to reset the device if the firmware execution fails to periodically service the watchdog timer
- Internal system reset to reset the device on demand using firmware
- Fault detection resets to reset the device if certain faults occur
- Clock-supervision resets to reset the device when clock-related errors occur
- Hibernate wakeup resets most logic to bring the device out of the Hibernate low-power mode

19.1 Reset Sources

The following sections provide a description of the reset sources available in Traveo II. [Table 19-1](#) shows the mapping of reset sources to the corresponding destinations that are affected by a reset event.

Table 19-1. Reset Cause Distribution

Reset Type/Source	High-voltage Reset Causes in RES_CAUSE	Low-voltage Reset Causes in RES_CAUSE	Data Registers in FAULT Structures	Debug	Hibernate Registers	RTC	CM0+	CM7_0/CM7_1	SRAM Retention
POR	x	x	x	x	x	x	x	x	No
XRES_L/PXRES		x	x	x	x		x	x	No
BOD		x	x	x	x		x	x	No
OVD		x	x	x	x		x	x	No
OCD		x	x	x	x		x	x	No
HIB WAKEUP		x	x	x			x	x	No
AIRCR.SYSRESETREQ ^a							x	x	b
CDBGIRSTREQ				x			x	x	b
FAULT				c			x	x	b
WDT		x	x	x	x		x	x	d
MCWDT				c			x	x	b
CSV HF							x	x	e
CSV REF							x	x	No

a. AIRCR.SYSRESETREQ is software reset.

b. Yes, if there is an orderly shutdown of the RAM.

c. Reset occurs if the source triggers during DeepSleep.

d. Yes, if there is an orderly shutdown of the RAM only during a warning interrupt.

e. Yes if there is an orderly shutdown of the RAM and the CSV reset is not from CSV_HF0.

Note: The SRAM region from (SRAM size minus 6 KB) to (SRAM size minus 2 KB) is used by the Cypress firmware during boot operation. Therefore, this region is available to the user; however, data retention across resets is not guaranteed in this area because it can be overwritten by the Cypress boot firmware. See [RAM Retention Configuration on page 131](#) for details.

19.1.1 Power-on Reset

Power-on reset keeps the system in a reset state during power-up. POR holds the device in reset until the supply voltage V_{DD} reaches a sufficient level to initialize the startup circuits. The POR activates automatically at power-up. All other circuits are disabled until POR releases. See the [Power Supply and Monitoring chapter on page 163](#) for more details.

19.1.2 Brownout Detection Reset

Brownout detection circuits monitor the device digital voltage supply V_{DD} , device analog voltage supply V_{DDA} , and internally-generated supply voltage V_{CCD} and generate a reset if they fall below their voltage threshold. The device stays in reset until all brownout detectors release. This also occurs during an initial power ramp, but is not recorded as brownout reset. See the [Power Supply and Monitoring chapter on page 163](#) for more details.

19.1.3 Over-Voltage Detection Reset

Over-voltage circuits monitor the device digital voltage supply V_{DD} , device analog voltage supply V_{DDA} , and internally-generated supply voltage V_{CCD} and generate a reset if they rise above their voltage threshold. The device stays in reset until all over-voltage detectors release. See the [Power Supply and Monitoring chapter on page 163](#) for more details.

19.1.4 Over-Current Reset

Over-currents of the internally-generated supply voltage V_{CCD} are detected and cause a reset. The observation is done in Active and DeepSleep power modes. See the [Power Supply and Monitoring chapter on page 163](#) for more details.

19.1.5 External Reset

External reset (XRES_L) is a reset triggered by an external signal that causes immediate system reset when asserted. The XRES_L pin is active low – a logic ‘1’ on the pin has no effect and a logic ‘0’ causes reset. The pin is pulled to logic ‘1’ inside the device. XRES_L is available as a dedicated pin. For the detailed pinout, refer to the pinout section of the device datasheet.

The XRES_L pin holds the device in reset as long as the pin input is ‘0’. When the pin is released (changed to logic ‘1’), the device goes through a normal boot sequence. The logical thresholds for XRES_L and other electrical characteristics are listed in the Electrical Specifications section of the device datasheet. XRES_L is available in all power modes.

19.1.6 Programmable Reset

PXRES is a programmable XRES_L trigger. This is a software-accessible register that triggers a full-scope reset equivalent to XRES_L, except it records in a different RES_CAUSE flag, RESET_PXRES. Hardware clears this bit during POR.

19.1.7 Watchdog Timer Reset

Watchdog timer reset causes a reset if the WDT or MCWDTs are not serviced by the firmware within a specified time limit or it is serviced too early in case of window mode.

For details, see the [Watchdog Timer chapter on page 222](#).

19.1.8 Internal System Reset

The internal system reset is a mechanism that allows software running on any of the CPUs or a connected debugger to request a system reset. The Cortex-M0+ and Cortex-M7 Application Interrupt and Reset Control registers (CM0P_SCS_AIRCR and CM7_0/CM7_1_SCS_AIRCR, respectively) can request a reset by writing a ‘1’ to the SYSRESETREQ bit of the respective registers.

Note that a value of 0x5FA should be written to the VECTKEY field of the AIRCR register before setting the SYSRESETREQ bit; otherwise, the processor ignores the write. See the [CPU Subsystem \(CPUSS\) chapter on page 38](#) for details.

19.1.9 Fault Detection Reset

The fault reporting structures in Traveo II can be configured to request a reset for user-configurable faults, such as uncorrectable ECC errors or protection violations.

Traveo II does not support direct handling of faults during DeepSleep mode. If a fault occurs during DeepSleep, it wakes the device and then triggers the Active mode fault detection reset. The DeepSleep mode fault detection reset is not used in Traveo II, so it cannot be set by a fault. Both bits remain set until cleared by firmware or until a POR reset.

Faults generated by clock supervision and MCWDTs can indicate that the fault system may also have failed. These circuits can be configured to directly cause a reset.

19.1.10 Clock-Supervision Reset

Clock-supervision logic initiates a reset when a monitored clock stops or is outside the configured relationship to a reference clock. Clock supervisors on the high-frequency clocks (HFCLKn) and the CSV reference clock supervisor can generate resets. Clock supervisors for the low-frequency clocks cannot trigger a direct reset, because the fault system and processor can safely convert these faults into resets.

The fault manager and processor clocks are derived from HFCLK0. It is recommended to configure the HFCLK0 CSV to generate a reset or fault-then-reset.

For more information on clocks, see the [Clocking System chapter on page 182](#).

19.1.11 Hibernate Wakeup Reset

Hibernate wakeup reset occurs when a wakeup source triggers an exit from Hibernate mode. The device returns to the Active power mode and the processors reboot. See the [Device Power Modes chapter on page 186](#) for details on Hibernate mode and the available wakeup sources.

TOKEN is an 8-bit field in the PWR_HIBERNATE register that is retained through a hibernate-wakeup sequence. The firmware can use this bit field to differentiate hibernate wakeup from a general reset event, such as XRES_L or POR. Similarly, the PWR_HIB_DATAx register retains its contents through a hibernate wakeup sequence.

19.2 Identifying Reset Sources

When the device comes out of reset, it is often useful to know the cause of the reset. Reset causes are recorded in the RES_CAUSE and RES_CAUSE2 registers. The bits in these registers are set on the occurrence of the corresponding reset and remain set until cleared by the firmware or a POR reset.

An internal reset that occurred due to hibernate wakeup can be detected by examining the TOKEN field in the PWR_HIBERNATE register as described previously. Hibernate exit caused by an XRES_L is recorded as an external reset. The reset causes in the RES_CAUSE and RES_CAUSE2 registers are shown in [Table 19-2](#).

After identifying and evaluating the reset cause, clear the RESET_CAUSE and RESET_CAUSE2 register. This procedure is required to capture the next reset cause.

Table 19-2. Reset Cause Bits to Detect Reset Source

Register [Bit_Pos]	Bit Field	Description
RES_CAUSE [0]	RESET_WDT	A basic WDT reset.
RES_CAUSE [1]	RESET_ACT_FAULT	Fault logging system requested a reset from its Active logic.
RES_CAUSE [2]	RESET_DPSLP_FAULT	Fault logging system requested a reset from its DeepSleep logic.
RES_CAUSE [3]	RESET_TC_DBGRESET	Test controller or debugger asserted reset. Only resets debug domain.
RES_CAUSE [4]	RESET_SOFT	A CPU requested a system reset through its SYSRESETREQ.
RES_CAUSE [5]	RESET_MCWDT0	MCWDT reset #0.
RES_CAUSE [6]	RESET_MCWDT1	MCWDT reset #1.
RES_CAUSE [7]	RESET_MCWDT2	MCWDT reset #2.
RES_CAUSE [8]	RESET_MCWDT3	MCWDT reset #3.
RES_CAUSE [16]	RESET_XRES	External XRES_L pin is asserted.
RES_CAUSE [17]	RESET_BODVDDD	External V _{DDD} supply crossed the brownout limit.
RES_CAUSE [18]	RESET_BODVDDA	External V _{DDA} supply crossed the brownout limit.
RES_CAUSE [19]	RESET_BODVCCD	External V _{CCD} supply crossed the brownout limit.
RES_CAUSE [20]	RESET_OVDVDDD	Overvoltage detection on the external V _{DDD} supply.
RES_CAUSE [21]	RESET_OVDVDDA	Overvoltage detection on the external V _{DDA} supply.
RES_CAUSE [22]	RESET_OVDVCCD	Overvoltage detection on the internal core V _{CCD} supply.
RES_CAUSE [23]	RESET_OCD_ACT_LINREG	Overcurrent detection on the internal V _{CCD} supply when supplied by the Active power mode linear regulator.
RES_CAUSE [24]	RESET_OCD_DPSLP_LINREG	Overcurrent detection on the internal V _{CCD} supply when supplied by the Deep-Sleep power mode linear regulator.
RES_CAUSE [28]	RESET_PXRES	CPU requested a programmable reset through PXRES.
RES_CAUSE[30]	RESET_PORVDDD	Indicator that a POR occurred. This is a high-voltage cause bit, and hardware clears the other bits when this bit is set. It does not block further recording of other high-voltage causes.
RES_CAUSE2 [15:0]	RESET_CSV_HF	Clock supervision logic requested a reset due to loss or frequency violation of a high-frequency clock. Each bit index K corresponds to a HFCLK<K>.
RES_CAUSE2 [16]	RESET_CSV_REF	Clock supervision logic requested a reset due to loss or frequency violation of the reference clock source that is used to monitor other high-frequency clock sources.

For more information, see the RES_CAUSE and RES_CAUSE2 registers in the *Traveo II Body Controller High Registers TRM*.

19.3 Register List

Table 19-3. Reset System Register List

Register	Name	Description
RES_CAUSE	Reset Cause Observation Register	Indicates the cause for the latest reset(s) that occurred in the system.
RES_CAUSE2	Reset Cause Observation Register 2	Indicates the cause for the latest reset(s) that occurred in the system.
PWR_HIBERNATE	Hibernate mode register	This register controls entry/exit from Hibernate power mode.
PWR_HIB_DATAx	Hibernate data register	This register retains its contents through Hibernate wakeup reset. 'x' signifies the number of such DATA registers. Refer to the Traveo II Registers TRM for more information.
CM7_0/ CM7_1_SCS_AIRCR	Application Interrupt and Reset Control Register of Cortex-M7	Application interrupt and reset control register specific to Cortex-M7
CM0P_SCS_AIRCR	Application Interrupt and Reset Control Register of Cortex-M0+	Application interrupt and reset control register specific to Cortex-M0+

20. Watchdog Timer



The watchdog timer (WDT) in Traveo II includes a hardware timer that automatically resets the device in the event of an unexpected firmware execution path. It uses LFCLK (ILO0, ILO1, WCO, or ECO) as the input clock. The WDT, if enabled, must be serviced periodically in firmware to avoid a reset. Otherwise, the timer will elapse and generate a device reset. In the window function mode, the WDT can generate a reset if it is serviced too early or not serviced at all before a timeout is reached. In addition, the WDT can be used as an interrupt source or a wakeup source in low-power modes.

Traveo II includes one 32-bit free-running basic WDT supporting window mode and up to four multi-counter watchdog timers (MCWDT). Each MCWDT includes three subcounters – two 16-bit timers supporting window mode and one 32-bit free-running timer. All counters with window mode functionality can generate a reset and a WARN interrupt; the MCWDT counters can also generate a FAULT condition. Each MCWDT is independent; it is recommended to assign one MCWDT per processor.


20.1 Features

The Traveo II watchdog timer supports the following features:

- One 32-bit free-running basic WDT with:
 - ILO0 as the input clock source
 - Programmable early threshold, warning threshold, and timeout threshold
 - Device reset generation if not serviced within a configurable interval
 - Warning threshold generates an interrupt to request servicing
 - Interrupt/wakeup generation in Active, Sleep, DeepSleep, and Hibernate power modes
- Up to four MCWDTs, each supporting:
 - LFCLK (ILO0, ILO1, WCO, or ECO) as the input clock source
 - Fault and device reset generation if not serviced within a configurable interval
 - Periodic interrupt/wakeup generation in Active, Sleep, and DeepSleep power modes
 - Three independent counters: two 16-bit counters and one 32-bit counter
 - Warning threshold generates an interrupt to request servicing
- Both watchdog timer types support:
 - Window mode
 - Running and freezing timers during DeepSleep mode
 - Debug

20.2 Block Diagram

Figure 20-1. Watchdog Timer Block Diagram



20.3 Basic Watchdog Timer


20.3.1 Overview

The WDT is a free-running up-counter with programmable limit values, a maximum of 32-bit resolution, and a clock from the ILO0. Servicing the watchdog clears and restarts the counter at zero.

The WDT can be configured to act on different counter limits where a reset is triggered if the watchdog is not serviced before the upper limit. In the window mode, a reset is triggered if the servicing occurs before the lower limit is reached. The warning limit triggers an interrupt to request servicing. Each of these actions can be activated independently. The WDT is enabled and specific registers are locked by default. An unlocking sequence is required to prevent accidental accesses. The WDT operates in Active, Sleep, DeepSleep, and Hibernate modes. Hibernate mode operation is possible because the logic and clock source are powered by the external high-voltage supply (V_{DD}). After a WDT reset the device returns to Active mode.

Figure 20-2 shows the functional overview of the WDT.

Figure 20-2. Basic WDT Functional Diagram



When enabled, the WDT counts up on each rising edge of the ILO0 clock. When the counter value (WDT_CNT register) equals the warning threshold value stored in WDT_WARN_LIMIT [31:0], an interrupt is generated if the WARN_ACTION [8] bit is set to '1' in the WDT_CONFIG register. The warn event will not reset the WDT counter and the WDT continues counting until it reaches the timeout threshold value stored in UPPER_LIMIT [31:0]; it generates a reset if the UPPER_ACTION [4] bit is set to '1' in the WDT_CONFIG register. If no action is taken on the upper threshold, the counter increments to the 32-bit boundary and then wraps around to '0' and counts up. In the window mode, an early threshold stored in LOWER_LIMIT [31:0] can be used if the LOWER_ACTION [0] bit is set to '1' in the WDT_CONFIG register for generating a reset if the counter is serviced too early. The watchdog counter is serviced by the SERVICE [0] bit in the WDT_SERVICE register. If this bit is set to '1' the watchdog counter is set to zero.

The WDT [0] bit in the WDT_INTR register is set whenever the WDT counter matches with the WARN_LIMIT and an interrupt is requested by the CPU. This interrupt must be cleared by writing a '1' to the same bit (WDT bit of WDT_INTR). Clearing the interrupt does not reset the watchdog counter.

The WDT can be enabled or disabled using the ENABLE [31] bit of the WDT_CTL register. The actual status of the

counter is indicated by the ENABLED [0] bit of the WDT_CTL register.

The WDT provides a mechanism to lock WDT configuration registers. The WDT_LOCK bits [1:0] control the lock status of WDT-related registers. These are special bits, which can enable the lock in a single write (WDT_LOCK = 3); to release the lock, two different write accesses are required (WDT_LOCK = 1 to clear WDT_LOCK [0] and WDT_LOCK = 2 to clear WDT_LOCK [1]). When the WDT_LOCK bits are not equal to '0' the write accesses to the CTL, LOWER_LIMIT, WARN_LIMIT, UPPER_LIMIT, CNT, and SERVICE registers are prohibited. Note that this field is two bits to force multiple writes only. It represents only a single write protect signal protecting all those registers at the same time. WDT will lock and enable on any reset. This field is not retained during DeepSleep or Hibernate mode, so the WDT will be locked after wakeup from these modes.

Note: The lock mechanism is an additional safety opportunity, which requires to unlock/lock the SERVICE register when servicing each watchdog counter. Alternatively, the WDT registers can also be protected by the PPU, which allows to keep the WDT registers unlocked.

When the watchdog counter is disabled and unlocked, the count value can be written for verification and debugging purposes. Software writes are always ignored when the counter is enabled.

Table 20-1 explains various registers and bit fields used to configure and use the WDT.

Table 20-1. Basic Watchdog Timer Configuration Options

Register [Bit_Pos]	Bit Name	Description
WDT_CTL[31]	ENABLE	Enable or disable the watchdog counter ■ 0: Counter is disabled (not clocked) ■ 1: Counter is enabled (counting up)
WDT_CTL[0]	ENABLED	Indicates actual state of watchdog
WDT_LOCK[1:0]	WDT_LOCK	Prohibits writing control and configuration registers related to this MCWDT when not equal to 0 ■ 0: No effect ■ 1: Clear bit 0 ■ 2: Clear bit 1 ■ 3: Set both bit 0 and 1 (lock enabled)
WDT_CNT[31:0]	CNT	Current value of WDT counter
WDT_LOWER_LIMIT[31:0]	LOWER_LIMIT	Lower limit for watchdog
WDT_UPPER_LIMIT[31:0]	UPPER_LIMIT	Upper limit for watchdog
WDT_WARN_LIMIT[31:0]	WARN_LIMIT	Warn limit for watchdog
WDT_CONFIG[0]	LOWER_ACTION	Action taken if this watchdog is serviced before LOWER_LIMIT is reached ■ 0: Do nothing ■ 1: Trigger a reset
WDT_CONFIG[4]	UPPER_ACTION	Action taken if this watchdog is not serviced before UPPER_LIMIT is reached ■ 0: Do nothing ■ 1: Trigger a reset
WDT_CONFIG[8]	WARN_ACTION	Action taken when the count value reaches WARN_LIMIT ■ 0: Do nothing ■ 1: Trigger an interrupt

Table 20-1. Basic Watchdog Timer Configuration Options

Register [Bit_Pos]	Bit Name	Description
WDT_CONFIG[12]	AUTO_SERVICE	Automatically service when the count value reaches WARN_LIMIT. This allows creation of a periodic interrupt if this counter is not needed as a watchdog.
WDT_CONFIG[28]	DEBUG_TRIGGER_EN	Enables the trigger input for the WDT to pause the counter in debug mode. 0: Pauses the counter when a debug probe is connected. 1: Pauses the counter when a debug probe is connected and the trigger input is high.
WDT_CONFIG[29]	DPSLP_PAUSE	Pauses/runs this counter when the system is in DeepSleep 0: Counter behaves normally during DeepSleep 1: Counter pauses during DeepSleep
WDT_CONFIG[30]	HIB_PAUSE	Pauses/runs this counter when the system is in Hibernate 0: Counter behaves normally during Hibernate 1: Counter pauses during Hibernate
WDT_CONFIG[31]	DEBUG_RUN	Pauses/runs this counter while a debugger is connected 0: Counter pauses according to DEBUG_TRIGGER_EN configuration 1: Counter runs normally when debugger connected
WDT_INTR[0]	WDT	WDT Interrupt Request. This bit is set as configured by WDT action and limits. The WDT interrupt is cleared by writing a '1' to this bit.
WDT_INTR_SET[0]	WDT	WDT Interrupt set register. Can be used to set interrupts for firmware testing.
WDT_INTR_MASK[0]	WDT	Mask for the WDT interrupt 0: WDT interrupt is masked to CPU 1: WDT interrupt is not masked to CPU
WDT_INTR_MASKED[0]	WDT	Logical AND of corresponding request and mask bits

Note: WDT configuration registers are in a separate protection region from the register used to service it. The protection regions are handled by the peripheral protection unit (PPU). Refer to the [CPU Subsystem \(CPUSS\) chapter on page 37](#) for more information.

20.3.2 Watchdog Reset

A watchdog is typically used to protect the device against firmware/system crashes or faults. When the WDT is used to protect against system crashes, the WDT counter should be cleared by writing a '1' to the SERVICE [0] bit in the SERVICE register from a portion of the code that is not directly associated with the WDT interrupt. Otherwise, even if the main function of the firmware crashes or is in an endless loop, the WDT interrupt vector can still be intact and feed the WDT periodically.

The safest way to use the WDT against system crashes is to:

- Configure the UPPER_LIMIT such that firmware is able to reset the watchdog at least once during the period, even along the longest firmware delay path
- In window mode, configure the LOWER_LIMIT to serve the watchdog counter not too early, even along the shortest firmware delay path.
- Reset (feed) the watchdog for clearing the counter regularly in the main body of the firmware code by setting the SERVICE [0] bit to '1' in SERVICE register.

It is not recommended to reset the watchdog counter in the WDT interrupt service routine (ISR), if WDT is being used as a reset source to protect the system against crashes. If necessary, use the warning interrupt to set a flag in the ISR. Local processing loops can observe that flag and break out of their loop. This allows the main loop to reach the servicing code (and clear the flag for the next pass through the main loop).

Recommended steps to use WDT as a reset source are as follows:

1. Write UPPER_LIMIT value to define the timeout period for reset generation. Set UPPER_ACTION [4] bit to '1' in the WDT_CONFIG register to enable a reset trigger when the watchdog counter reaches the UPPER_LIMIT.
2. If required, write the WARN_LIMIT to generate an interrupt before reaching the UPPER_LIMIT threshold. Do not use the ISR to feed the WDT; instead, use this interrupt to indicate that there is a firmware delay path,


which is already critical. Use a warn level that is close enough to the UPPER_LIMIT but consider also the delay to handle the ISR and return to your main body functions for serving the watchdog counter. Set WARN_ACTION [8] bit to '1' in the WDT_CONFIG register to enable a watchdog warn interrupt when the watchdog counter matches with the WARN_LIMIT.

3. In window mode, define an adequate LOWER_LIMIT, which cannot be violated by the shortest firmware delay path. Set the LOWER_ACTION [0] bit to '1' in the WDT_CONFIG register to enable a reset trigger when the watchdog counter is serviced before the counter reaches the LOWER_LIMIT.
4. Set the WDT [0] bit in the WDT_INTR register to clear any pending WDT interrupt.
5. Set the ENABLE [31] bit in the CLK_ILO0_CONFIG register to enable the ILO0 clock.
6. Enable the WDT by setting the ENABLE [31] bit in the WDT_CTL register.
7. In the firmware, write '1' to the SERVICE [0] bit in the SERVICE register to feed (reset) the watchdog.
8. Lock the WDT configuration by writing '3' to the WDT_LOCK bits.

Figure 20-3 shows all scenarios of the WDT operation while LOWER_ACTION, WARN_ACTION, and UPPER_ACTION are enabled.

- Counter is serviced between LOWER_LIMIT and WARN_LIMIT: This is the regular behavior of the WDT. No WARN interrupt is issued and no RESET is done
- Counter is serviced between WARN_LIMIT and UPPER_LIMIT: The service is done late, a WARN interrupt is issued but no RESET is done.
- Counter is not serviced at all: WARN interrupt is issued but the SERVICE bit is not set. When the counter reaches the UPPER_LIMIT a reset is executed.
- Counter is serviced before the LOWER_LIMIT is reached: The counter is serviced too early; a reset is executed because the counter is cleared outside of the window.

Figure 20-3. WDT Counter Operation in Window Mode



Note: This figure illustrates the different scenarios with or without servicing the watchdog counter. It does not consider the WDT configuration, especially after a reset.

20.3.3 Watchdog Interrupt

In addition to generating a device reset, the WDT can be used to generate interrupts. The watchdog counter can send interrupt requests to the CPU in Active power modes and to the wakeup interrupt controller (WIC) in Sleep and DeepSleep power modes. In addition, the watchdog is capable of waking up the device from Hibernate power mode. It works as follows:

- **Active Mode:** In this mode, the WDT can send the interrupt to the CPU. The CPU acknowledges the interrupt request and executes the ISR. The interrupt must be cleared after entering the ISR in firmware.
- **Sleep or DeepSleep Mode:** In these modes, the CPU subsystem is powered down. Therefore, the interrupt request from the WDT is directly sent to the WIC, which will then wake up the CPU. The CPU acknowledges the interrupt request and executes the ISR. The interrupt must be cleared after entering the ISR in firmware.
- **Hibernate Mode:** In this mode, the entire device except a few peripherals (such as WDT) are powered down. Any interrupt to wake up the device in this mode results in a device reset. Hence, there is no interrupt service routine or mechanism associated with this mode.

For more details on device power modes, see the [Device Power Modes chapter on page 186](#). Because of its free-running nature, it is not recommended to use the WDT for periodic interrupt generation. The MCWDT counters can be used to generate periodic interrupts. If absolutely required, follow these steps to use the WDT as a periodic interrupt generator:

1. Write the `WARN_LIMIT` to set the interrupt period. If the WDT is not serviced, the counter will continue to count


up until the maximum counter level of `0xFFFFFFFF` is reached and then the counter starts from zero.

2. Set the `WARN_ACTION` [8] bit to '1' in the `WDT_CONFIG` register to enable a watchdog warn interrupt when the watchdog counter matches with the `WARN_LIMIT`.
3. Set the `WDT` [0] bit in the `WDT_INTR` register to clear any pending WDT interrupt.
4. Enable the WDT interrupt to CPU by setting the `WDT` [0] bit in the `WDT_INTR_MASK` register.
5. Enable SRSS interrupt to the CPU by configuring the appropriate `ISER` register, see the [Interrupts chapter on page 144](#) for details.
6. In the ISR, clear the WDT interrupt; if required, clear the watchdog timer by writing '1' to the `SERVICE` [0] bit in the `SERVICE` register. Servicing the WDT allows to generate various interrupt periods, which can be defined by the `WARN_LIMIT`. Alternatively, set the `AUTO_SERVICE`[12] bit to '1' in the `CONFIG` register to automatically service the WDT when the count value reaches `WARN_LIMIT`.

Waking up from DeepSleep mode requires to execute an unlock sequence by writing the value '1' to the `WDT_LOCK` [1:0] bits in the `WDT_LOCK` register followed by writing '2' to the same bit field.

[Figure 20-4](#) shows the behavior of the WDT counter in interrupt mode. LOWER and UPPER actions are disabled. An interrupt is issued each time the counter matches the `WARN_LIMIT` and continuous to count up to the 32-bit maximum value. The interrupt period is calculated by $2^{32} \times \text{ILO0 clock cycles}$. The WDT does not provide an automatic counter clear function; therefore, the counter must be cleared manually by writing '1' to the `SERVICE`[0] bit in the `WDT_SERVICE` register.

Figure 20-4. WDT Counter Operation with WARN Interrupt Only



20.4 Multi-Counter Watchdog Timer

20.4.1 Overview

Figure 20-5 shows the functional overview of a single MCWDT block. Depending on the device, Traveo II includes up to four MCWDT blocks. Each MCWDT block includes two 16-bit counters, Subcounter 0 (MCWDTx_CNT0) and Subcounter 1 (MCWDTx_CNT1), which include the same window and threshold concept described for the WDT, and one 32-bit counter, Subcounter 2 (MCWDTx_CTR2_CNT). Cascading of these counters is not supported. These counters work independently.

The Subcounters 0 and 1 have the ability to generate a FAULT when the MCWDTx_CTRy_LOWER_LIMIT or MCWDTx_CTRy_UPPER_LIMIT is violated. The fault structure can convert this to an interrupt (such as a high-priority NMI) that gives the processor an opportunity to return to a safe state, such as halting memory writes and releasing peripherals. It can then clear the fault and trigger its own local reset. If the fault is not cleared within a fixed number of LFCLK cycles, MCWDT will trigger a system-wide reset as a failsafe.

Note: If a single MCWDT triggers additional fault actions while transferring fault data to the fault manager, then the fault manager receives only the first action in the fault data. The additional overlapping fault actions do not cause another fault report and are lost. Faults are transferred

correctly, if they occur when the MCWDT is not in the middle of transferring another fault. Faults generated by a different MCWDT are not affected.

A missed MCWDTx_CTRy_CONFIG.UPPER_ACTION fault can be detected during MCWDT fault processing. For counter values of both subcounter 0 and subcounter 1, check whether the condition $CNT \geq MCWDTx\_CTRy\_UPPER\_LIMIT$ is valid. There is no known method to detect a missed MCWDTx_CTRy_CONFIG.LOWER_ACTION fault.


The 32-bit counter can only generate interrupt after a programmed bit position toggles.

All the counters are clocked by LFCLK and operate in Active, Sleep, and DeepSleep modes. Hibernate mode is not supported. After a MCWDT reset, the chip is recovered to Active mode. Servicing a counter clears and restarts the related counter at zero. The MCWDT is disabled and unlocked by default.

Note: Because Traveo II includes two CPUs (Cortex-M0+ and Cortex-M7), it is recommended to associate one MCWDT block to only one CPU during runtime. Although both the MCWDT blocks are available to both CPUs, a single MCWDT is not intended to be used by multiple CPUs simultaneously.

Register protection is handled by putting the Subcounter 0 and Subcounter 1 configuration registers in a protection region. A separate protection region is used for the registers related to servicing and Subcounter 2.

Figure 20-5. Multi Counter Watchdog Timer Functional Diagram



20.4.2 How It Works

20.4.2.1 Subcounter 0/1 Operation

Subcounter 0 (MCWDTx_CNT0) and Subcounter 1 (MCWDTx_CNT1) are independent 16-bit up-counters. If enabled, they can count up on each rising edge of the LFCLK clock. ILO0, ILO1, WCO, or ECO can be configured as a clock source. See the [Clocking System chapter on page 197](#).

When a counter value (MCWDTx_CTRy_CNT¹ register) equals the warning threshold value stored in MCWDTx_CTRy_WARN_LIMIT [15:0], an interrupt is generated if the MCWDTx_CTRy_CONFIG.WARN_ACTION [8] bit is set to '1' in the MCWDTx_CTRy_CONFIG register. Both counters can be cleared automatically by each warn event when the MCWDTx_CTRy_CONFIG.AUTO_SERVICE [12] bit in the MCWDTx_CTRy_CONFIG register is set to '1' when both MCWDTx_CTRy_CONFIG.UPPER_ACTION==NOTHING && MCWDTx_CTRy_CONFIG.LOWER_ACTION==NOTHING. This allows creation of a periodic interrupt if this counter is not needed as a watchdog. The CTR0_INT [0] or CTR1_INT [1] bits in the MCWDTx_INTR register are set whenever the corresponding MCWDT counter matches with the related WARN_LIMIT [15:0] and an interrupt occurs. This interrupt must be cleared by writing a '1' to the same bit.

If no automatic service is enabled the match event will keep the MCWDT counting until it reaches the timeout threshold value stored in MCWDTx_CTRy_UPPER_LIMIT [15:0]; this generates a FAULT if the MCWDTx_CTRy_CONFIG.UPPER_ACTION [5:4] bits are set to '1' in the MCWDTx_CTRy_CONFIG register. In window mode, an early threshold stored in MCWDTx_CTRy_LOWER_LIMIT [15:0] can be used if the MCWDTx_CTRy_CONFIG.LOWER_ACTION [1:0] bit is set to '1' in the MCWDTx_CTRy_CONFIG register to generate a FAULT if the counter is serviced too early.

All four faults for each counter (early and timeout for each 16-bit subcounter) are combined into a single fault triggered, so the fault structure can record the correct fault cause. The fault structure can convert this to an interrupt (such as a high-priority NMI) that gives the processor an opportunity to return to a safe state, such as halting memory writes and releasing peripherals. It can then clear the FAULT and trigger its own local reset. If the FAULT is not cleared within a fixed number of LFCLK cycles, MCWDT will trigger a system-wide reset as a failsafe if the value '2' is written to MCWDTx_CTRy_CONFIG.LOWER_ACTION [1:0] or MCWDTx_CTRy_CONFIG.UPPER_ACTION [5:4] bits.

1. Subcounter 0 and Subcounter 1 have its own register sets. For simplification MCWDTx prefix is used for both register sets, MCWDT0 and MCWDT1. In all cases when Subcounter 2 is used, MCWDT2 register set is mentioned.

Note: MCWDT does not report overlapping faults generated by other subcounter actions. If a single MCWDT triggers additional fault action(s) while it is transferring fault data to the fault manager, then the fault manager receives only the first action in the fault data. The additional overlapping fault action(s) do not cause another fault report and are lost. Faults are transferred properly if they occur when the MCWDT is not transferring another fault. When processing an MCWDT fault, a missed MCWDTx_CTRy_CONFIG.UPPER_ACTION fault can be detected. For each subcounter 0 and 1, if MCWDTx_CTRy_CONFIG.UPPER_ACTION is configured for FAULT or FAULT_THEN_RESET, and CNT ≥ MCWDTx_CTRy_UPPER_LIMIT, then process it as a fault even if it is not present in the fault data. There is no known method to detect a missed MCWDTx_CTRy_CONFIG.LOWER_ACTION fault.

If no action is taken on the upper threshold the counter increments up to the 16-bit boundary at which point, it wraps around to 0 and counts up.

The watchdog counters are serviced by dedicated service bits. CTR0_SERVICE [0] bit is related to Subcounter 0 and CTR1_SERVICE [1] is related to Subcounter 1. Both bits are located in the MCWDTx_SERVICE register. If this bit is set to '1' the watchdog counter is set to zero.

Note: When the software writes the MCWDT SERVICE bit in the MCWDTx_SERVICE register just before updating a limit register in an enabled MCWDT counter, the limit update may take effect before the service clears the counter. The new limit may trigger actions when they are compared to the uncleared counter value. For example, this can happen if the value in the MCWDTx_CTRy_LOWER_LIMIT register is changed to a value smaller than the existing CNT value. An unexpected fault or reset can occur during update, depending on the MCWDT configuration. To avoid this issue, make sure that a pending service is completed by waiting until the SERVICE bit value is read '0' before writing the limit registers. It can take up to three LFCLK cycles for the service to complete.

Subcounter 0 and Subcounter 1 can be enabled or disabled using the ENABLE [31] bit of the MCWDTx_CTRy_CTL register. The actual status of the counter is indicated by the ENABLED [0] bit of the MCWDTx_CTRy_CTL register.

Both subcounters have the same mechanism to lock the MCWDT configuration registers as provided by the basic WDT. When the MCWDT_LOCK[1:0] bits in the MCWDTx_LOCK register are not equal to '0' the write access to the MCWDTx_CTRy_CTL, MCWDTx_CTRy_LOWER_LIMIT, MCWDTx_CTRy_UPPER_LIMIT, MCWDTx_CTRy_WARN_LIMIT, MCWDTx_CTRy_CONFIG, MCWDTx_CTRy_CNT, and MCWDTx_SERVICE registers is prohibited. MCWDT will be unlocked and disabled on any reset.

Note: The lock mechanism is an additional safety opportunity, which requires to unlock/lock the SERVICE register when servicing each watchdog counter. Alternatively, the MCWDT registers can also be protected by the PPU, which allows to keep these registers unlocked.

When the watchdog counter is disabled and unlocked, the count value can be written for verification and debugging purposes. Software writes are always ignored when the counter is enabled.

Figure 20-6 shows the operation of the 16-bit subcounters. If the MCWDTx_CTRy_CONFIG.WARN_ACTION is activated, the counter can be used for interrupt generation

exclusively. The counter continues to increment after the counter value matches the WARN_LIMIT until the 16-bit maximum value is reached. Then the counter restarts at zero. Note that the interrupt period is fixed in this use case. For various interrupt timing the counter must be serviced regularly.

To clear the counter manually within an ISR, the CTR0_SERVICE[0] or CTR1_SERVICE[1] bit in the MCWDTx_SERVICE register should be set to '1'. Alternatively, the counter can be serviced automatically by setting MCWDTx_CTRy_CONFIG.AUTO_SERVICE[12] bit in the CONFIG register to '1'.

Figure 20-6. Subcounter 0/1 Operation with WARN Interrupt Only (MCWDTx_CTRy_CONFIG.AUTO_SERVICE = 0)





Figure 20-7 illustrates the interrupt mode with enabled automatic service (MCWDTx_CTRy_CONFIG.AUTO_SERVICE = 1). Whenever the counter matches the WARN_LIMIT value, an interrupt is issued and the counter is restarted with zero.

Note: The MCWDTx_CTRy_CONFIG.AUTO_SERVICE bit is ignored when either MCWDTx_CTRy_CONFIG.LOWER_ACTION or MCWDTx_CTRy_CONFIG.UPPER_ACTION is enabled.

Figure 20-7. Subcounter 0/1 Operation with WARN Interrupt Only (MCWDTx_CTRy_CONFIG.AUTO_SERVICE = 1)




In Figure 20-8 the window mode is shown when FAULT_THEN_RESET function is enabled. Four scenarios can happen while MCWDTx_CTRy_CONFIG.LOWER_ACTION, MCWDTx_CTRy_CONFIG.WARN_ACTION, and MCWDTx_CTRy_CONFIG.UPPER_ACTION are activated:

- Counter is serviced between MCWDTx_CTRy_LOWER_LIMIT and WARN_LIMIT: This is the regular behavior of the MCWDT. No WARN interrupt is issued and no RESET is done.

- Counter is serviced between WARN_LIMIT and MCWDTx_CTRy_UPPER_LIMIT: The service is done late; a WARN interrupt is issued but no RESET is done.
- Counter is not serviced at all: WARN interrupt is issued but the CTR0_SERVICE or CTR1_SERVICE bit is not set. When the counter reaches the MCWDTx_CTRy_UPPER_LIMIT, a FAULT is issued. If the firmware does not handle this FAULT to bring the system back into a safe state, a RESET is issued after a fixed number of LFCLK cycles.
- Counter is serviced before the MCWDTx_CTRy_LOWER_LIMIT is reached: The counter is serviced too early; a FAULT is issued followed by a RESET in case the FAULT is not handled in time by the firmware.

Figure 20-8. Subcounter 0/1 Operation in Window Mode with FAULT and RESET Action



Note: This figure illustrates the different scenarios with or without servicing the watchdog counter. It does not consider the WDT configuration, especially after a reset.

Table 20-2. MCWDT Subcounter 0 and Subcounter 1 Configuration Options

Register [Bit_Pos]	Bit Name	Description
MCWDTx_CTRy_CONFIG[1:0]	LOWER_ACTION	Action taken if this watchdog is serviced before MCWDTx_CTRy_LOWER_LIMIT is reached ■ 0: Do nothing ■ 1: FAULT ■ 2: FAULT_THEN_RESET
MCWDTx_CTRy_CONFIG[5:4]	UPPER_ACTION	Action taken if this watchdog is not serviced before MCWDTx_CTRy_UPPER_LIMIT is reached ■ 0: Do nothing ■ 1: FAULT ■ 2: FAULT_THEN_RESET
MCWDTx_CTRy_CONFIG[8]	WARN_ACTION	Action taken when the count value reaches WARN_LIMIT ■ 0: Do nothing ■ 1: Interrupt
MCWDTx_CTRy_CONFIG[12]	AUTO_SERVICE	Automatically service when the count value reaches WARN_LIMIT
MCWDTx_CTRy_CONFIG[28]	DEBUG_TRIGGER_EN	Enables the trigger input for the MCWDT to pause the counter in debug mode. ■ 0: Pauses the counter when a debug probe is connected. ■ 1: Pauses the counter when a debug probe is connected and the trigger input is high.

Table 20-2. MCWDT Subcounter 0 and Subcounter 1 Configuration Options

Register [Bit_Pos]	Bit Name	Description
MCWDTx_CTRy_CONFIG[30]	SLEEPDEEP_PAUSE	Pauses/runs this counter when the corresponding processor is in SLEEPDEEP 0: Counter runs normally regardless of processor mode. 1: Counter pauses when corresponding processor is in SLEEPDEEP.
MCWDTx_CTRy_CONFIG[31]	DEBUG_RUN	Pauses/runs this counter while a debugger is connected 0: Counter pauses according to DEBUG_TRIGGER_EN configuration. 1: Counter runs normally when debugger connected.
MCWDTx_CTRy_CTL[31]	ENABLE	Enable or disable the watchdog reset. 0: Counter is disabled (not clocked) 1: Counter is enabled (counting up)
MCWDTx_CTRy_CTL[0]	ENABLED	Indicates actual state of watchdog
MCWDTx_CTRy_CNT[15:0]	CNT	Current value of subcounter for this MCWDT
MCWDTx_CTRy_LOWER_LIMIT[15:0]	LOWER_LIMIT	Lower limit for watchdog
MCWDTx_CTRy_UPPER_LIMIT[15:0]	UPPER_LIMIT	Upper limit for watchdog
MCWDTx_CTRy_WARN_LIMIT[15:0]	WARN_LIMIT	Warn limit for watchdog
MCWDTx_LOCK[1:0]	MCWDT_LOCK	Prohibits writing control and configuration registers related to this MCWDT when not equal to 0. 0: No effect 1: Clear bit 0 2: Clear bit 1 3: Set both bit 0 and 1 (lock enabled)
MCWDTx_INTR[0]	CTR0_INT	MCWDT Interrupt Request for Subcounter 0
MCWDTx_INTR[1]	CTR1_INT	MCWDT Interrupt Request for Subcounter 1
MCWDTx_INTR_MASK[0]	CTR0_INT	Mask for Subcounter 0 for warning interrupt 0: MCWDT interrupt is masked to CPU. 1: MCWDT interrupt is not masked to CPU.
MCWDTx_INTR_MASK[1]	CTR1_INT	Mask for Subcounter 1 for warning interrupt 0: MCWDT interrupt is masked to CPU. 1: MCWDT interrupt is not masked to CPU.

20.4.2.2 32-bit Counter Operation

The Subcounter 2 (MCWDTx_CNT2) is a 32-bit free-running counter that can be configured to generate an interrupt. The MCWDTx_CTR2_CNT register holds the current value of Subcounter 2. Subcounter 2 does not support the window mode. However, it can be configured to generate an interrupt when one of the counter bits toggles. The BITS[20:16] bit field of the MCWDTx_CTR2_CONFIG register selects the bit on which the Subcounter 2 interrupt is asserted. ACTION bit [0] of the MCWDTx_CTR2_CONFIG register decides whether to assert an interrupt on bit toggle or not. [Figure 20-9](#) shows the Subcounter 2 counter operation.

Figure 20-9. Subcounter 2 Operation




Table 20-3. MCWDT Subcounter 2 Configuration Options

Register [Bit_Pos]	Bit Name	Description
MCWDTx_CTR2_CONFIG[0]	ACTION	Action taken when the specified BIT toggles 0: Nothing 1: Trigger an interrupt
MCWDTx_CTR2_CONFIG[20:16]	BITS	Bit to observe for a toggle: 0: Do ACTION after CTR2_CNT[0] toggles (every tick) ... 31: Do ACTION after CTR2_CNT[31] toggles (every 2³¹ ticks)
MCWDTx_CTR2_CONFIG[28]	DEBUG_TRIGGER_EN	Enables the trigger input for the MCWDT to pause the counter in debug mode. 0: Pauses the counter when a debug probe is connected. 1: Pauses the counter when a debug probe is connected and the trigger input is high.
MCWDTx_CTR2_CONFIG[30]	SLEEPDEEP_PAUSE	Pauses/runs this counter when the corresponding processor is in SLEEPDEEP 0: Counter runs normally regardless of processor mode. 1: Counter pauses when corresponding processor is in SLEEPDEEP.

Table 20-3. MCWDT Subcounter 2 Configuration Options

Register [Bit_Pos]	Bit Name	Description
MCWDTx_CTR2_CONFIG[31]	DEBUG_RUN	Pauses/runs this counter while a debugger is connected 0: Counter pauses according to DEBUG_TRIGGER_EN configuration. 1: Counter runs normally when debugger connected.
MCWDTx_INTR[2]	CTR2_INT	MCWDT Interrupt Request for Subcounter 2
MCWDTx_INTR_MASK[2]	CTR2_INT	MCWDT Interrupt Mask Register for Subcounter 2 0: MCWDT interrupt is masked to CPU. 1: MCWDT interrupt is not masked to CPU.
MCWDTx_INTR_MASKED[2]	CTR2_INT	MCWDT Interrupt Masked Register for Subcounter 2. Logical AND of corresponding request and mask bits

20.4.3 Enabling and Disabling MCWDT

The MCWDT counters are enabled by setting the ENABLE[31] bit in the MCWDTx_CTRy_CTL and MCWDTx_CTR2_CTL registers and are disabled by clearing it. Enabling or disabling a MCWDT counter requires two LFCLK cycles to come into effect. Therefore, the ENABLE bit value must not be changed more than once in that period and the ENABLED[0] bit of the MCWDTx_CTRy_CTL and MCWDTx_CTR2_CTL registers can be used to monitor the enabled/disabled state of the counter. The CTR0_SERVICE[0] and CTR1_SERVICE[1] bits of the MCWDTx_SERVICE register clears the corresponding subcounter when set in firmware. The hardware clears the bit after the MCWDT counter resets. This option is useful when Subcounter 0 or Subcounter 1 is configured to generate a device reset after a FAULT event. After the MCWDT counter is enabled, it is not recommended to write to the MCWDT configuration (MCWDTx_CTRy_CONFIG and MCWDTx_CONFIG) and control (MCWDTx_CTRy_CTL and MCWDTx_CTR2_CTL) registers. Accidental corruption of MCWDT registers can be prevented by setting the MCWDT_LOCK[1:0] bit of the MCWDTx_LOCK register. If the application requires updating any register while the WDT is running, the MCWDT_LOCK bits must be cleared. The MCWDT_LOCK bits require two different writes to clear both the bits. Writing a '1' to the bits clears bit 0. Writing a '2' clears bit 1. Writing a '3' sets both the bits and writing '0' does not have any effect. Note that the MCWDT_LOCK bits are only protecting following registers:

- MCWDTx_CTRy_CTL
- MCWDTx_CTRy_LOWER_LIMIT
- MCWDTx_CTRy_UPPER_LIMIT
- MCWDTx_CTRy_WARN_LIMIT
- MCWDTx_CTRy_CONFIG
- MCWDTx_CTRy_CNT
- MCWDTx_CTR2_CTL
- MCWDTx_CTR2_CONFIG
- MCWDTx_CTR2_CNT
- MCWDTx_SERVICE

Table 20-4. Watchdog Configuration Options

Register [Bit_Pos]	Bit Name	Description
MCWDTx_CTRy_CTL[31] MCWDTx_CTL[31]	ENABLE ENABLE	Enable or disable the watchdog reset ■ 0: Counter is disabled (not clocked) ■ 1: Counter is enabled (counting up)
MCWDTx_CTRy_CTL[0] MCWDTx_CTL[0]	ENABLED ENABLED	Indicates actual state of watchdog
MCWDTx_LOCK[1:0]	MCWDT_LOCK	Locks or unlocks write access to the MCWDT registers. When the bits are set, the lock is enabled. ■ 0: No effect ■ 1: Clears bit 0 ■ 2: Clears bit 1 ■ 3: Sets both bit 0 and 1 (lock enabled)
MCWDTx_SERVICE[0] MCWDTx_SERVICE[1]	CTR0_SERVICE CTR1_SERVICE	Services Subcounter 0. This resets the count value for Subcounter 0 to zero Services Subcounter 1. This resets the count value for Subcounter 1 to zero

Note: When the watchdog counters are configured to generate an interrupt every LFCLK cycle, make sure you read the MCWDTx_INTR register after clearing the watchdog interrupt (setting the CTR0_INT, CTR1_INT, and CTR2_INT bits in the MCWDTx_INTR register). Failure to do this may result in missing the next interrupt. Hence, the interrupt cycle will become LFCLK/2.

20.4.4 Watchdog Reset

Subcounter 0 and Subcounter 1 can be configured to generate a device reset similar to the basic WDT reset. Follow these steps to use Subcounter 0 or Subcounter 1 of an MCWDT block to generate a system reset. Note that a reset is asserted after an unhandled FAULT condition. The subcounters can be individually configured whether to generate only a FAULT, or a reset after a FAULT event.

1. Configure the MCWDT to generate a reset by setting MCWDTx_CTRy_CONFIG.LOWER_ACTION[1:0] or MCWDTx_CTRy_CONFIG.UPPER_ACTION[5:4] bits in the MCWDTx_CTRy_CONFIG register to '2'.
2. Calculate the watchdog reset period such that firmware is able to reset the watchdog at least once during the period, even along the longest firmware delay path, and write the value into the MCWDTx_CTRy_UPPER_LIMIT register. In window mode define an adequate MCWDTx_CTRy_LOWER_LIMIT, which cannot be violated by the shortest firmware delay path.
3. Enable MCWDT by setting the ENABLE[31] bit in the MCWDTx_CTRy_CTL register. Wait until the ENABLED[0] bit is set.
4. Lock the MCWDT configuration by setting the MCWDT_LOCK bits of the MCWDTx_LOCK register to '3'.
5. In the firmware, feed (reset) the watchdog by writing '1' into the CTR0_SERVICE[0] or CTR1_SERVICE[1] bit in the MCWDTx_SERVICE register.

It is not recommended to reset watchdog in the MCWDT ISR.

20.4.5 Watchdog Interrupt

When configured to generate an interrupt, the CTR0_INT (Subcounter 0), CTR1_INT (Subcounter 1), and CTR2_INT (Subcounter 2) bits of the MCWDTx_INTR register provide the status of any pending watchdog interrupts. The firmware must clear the interrupt by setting the same bit to '1'. The CTR0_INT, CTR1_INT, and CTR2_INT bits of the MCWDTx_INTR_MASK register unmask the corresponding MCWDT interrupt to the CPU.

Follow these steps to use MCWDT as a periodic interrupt generator:

1. Write the desired warning threshold value to the WARN_LIMIT register for Subcounter 0 and Subcounter 1 or the BITS[20:16] value to the MCWDTx_CTR2_CONFIG register for Subcounter 2.
2. For Subcounter 0 and Subcounter 1 configure the MCWDT to generate an interrupt using the MCWDTx_CTRy_CONFIG.WARN_ACTION[8] bit in MCWDTx_CTRy_CONFIG register. For Subcounter 2, set the ACTION[0] bit in the MCWDTx_CTR2_CONFIG register.
3. Set the CTR0_INT, CTR1_INT, and CTR2_INT bits in MCWDTx_INTR to clear any pending interrupt.
4. Set the MCWDTx_CTRy_CONFIG.AUTO_SERVICE[12] bit in MCWDTx_CTRy_CONFIG for Subcounter 0 and Subcounter 1 to reset the corresponding watchdog counter to '0' on a warning interrupt event. **Note:** For Subcounter 2, no automatic counter clearing is supported.
5. Unmask the MCWDT interrupt to the CPU by setting the CTRx_INT bit in the MCWDTx_INTR_MASK register.

6. Enable MCWDT by setting the ENABLE[31] bit in the MCWDTx_CTRy_CTL register. Wait until the ENABLED[0] bit is set.
7. Enable the MCWDT interrupt to the CPU by configuring the appropriate ISER register. Refer to the [Interrupts chapter on page 144](#).
8. In the ISR, clear the MCWDT interrupt by setting the CTRx_INT bit in the MCWDTx_INTR register.

Note that interrupts from all three subcounters of the MCWDT block are mapped as a single interrupt to the CPU.

In the interrupt service routine, the CTRx_INT bits of the MCWDTx_INTR register can be read to identify the interrupt source. However, each MCWDT block has its own interrupt to the CPU. For details on interrupts, see the [Interrupts chapter on page 144](#). The MCWDT block can send interrupt requests to the CPU in Active power mode and to the WIC in Sleep and DeepSleep power modes. It works similar as the basic WDT.

The hardware does not support changing the timeout for DeepSleep mode. However, Subcounters 0 and Subcounter 1 can work together to get a similar behavior. Subcounter 0 can be configured with a timeout threshold suitable to protect running firmware and configured to stop during DeepSleep. Subcounter 1 can be configured with a longer timeout that continues to operate in DeepSleep. For this usage example to work, the early window thresholds should be the same (if window mode is enabled) and firmware should service both these subcounters at the same time.

20.5 Reset Cause Detection

The RESET_WDT bit [0] in the RES_CAUSE register indicates the reset generated by the basic WDT. The RESET_MCWDT0 [5], RESET_MCWDT1 [6], RESET_MCWDT2 [7], and RESET_MCWDT3 [8] bits in the RES_CAUSE register indicate the reset generated by the MCWDTx block. These bits remain set until cleared or until a power-on reset (POR), brownout reset (BOD), or external reset (XRES_L) occurs. All other resets leave the bits unaltered. For more details, see the [Reset System chapter on page 217](#).

20.6 Debug Mode

For both types of WDTs, watchdog resets are automatically blocked by hardware during debugging, and window mode is automatically paused. By default, all the WDTs also stop counting. Two configuration bits (per WDT) configure the behavior of the counter when the debugger is connected. The recommended procedure to disconnect a debug probe is to service any active watchdog timers using the debug probe, then disconnect the probe. The firmware begins running again and the next service will realign the window and resume normal window operation.

In a multi-core environment, 'debug state' indicates that at least one of the CPUs is in debug state. If one CPU is debugged but another or multiple other CPUs are continuously running, then the user can configure the counter via the debugger to continue or pause depending on which CPU is using the counter.

The configuration is done with DEBUG_TRIGGER_ENABLE[28] and DEBUG_RUN[31] bits, which are both located in the related CONFIG register for basic WDT and MCWDT. [Table 20-5](#) shows the configuration options.

Table 20-5. Debug Modes

DEBUG_RUN	DEBUG_TRIGGER_ENABLE	Description
0	0	Counter is stopped when a debugger is connected.
0	1	Counter is stopped only when a debugger is connected and the CPU is halted during a breakpoint.
1	x	Counter is running when debugger is connected. No reset is issued when the CPU is halted during a breakpoint but the counter is not stopped.

Note that in each case, no reset and no FAULT is issued when the debugger is connected to the target system.

To pause at a breakpoint while debugging, configure the trigger matrix to connect the related CPU halted signal to the trigger input for the related watchdog timer. It takes up to two LFCLK cycles for the trigger signal to be processed. Triggers that are less than two LFCLK cycles may be missed. Synchronization errors can accumulate each time it is halted.

Note that it may take up to two ILO0 (or LFCLK for MCWDT) clock cycles for the counter to pause, due to internal synchronization. After the debugger is disconnected, the MCWDTx_CTRy_CONFIG.LOWER_ACTION is ignored until after the first service. This prevents an unintentional trigger of the MCWDTx_CTRy_CONFIG.LOWER_ACTION before the firmware realigns the servicing period. After the first service, MCWDTx_CTRy_CONFIG.LOWER_ACTION behaves as configured.

20.7 CPU Select

In a multi-core system it is recommended to assign one MCWDT to a dedicated CPU to select the SLEEPDEEP signal to be used to control the counter in SleepDeep power mode. The counter pauses in SleepDeep power mode in case SLEEPDEEP_PAUSE[30] bit is set to '1' in CTR2_CONFIG register.

A single MCWDT is not intended to be used simultaneously by multiple CPUs because of the complexity involved in coordination.

CPU_SEL[1:0] bits in the CPU_SELECT register are defined in [Table 20-6](#).

Table 20-6. MCWDT Assignment to the Cores

CPU_SEL[1:0]	CPU
0	CM0+
1	CM7_0
2	CM7_1

20.8 Register List

Table 20-7. WDT Registers

Register	Name	Description
WDT_CTL	Watchdog Control Register	Control register for the basic WDT.
WDT_LOWER_LIMIT	WDT Lower Limit Register	Lower limit for the basic WDT.
WDT_UPPER_LIMIT	WDT Upper Limit Register	Upper limit for the basic WDT.
WDT_WARN_LIMIT	WDT Warn Limit Register	Warn limit for the basic WDT.
WDT_CONFIG	WDT Configuration Register	Configuration for the basic WDT. Includes the ACTION configuration for Upper, Lower, and Warn limits, auto-servicing, and pause settings in low-power and debug modes.
WDT_CNT	WDT Count Register	Count value for the basic WDT.
WDT_LOCK	WDT Lock Register	Lock or unlock the basic WDT registers.
WDT_SERVICE	WDT Service Register	Clears the basic WDT counter.
WDT_INTR	WDT Interrupt Register	Interrupt signal from basic WDT
WDT_INTR_SET	WDT Interrupt Set Register	Sets interrupts for firmware testing.
WDT_INTR_MASK	WDT Interrupt Mask Register	Controls whether interrupt is forwarded to CPU. All masks block the interrupt when 0 and forward the interrupt when 1.
WDT_INTR_MASKED	WDT Interrupt Masked Register	Bitwise AND between the interrupt request and mask registers so firmware can read the status of all mask enabled interrupt causes with a single load operation
MCWDTx_CTRy_CTL	MCWDT Subcounter 0/1 Control Register	Control register for MCWDT subcounter.
MCWDTx_CTRy_LOWER_LIMIT	MCWDT Subcounter 0/1 Lower Limit Register	Lower limit for this MCWDT subcounter.
MCWDTx_CTRy_UPPER_LIMIT	MCWDT Subcounter 0/1 Upper Limit Register	Upper limit for this MCWDT subcounter.
MCWDTx_CTRy_WARN_LIMIT	MCWDT Subcounter 0/1 Warn Limit Register	Warn limit for this MCWDT subcounter.
MCWDTx_CTRy_CONFIG	MCWDT Subcounter 0/1 Configuration Register	Configuration for this MCWDT subcounter. Includes the ACTION configuration for Upper, Lower, and Warn limits
MCWDTx_CTRy_CNTy	MCWDT Subcounter 0/1 Count Register	Count value for this MCWDT subcounter.
MCWDTx_CTR2_CTL	MCWDT Subcounter 2 Control Register	Control register for MCWDT subcounter 2.
MCWDTx_CTR2_CONFIG	MCWDT Subcounter 2 Configuration Register	Configuration for MCWDT subcounter 2.

Table 20-7. WDT Registers

Register	Name	Description
MCWDTx_CTR2_CNT	MCWDT Subcounter 2 Count Register	Count value for this MCWDT subcounter 2.
MCWDTx_LOCK	MCWDT Lock Register	Lock or unlock the respective configuration registers of subcounters 0/1/2 of this MCWDT.
MCWDTx_SERVICE	MCWDT Service Register	Includes service bits to clear subcounter 0/1 of this MCWDT.
MCWDTx_INTR	MCWDT Interrupt Register	Interrupt status register for subcounters 0/1/2 for this MCWDT.
MCWDTx_INTR_SET	MCWDT Interrupt Set Register	Triggers an interrupt for firmware testing.
MCWDTx_INTR_MASK	MCWDT Interrupt Mask Register	Controls whether a subcounter interrupt is forwarded to the corresponding processor. All masks block the interrupt when 0 and forward the interrupt when 1.
MCWDTx_INTR_MASKED	MCWDT Interrupt Masked Register	Bitwise AND between the interrupt request and mask registers so firmware can read the status of all mask enabled interrupt causes with a single load operation.
CLK_SELECT	Clock Selection Register	Clock source selection register.
CLK_ILO0_CONFIG	ILO0 Configuration	ILO0 configuration
RES_CAUSE	Reset Cause Observation Register	Reset cause observation register

Note: In MCWDTx_CTRy, 'x' signifies the instance and 'y' signifies the subcounter (0/1). Refer to the device datasheet or the Registers TRM for more information.

21. Real-Time Clock



The Real-Time Clock (RTC) system is an “always-on” function, which is a part of the Backup domain. It contains a real-time clock with alarm feature, supported by a 32768-Hz watch crystal oscillator (WCO), low-power external crystal oscillator (LPECO)¹ for 4 MHz to 8 MHz crystal and Backup registers.

Backup is not a power mode; the Backup domain always runs on VDDD. For more details, see the [Power Supply and Monitoring chapter on page 163](#), the [Device Power Modes chapter on page 186](#), and the [Clocking System chapter on page 197](#) for WCO and LPECO.


21.1 Features

- Fully-featured RTC
 - Year/Month/Date, Day-of-Week, Hour : Minute : Second fields (All fields Integer)
 - Supports both 12-hour and 24-hour formats
 - Automatic leap year correction
- Configurable alarm function
 - Alarm on Month/Date, Day-of-Week, Hour : Minute : Second fields
 - Two independent alarms
- Calibration for 32768-Hz WCO and 4 MHz to 8 MHz LPECO
- Calibration waveform output
 - Supports 512 Hz, 1 Hz, and 2 Hz
- Backup registers

1. See the device datasheet to confirm whether LPECO is present.

21.2 Block Diagram

Figure 21-1. Block Diagram



The RTC system includes an accurate WCO or LPECO that can generate the required clock with the help of an external crystal or external clock inputs. The RTC has a programmable alarm feature, which can generate interrupts to the CPU. An AHB-Lite interface provides firmware access to registers interface in the Backup domain.

The VDDBAK in the Backup domain is always supplied from VDDD.

The domain also has a backup registers block, which can retain its contents even when the device enters Hibernate or XRES mode. The RTC system can also output a calibration waveform.

21.3 Power Supply

Power to the RTC system is VDDD (unregulated main supply). See the [Power Supply and Monitoring chapter on page 163](#) for more details.

It is possible to monitor the Backup domain voltage (VDDD) using the low-voltage detect (LVD) feature of Traveo II. For more information on LVD, see the [Power Supply and Monitoring chapter on page 163](#).

21.4 Clocking

The RTC primarily runs from a 32768-Hz clock, after it is scaled down to one-second ticks. This clock signal can come from either of these internal sources:

- Watch-crystal oscillator (WCO). This is a high-accuracy clock generator that is suitable for RTC applications and requires a 32768-Hz external crystal populated on the application board. WCO can also operate without crystal, using external clock/sine wave inputs. These additional operating modes are explained in the [Clocking System chapter on page 197](#). WCO is supplied by the Backup domain.
- Low-power external crystal oscillator (LPECO). This is a 4-8 MHz crystal oscillator that can be fractionally divided to 32768 Hz, and then used as a replacement for WCO. The LPECO key specifications are explained in the [Clocking System chapter on page 197](#).
- Alternate Backup Clock (ALTBK): This option allows the use of CLK_LF generated by the SRSS as the Backup domain clock. Note that CLK_LF is not always available in all device power modes. See the [Device Power Modes chapter on page 186](#) for more details. CLK_LF is described in the [Clocking System chapter on page 182](#). Clock glitches can propagate into the RTC system when CLK_LF is enabled or disabled by the SRSS. In addition, CLK_LF may not be as accurate as WCO or LPECO depending on the actual source of CLK_LF. Because of these reasons, CLK_LF is not recommend for RTC applications. Also, if the WCO or LPECO is intended as the clock source then choose it directly instead of routing through CLK_LF.
- Internal Low-frequency Oscillator (ILO): This option allows the use of ILO0. ILO0 is described in the [Clocking System chapter on page 197](#).

For more details on these clocks and calibration, see the [Clocking System chapter on page 197](#).

The RTC clock source can be selected using the `BACKUP_CTL.CLK_SEL` bit. The `BACKUP_CTL.WCO_EN` bit can be used to enable or disable the WCO. If the WCO operates with an external crystal, make sure the `BACKUP_CTL.WCO_BYPASS` bit is cleared before enabling the WCO. In addition, the `BACKUP_CTL.PRESCALER` bit must be configured for a prescaler value of 32768. The `BACKUP_LPECO_CTL.LPECO_EN` can be used to enable or disable the LPECO.

Note: External crystal and bypass capacitors of proper values must be connected to `WCO_IN` and `WCO_OUT` pins or `LPECO_IN` and `LPECO_OUT` pins. See the device datasheet for details of component values and electrical connections. In addition, GPIOs must be configured for `WCO_OUT` and `WCO_IN` signals or `LPECO_OUT` and `LPECO_IN` signals. See the [I/O System chapter on page 246](#) to know how to configure the GPIOs.

Note: If WCO is used as an RTC clock, then it is important to make sure that the WCO is running stable; that is, wait for `BACKUP_STATUS.WCO_OK`, before writing to the RTC registers.

Note: If LPECO is used as an RTC clock, make sure that the LPECO is running stable; that is, wait for `BACKUP_LPECO_STATUS.LPECO_READY`, before writing to the RTC registers.

21.5 Reset

To keep the RTC operating through resets, the Backup domain should not be reset under most circumstances.

The RTC initializes itself at power up; it cannot be reset by other internal and external resets such as BOD reset, OVD, OCD, WDT, and `XRES_L`. See the [Reset System chapter on page 217](#) for more details.

The RTC system is reset only when all the power supplies are removed from the Backup domain. Also, user firmware can reset the RTC system logic by using `BACKUP_RESET.RESET`.

If `RES_CAUSE` reports a BOD/OVD/OCD event, user firmware should initialize the RTC system by writing `BACKUP_RESET.RESET=1` because faulty supplies may have corrupted the Backup domain contents.

If `RES_CAUSE` reports a `XRES_L` or WDT event, it is an application-specific decision whether to trust the Backup domain contents.

Although rare, `XRES_L`/WDT may mean that the Backup domain contents were corrupted by faulty user firmware execution or by interrupting an AHB write to the backup logic.

21.6 Real-Time Clock

The RTC consists of seven integer fields and one control bit as shown in the following table:

Table 21-1. RTC Fields

Bit Field Name	Description
RTC_SEC	Calendar seconds, value range = 0-59
RTC_MIN	Calendar minutes, value range = 0-59
RTC_HOUR	Calendar hours, value depends on 12-hour or 24-hour format set in the <code>BACKUP_RTC_TIME.CTRL_12HR</code> bit. In 12-hour mode, bit <code>BACKUP_RTC_TIME.RTC_HOUR[4]</code> = 0 for AM and 1 for PM, bits <code>BACKUP_RTC_TIME.RTC_HOUR[3:0]</code> = 1–12 In 24-hour mode, bits <code>BACKUP_RTC_TIME.RTC_HOUR[4:0]</code> = 0–23
CTRL_12HR	Select the 12-hour or 24-hour mode: 1=12HR, 0=24HR
RTC_DAY	Calendar day of the week, value range = 1-7 The user should define the meaning of the values
RTC_DATE	Calendar day of the month, value range = 1-31 Automatic leap year correction until 2400
RTC_MON	Calendar month, value range = 1-12
RTC_YEAR	Calendar year, value range = 0-99

RTC value fields indicate an integer format. Constant bits are omitted in the RTC implementation. For example, the maximum `BACKUP_RTC_TIME.RTC_SEC` is 59, which can be represented as one byte 0b00111011. However, the most significant bit is always zero and is therefore omitted, making the `BACKUP_RTC_TIME.RTC_SEC` a 6-bit field.

The RTC supports both 12-hour format with AM/PM flag, and 24-hour format for the “hours” field. The RTC also includes a “day of the week” field, which counts from 1 to 7. The user should define which weekday is represented by a value of ‘1’.

The RTC implements automatic leap year correction for the Date field (day of the month). If the Year is divisible by four, the month of February (Month=2) will have 29 days instead of 28. When the Year field rolls over from 99 to 00, the firmware should update the otherwise static century value and therefore an interrupt is raised. This interrupt is called the century interrupt.

User registers containing these bit fields are `BACKUP_RTC_TIME` and `BACKUP_RTC_DATE`. See the corresponding register descriptions in the *Traveo II Body Controller High Registers TRM* for details.

As the user registers are in the high-frequency bus-clock domain and the actual RTC registers run from the

low-frequency 32768-Hz clock, reading and writing RTC registers require special care. These processes are explained in the following sections.

21.6.1 Reading RTC User Registers

To start a read transaction, the firmware should set the `BACKUP_RTC_RW.READ` bit. When this bit is set, the RTC registers will be copied to user registers and frozen so that a coherent RTC value can safely be read by the firmware. The read transaction is completed by clearing the `BACKUP_RTC_RW.READ` bit.

`BACKUP_RTC_RW.READ` bit cannot be set if:

- RTC is still busy with a previous operation (that is, the `BACKUP_STATUS.RTC_BUSY` bit is set)
- The `BACKUP_RTC_RW.WRITE` bit is set

The firmware should verify that the above bits are not set before setting the `BACKUP_RTC_RW.READ` bit.

21.6.2 Writing to RTC User Registers

When the `BACKUP_RTC_RW.WRITE` bit is set, data can be written into the RTC user registers; otherwise, writes to the RTC user registers are ignored. When all the RTC writes are done, the firmware needs to clear the `BACKUP_RTC_RW.WRITE` bit for the RTC update to take effect. After the `BACKUP_RTC_RW.WRITE` bit is cleared, the hardware will copy all the new data on one single WCO clock edge to ensure coherency to the actual RTC registers.

`BACKUP_RTC_RW.WRITE` bit cannot be set if:

- RTC is still busy with a previous operation (that is, the `BACKUP_STATUS.RTC_BUSY` bit is set)
- `BACKUP_RTC_RW.READ` bit is set

The firmware should make sure that the values written to the RTC fields form a coherent legal set. The hardware does not check the validity of the written values. Writing illegal values results in undefined behavior of the RTC.

When in the middle of an RTC update with the `BACKUP_RTC_RW.WRITE` bit set, and a brownout, reset, or entry to DeepSleep or Hibernate mode occurs, the write operation will not be complete. This is because the `BACKUP_RTC_RW.WRITE` bit will be cleared by a reset, and the RTC update is only triggered when this bit is cleared by an AHB WRITE transaction. If the write operation is in progress (`BACKUP_STATUS.RTC_BUSY`), data corruption can occur if the system is reset or enters DeepSleep or Hibernate mode.

To update only one or a few of the RTC fields, for example, when the RTC is adjusted for daylight saving time (DST), then only the Hour field needs an update, although the Seconds and Minutes fields should not be disturbed – they should continue running. For that reason, an ‘Update’ flag is maintained for each RTC field. Only those fields that have been updated will be copied to the actual RTC when the `BACKUP_RTC_RW.WRITE` bit is cleared.

21.7 WCO/LPECO Calibration

It is possible to improve the accuracy of the RTC by calibrating the WCO or LPECO. The `CLK_LF` can also be calibrated. See the [Clocking System chapter on page 197](#) for details.

The WCO or LPECO accuracy is affected by an absolute crystal accuracy. This occurs because the crystal itself oscillates slightly faster or slower due to imperfect manufacturing. The user firmware can calibrate the RTC accuracy. The calibration bit fields are as follows:

Table 21-2. Calibration Bit Fields

Bit Field Name	Description
<code>CALIB_VAL</code>	Calibration value for absolute frequency. Each step causes 128 ticks to be added or removed each hour.
<code>CALIB_SIGN</code>	0: Negative sign: remove pulses (it takes more clock ticks to count one second) 1: Positive sign: add pulses (it takes less clock ticks to count one second)
<code>CAL_SEL</code>	Select calibration wave output signal 0: 512-Hz wave, not affected by calibration setting (not supported for 50/60-Hz input clock) 1: Reserved 2: 2-Hz wave, includes the effect of the calibration setting (not supported for 50/60-Hz input clock) 3: 1-Hz wave, includes the effect of the calibration setting (supported for all input clocks)
<code>CAL_OUT</code>	Output enable for wave signal for calibration, and allow <code>BACKUP_CAL_CTL.CALIB_VAL</code> to be written.

21.7.1 Absolute Accuracy Calibration

To measure the WCO or LPECO error, the `CAL_OUT` bit must be set; this will cause a clock derived from the RTC system to be output on the `RTC_CAL` pin. The user should measure the deviation from 512 Hz, convert that to a ppm

value, and derive the calibration settings to be used to correct the error.

The calibration correction is done by either adding or removing pulse counts from the oscillator divider each hour, which respectively speeds up or slows down the clock. After a calibration starts, it is performed hourly; it is applied as 64

ticks every 30 seconds until there are $2 \times \text{BACKUP\_CAL\_CTL.CALIB\_VAL}$ adjustments.

Because this is digital calibration, changing the calibration value does not affect the 512 Hz calibration output clock signal. Traveo II supports two others calibration waveform frequencies; 1 Hz and 2 Hz. However, those calibration waveforms are affected by the current calibration.

The calibration register can only be written when the `BACKUP_RTC_RW.WRITE` bit is set. See [21.6.2 Writing to RTC User Registers](#).

21.8 Alarm Feature

The Alarm feature allows the RTC to generate an interrupt, which may be used to wake up the system from Sleep, DeepSleep, and Hibernate power modes. The Alarm feature consists of six fields corresponding to the fields of the RTC: Month/Date, Day-of-Week, and Hour: Minute: Second. Each Alarm field has an enable bit that needs to be set to enable matching; if the bit is cleared, then the field will be ignored for matching. [Table 21-3](#) shows the Alarm bit fields.

Table 21-3. Alarm Bit Fields

Bit Field Name	Description
ALM_SEC	Alarm seconds, value range = 0-59
ALM_SEC_EN	Alarm second enable: 0=disable, 1=enable
ALM_MIN	Alarm minutes, value range = 0-59
ALM_MIN_EN	Alarm minutes enable: 0=disable, 1=enable
ALM_HOUR	Alarm hours, value depending on the 12-hour or 24-hour mode. In 12-hour mode, bit <code>BACKUP_ALMx_TIME.ALM_HOUR[4]</code> = 0 for AM and 1 for PM, bits <code>BACKUP_ALMx_TIME.ALM_HOUR[3:0]</code> = 1–12 In 24-hour mode, bits <code>BACKUP_ALMx_TIME.ALM_HOUR[4:0]</code> = 0–23
ALM_HOUR_EN	Alarm hour enable: 0=disable, 1=enable
ALM_DAY	Calendar day of the week, value range = 1-7 The user should define the meaning of the values
ALM_DAY_EN	Alarm day of the week enable: 0=disable, 1=enable
ALM_DATE	Alarm day of the month, value range = 1-31
ALM_DATE_EN	Alarm day of the month enable: 0=disable, 1=enable
ALM_MON	Alarm month, value range = 1-12
ALM_MON_EN	Alarm month enable: 0=disable, 1=enable
ALM_EN	Master enable for alarm. 0: Alarm is disabled. Fields for date and time are ignored. 1: Alarm is enabled. If none of the date and time fields are enabled, then this alarm triggers once every second.

If the master enable (`BACKUP_ALMx_DATE.AL_M_EN`) is set, but all alarm fields for date and time are disabled, an alarm interrupt will be generated once every second. Note that there is no alarm field for Year because the life expectancy of a chip is about 20 years. Thus, setting an alarm for a certain year indicates that the alarm matches either once or never in the lifetime of the chip.

Traveo II has two independent alarms. See the `BACKUP_ALM1_TIME`, `BACKUP_ALM1_DATE`, `BACKUP_ALM2_TIME`, and `BACKUP_ALM2_DATE` registers in the *Traveo II Body Controller High Registers TRM* for details.

Note that the alarm user registers, similar to RTC user registers, require the same steps for read/write operations, as explained in [21.6.1 Reading RTC User Registers](#) and [21.6.2 Writing to RTC User Registers](#).

Interrupts must be properly configured for the RTC to generate interrupts/wakeup events. Also, to enable RTC interrupts to wake up the device from Hibernate mode, the `PWR_HIBERNATE.MASK_HIBALARM` bit must be set. See the [Device Power Modes chapter on page 186](#) and the [Interrupts chapter on page 144](#) for details.

`BACKUP_INTR_MASK` register can be used to disable certain interrupts from the RTC system.

Table 21-4. Interrupt Mask Bits

Bit Name	Description
ALARM1	Mask bit for interrupt generated by ALARM1
ALARM2	Mask bit for interrupt generated by ALARM2
CENTURY	Mask bit for century interrupt, generated when the Year field rolls over from 99 to 00

21.9 Backup Registers

The RTC system has several registers (BACKUP_BREGx), which can be used to store important information/flags. This includes information that need to be retained when the device enters Hibernate mode. For the number of BACKUP_BREGx registers, see the *Traveo II Registers TRM*.

21.10 Real Time Clock Registers

Table 21-5. Backup Registers

Register	Name	Description
BACKUP_CTL	Control register	This register provides several settings of RTC operation. This register is hold in all device power modes including Sleep, Low-Power Sleep, DeepSleep and Hibernate.
BACKUP_RTC_RW	RTC read write register	This register provides read and write control function. This register is reset in DeepSleep.
BACKUP_CAL_CTL	Oscillator calibration control register	This register provides oscillator calibration for absolute frequency.
BACKUP_STATUS	Status register	This register provides status of the RTC System. Firmware must monitor these bits to execute some operation. This register is hold in all device power modes including Sleep, Low-Power Sleep, DeepSleep and Hibernate.
BACKUP_RTC_TIME	RTC time register	This register provides calendar seconds, minutes, hours, and day of week.
BACKUP_RTC_DATE	RTC date register	This register provides calendar day of month, month, and year.
BACKUP_ALM1_TIME	Alarm1 time register	This register provides Alarm 1 seconds, minute, hours, and day of week.
BACKUP_ALM1_DATE	Alarm1 date register	This register provides Alarm 1 day of month, and month.
BACKUP_ALM2_TIME	Alarm2 time register	This register provides Alarm 2 seconds, minute, hours, and day of week.
BACKUP_ALM2_DATE	Alarm2 date register	This register provides Alarm 2 day of month, and month.
BACKUP_INTR	Interrupt request register	This register holds Interrupt signals. This register is sets by hardware if Interrupts condition occur. Firmware can clear these bits with writing '1'.
BACKUP_INTR_SET	Interrupt set request register	This register is for firmware testing. Interrupts occur if firmware set '1' to these bits. (For firmware testing purpose)
BACKUP_INTR_MASK	Interrupt mask register	This register provides Interrupt mask. When Mask bit is set, the interrupt is enabled.
BACKUP_INTR_MASKED	Interrupt masked request register	This register allows the firmware to read the status of all mask-enabled interrupt causes with a single load operation, rather than two load operations: one for the interrupt causes and one for the masks. This simplifies firmware development.
BACKUP_BREGx	Backup register	These registers provide backup register regions. 'x' signifies the number of backup registers.
BACKUP_RESET	RTC system reset register	This register is used to reset the RTC system from firmware.
BACKUP_LPECO_CTL	LPECO control register	This register configures LPECO.
BACKUP_LPECO_STATUS	LPECO status register	This register indicates status for LPECO.
BACKUP_LPECO_PRESCALE	LPECO prescaler register	This register configures LPECO prescaler.

Note: 'x' signifies the number of backup register. Refer to the Register TRM for more information.

Section D: Input/Output Subsystem Overview




This section encompasses the following chapters:

- [I/O System chapter on page 246](#)

Top Level Architecture

I/O System Block Diagram



22. I/O System



This chapter explains the TVII-B-H I/O system, its features, architecture, operating modes, and interrupts. The I/O system provides the interface between the CPU core and peripheral components. The flexibility of TVII-B-H devices and the capability of its I/O to route most signals to most pins simplifies circuit design and board layout. The GPIO pins are grouped into ports; a port can have a maximum of eight GPIOs.

This chapter describes the following:

- Features and overview
- I/O cell architecture
- GPIO port configuration, interrupt support, and software I/O functionality
- I/O subsystem
- Smart I/O

22.1 Features

The TVII-B-H family GPIOs have these features:

- Analog and digital input and output capabilities
- Eight drive strength modes
- Separate port read and write registers
- Edge-triggered interrupts on rising edge, falling edge, or on both the edges, on all GPIO
- Slew rate control
- Hold mode for latching previous state (used to retain the I/O state in DeepSleep mode)
- Selectable CMOS, TTL, and automotive input buffer mode
- Smart I/O provides the ability to perform Boolean functions in the I/O signal path

22.2 GPIO Interface Overview

Each of the GPIOs may fall into one of the following categories:


- GPIO cells provide a means for the CPU and peripherals to communicate off-chip. All GPIO cells are software-controllable and observable by the CPU. Some or all GPIO cells may be routed to one or more peripherals. A peripheral I/O signal may be routed to multiple GPIO cells; HSIOM control registers specify the active route connection.
- System function cells such as reset or power supplies.
- Application-specific I/O pins

Analog peripheral connectivity:

- Some GPIO cells have dedicated analog connections to programmable analog peripherals, such as SARMUX.

TVII-B-H is equipped with analog and digital peripherals. [Figure 22-1](#) shows an overview of the routing between the peripherals and pins.

Figure 22-1. GPIO Interface Overview




The device has several options for interfacing to external signals and devices operating from a different supply voltage.

TVII-B-H devices may optionally include a V_{DDIO} supply pin with a different voltage supply from the V_{DD} pin. Where included, the V_{DDIO} pin may be used to power some or all of the GPIO cells, providing a built-in level-translator capability.

22.3 I/O Cell Architecture

Figure 22-2 shows the I/O cell architecture present in every GPIO cell. It comprises an input buffer and an output driver that connect to the HSIOM multiplexers for digital input and output signals. Analog peripherals connect directly to the pin for point-to-point connections or use of the AMUXBUS.

Figure 22-2. GPIO Cell Architecture



The GPIO component provides the I/O cell configuration information through registers. These registers are retained in DeepSleep power mode, but are reset to their default value in Hibernate power mode. To allow for Hibernate Interrupt functionality, the I/O cells hold/freeze their configuration information when entering either DeepSleep or Hibernate power mode. As a result, the configuration signals can be routed in the Active power domain.

If the HSIOM makes a functional connection to an I/O cell, the GPIO provides the configuration information. If the HSIOM makes a test connection (scan, PTM or JTAG) to an I/O cell, the GPIO configuration information is ignored, and the HSIOM provides the required configuration information.

I/O cell configuration includes information such as drive mode (pull-up/pull-down) and drive strength. Configuration

information may be for a specific I/O pad: drive mode, drive strength, fast versus slow slew control transitioning, input buffer mode, and so on.

The I/Os in an I/O port are accessible by software to provide controllability of the I/O output signals and observability of the I/O input signals. Combined, controllability and observability provide software bit banging functionality.

Each I/O port has a `GPIO_PRTx_OUT` register field that specifies the data and data enable to be fed to the I/O cells output drivers. Each I/O cell has a dedicated 1-bit data and data enable field. Three additional registers are provided to ease/speedup software bit banging functionality. These registers allow software to manipulate individual I/O output signals without requiring a 'read modify-write' sequence. The `GPIO_PRTx_OUT_SET` register allows software to set

specific data/data enable fields to '1', without affecting the signal level of the other data fields. The GPIO_PRTx_OUT_CLR register allows software to set specific data/data enable fields to '0', without affecting the signal level of the other data fields. The MMIO_OUT_INV register allows software to invert the value of specific data/data enable fields, without affecting the signal level of the other data fields.

Note that the GPIO_PRTx_OUT_SET, GPIO_PRTx_OUT_CLR, and GPIO_PRTx_OUT_INV registers all operate on the OUT register data fields; no dedicated flip-flops are created for these registers.

Each I/O port has a GPIO_PRTx_IN (I/O cell input buffer state) register that reflects the I/O cells inputs. Note that the I/O cell inputs may be different from the data fed to the I/O cell output drives (GPIO_PRTx_OUT register).

The GPIO data input and data output/data output enable signals for I/O cells are on the HSIOM functional connections. The specific connection is under control of HSIOM register fields.

22.4 High Speed I/O (HSIO)

These types of I/O ports are designed for high-speed operations supporting interfaces such as QSPI/OSPI, HyperBus, SD standard, and Ethernet. Being optimized for high-speed operations, these ports do not offer slew rate control, deep-sleep operation, and analog connections.

HSIOs can be used as the standard GPIO in Active mode only. In low-power mode HSIO retains their state while the GPIO can toggle. Drive strength can be controlled using the GPIO_PRTx_CFG_OUT.DRIVE_SEL bits.

Note: Refer to the device datasheet for the availability of HSIO. Not all device support HSIO functionality.

22.5 Digital Input Buffer

The digital input buffer provides a high-impedance buffer for the external digital input. The buffer is enabled or disabled by the GPIO_PRTx_CFG.IN_ENy bit (where 'x' is the port number and 'y' is the pin number).

The input buffer is connected to the HSIOM for routing to the CPU port registers and selected peripherals. Writing to the HSIOM port select register (HSIOM_PRTx_PORT_SEL) selects the pin connection. See the device datasheet for the specific connections available for each pin. A port pin can be used as an input and output at the same time.

If a pin is only connected to an analog signal, the input buffer should be disabled to avoid crowbar currents.


Each pin's input buffer trip point and hysteresis are configurable for the following modes:

- CMOS + I²C
- TTL
- Automotive

CMOS and TTL buffer modes are selected by the GPIO_PRTx_CFG.IN.VTRIP_SELy_0 bit. To set the mode to Automotive use the GPIO_PRTx_CFG.IN.AUTOLVL.VTRIP_SELy_1 bit to enable/disable the mode.

Note: Set the GPIO_PRTx_CFG.IN mode to CMOS if enabling the Automotive mode. The trip levels of CMOS and Automotive are shown in Figure 22-3.

Figure 22-3. Input Buffer Mode's Tripping Levels



22.6 Digital Output Driver

Pins are driven by the digital output driver. It consists of circuitry to implement different drive modes and slew rate control for the digital output signals. The HSIOM selects the control source for the output driver. The two primary types of control sources are port configuration registers and fixed-function digital peripherals. A particular HSIOM connection is selected by writing to the HSIOM port select register (HSIOM_PRTx_PORT_SEL).

Each GPIO pin has ESD diodes to clamp the pin voltage to the I/O supply source. Ensure that the voltage at the pin does not exceed the I/O supply voltage $V_{DDIO}/V_{DDD}/V_{DDA}$ or drop below $V_{SSIO}/V_{SSD}/V_{SSA}$. For the absolute maximum and minimum GPIO voltage, see the device datasheet.

The digital output driver can be enabled or disabled in hardware by the output data register (GPIO_PRTx_OUT) associated with the output pin. Peripherals other than GPIO port, directly control both the output and output-enable of the output buffer.

Configuration register, GPIO_PRTx_CFG. [Table 22-1](#) lists the drive modes. Drive mode '1' is reserved and should not be used in most designs. CPU register connections support seven discrete drive modes to maximize design flexibility. Fixed-function digital peripherals, such as SCB and TCPWM blocks, support modified functionality for the same seven drive modes compatible with fixed peripheral signaling. [Figure 22-4](#) shows simplified output driver diagrams of the pin view for the CPU registers on each of the eight drive modes. [Figure 22-5](#) is a simplified output driver diagram that shows the pin view for fixed-function-based peripherals for each of the eight drive modes.

22.6.1 Drive Modes

Each I/O is individually configurable to one of eight drive modes by the DRIVE_MODE[7:0] field of the Port

Table 22-1. Drive Mode Settings

Drive Mode	Value	GPIO Port Configuration Register, AMUXBUS,				Fixed-Function Digital Peripheral			
		OUT_EN = 1		OUT_EN = 0		OUT_EN = 1		OUT_EN = 0	
		OUT = 1	OUT = 0	OUT = 1	OUT = 0	OUT = 1	OUT = 0	OUT = 1	OUT = 0
High Impedance	0	High Z	High Z	High Z	High Z	High Z	High Z	High Z	High Z
Reserved	1	Strong 1	Strong 0	High Z	High Z	Strong 1	Strong 0	Weak 1 and 0	Weak 1 and 0
Resistive Pull Up	2	Weak 1	Strong 0	High Z	High Z	Strong 1	Strong 0	Weak 1	Weak 1
Resistive Pull Down	3	Strong 1	Weak 0	High Z	High Z	Strong 1	Strong 0	Weak 0	Weak 0
Open Drain, Drives Low	4	High Z	Strong 0	High Z	High Z	Strong 1	Strong 0	High Z	High Z
Open Drain, Drives High	5	Strong 1	High Z	High Z	High Z	Strong 1	Strong 0	High Z	High Z
Strong	6	Strong 1	Strong 0	High Z	High Z	Strong 1	Strong 0	High Z	High Z
Resistive Pull Up and Down	7	Weak 1	Weak 0	High Z	High Z	Strong 1	Strong 0	Weak 1	Weak 0

Note: OUT_EN is not user configurable; its value is set according to the pin mode. For example, in GPIO mode OUT_EN = 1. See [Table 22-6](#).

Figure 22-4. GPIO Port, Drive Mode Block Diagram





Figure 22-5. Fixed-Function Peripheral I/O Drive Mode Block Diagrams



■ High-Impedance

High-impedance mode is the standard high-impedance (High-Z) state recommended for analog and digital inputs. For digital signals, the input buffer is enabled; for analog signals, the input buffer is typically disabled to reduce crowbar current and leakage in low-power designs. To achieve the lowest device current, unused GPIOs must be configured to the high-impedance drive mode with input buffer disabled. High-impedance drive mode with input buffer disabled is also the default pin reset state.

■ Resistive Pull-Up or Resistive Pull-Down

Resistive modes provide a series resistance in one of the data states and strong drive in the other. Pins can be used for either digital input or digital output in these modes. If resistive pull-up is required, a '1' must be written to that pin's Data Register bit. If resistive pull-down is required, a '0' must be written to that pin's Data Register. Interfacing mechanical switches is a common application of these drive modes. The resistive modes are also used to interface TVII-B-H with open drain drive lines. Resistive pull-up is used when the input is open drain low and resistive pull-down is used when the input is open drain high.

■ Open Drain Drives High and Open Drain Drives Low

Open drain modes provide high impedance in one of the data states and strong drive in the other. Pins are useful as digital inputs or outputs in these modes. Therefore, these modes are widely used in bi-directional digital communication. Open drain drive high mode is used when the signal is externally pulled down and open drain drive low is used when the signal is externally pulled high. A common application for the open drain drives low mode is driving I²C bus signal lines.

■ Strong Drive

The strong drive mode is the standard digital output mode for pins; it provides a strong CMOS output drive in both high and low states. Strong drive mode pins should not be used as inputs under normal circumstances. This mode is often used for digital output signals or to drive external devices.

■ Resistive Pull-Up and Resistive Pull-Down

In the resistive pull-up and pull-down mode, the GPIO will have a series resistance in both logic 1 and logic 0 output states. The high data state is pulled up while the low data state is pulled down. This mode is useful when the pin is driven by other signals that may cause shorts.

22.6.2 Slew Rate Control

Some GPIO pins have fast and slow output slew rate options for the strong drivers configured using the SLOW bit of the port output configuration register (GPIO_PRTx_CFG_OUT). By default, this bit is cleared and the port works in fast slew mode. This bit can be set if a slow slew rate is required. Slower slew rate results in reduced EMI and crosstalk and are recommended for low-frequency signals or signals without strict timing constraints.

When configured for fast slew rate, the drive strength can be set to one of four levels using the GPIO_PRTx_CFG_OUT.DRIVE_SELy. The drive strength field determines the active portion of the output drivers used and can affect the slew rate of output signals. Drive strength options are full drive strength (default), one-half strength, and one-quarter strength. Drive strength must be set to full drive strength when the slow slew rate bit (SLOW) is set.

Note: Only an enhanced I/O port will support slew rate control; for standard ports slew rate can be controlled using drive strength. Refer to the device datasheet for I/O ports with Enhanced functionality.

Table 22-2. Drive Select for GPIO_STD and GPIO_ENH

Drive Select	DRIVE_SEL[0:1]	Description
DRIVE_SEL_ZERO	00	Full drive strength: GPIO drives current at its maximum rated specification.
DRIVE_SEL_ONE	01	Full drive strength: GPIO drives current at its maximum rated specification
DRIVE_SEL_TWO	10	1/2 drive strength: GPIO drives current at one-half of its maximum rated specification.
DRIVE_SEL_THREE	11	1/4 drive strength: GPIO drives current at one-quarter of its maximum rated specification.

Table 22-3. Drive Select for HSIO_STD

Drive Select	DRIVE_SEL[0:1]	Description
DRIVE_SEL_ZERO	00	HSIO default mode
DRIVE_SEL_ONE	01	GPIO full drive strength
DRIVE_SEL_TWO	10	GPIO 1/2 drive strength
DRIVE_SEL_THREE	11	GPIO 1/4 drive strength

Table 22-4. Drive Select for HSIO_ENH

Drive Select	DS_TRIM[0:2]	Description
DEFAULT	000	Default (50 Ω)
DS_120OHM	001	120 Ω
DS_90OHM	010	90 Ω
DS_60OHM	011	60 Ω
DS_50OHM	100	50 Ω
DS_30OHM	101	30 Ω
DS_20OHM	110	20 Ω
DS_15OHM	111	15 Ω

Table 22-5. Drive Select for HSIO_STD_LN

DRIVE_SEL_EXT[0:2]	Description
000	4 mA drive strength, 1.8 V
001	3 mA drive strength, 1.8 V
010	2 mA drive strength, 1.8 V
011	1 mA drive strength, 1.8 V
100	4 mA drive strength, 1.2 V
101	3 mA drive strength, 1.2 V
110	2 mA drive strength, 1.2 V
111	1 mA drive strength, 1.2 V

22.7 High-Speed I/O Matrix

The high-speed I/O matrix (HSIOM) is a set of high-speed multiplexers that route internal CPU and peripheral signals to and from GPIOs. HSIOM allows GPIOs to be shared with multiple functions and multiplexes the pin connection to a particular peripheral selected by the user. The HSIOM_PRTx_PORT_SEL registers allow a single selection from up to 32 different connections to each pin as listed in [Table 22-6](#).

Table 22-6. HSIOM Connections

SELy_SEL	Name	Digital Driver Signal Source		Digital Input Signal Destination	Description
		OUT	OUT_EN		
0	GPIO	OUT Register	1	IN Register	GPIO_PRTx_OUT register controls "out"
1	Reserved	–	–	–	–
2	Reserved	–	–	–	–
3	Reserved	–	–	–	–
4	Reserved	–	–	–	–
5	Reserved	–	–	–	–
6	Reserved	–	–	–	–
7	Reserved	–	–	–	–
8	ACT_0	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 0 - See device datasheet for specific pin connectivity
9	ACT_1	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 1 - See device datasheet for specific pin connectivity
10	ACT_2	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 2 - See device datasheet for specific pin connectivity
11	ACT_3	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 3 - See device datasheet for specific pin connectivity
12	DS_0	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 0 - See device datasheet for specific pin connectivity
13	DS_1	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 1 - See device datasheet for specific pin connectivity
14	DS_2	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 2 - See device datasheet for specific pin connectivity
15	DS_3	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 3 - See device datasheet for specific pin connectivity
16	ACT_4	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 4 - See device datasheet for specific pin connectivity
17	ACT_5	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 5 - See device datasheet for specific pin connectivity
18	ACT_6	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 6 - See device datasheet for specific pin connectivity
19	ACT_7	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 7 - See device datasheet for specific pin connectivity
20	ACT_8	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 8 - See device datasheet for specific pin connectivity
21	ACT_9	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 9 - See device datasheet for specific pin connectivity
22	ACT_10	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 10 - See device datasheet for specific pin connectivity
23	ACT_11	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 11 - See device datasheet for specific pin connectivity

Table 22-6. HSIOM Connections

SELy_SEL	Name	Digital Driver Signal Source		Digital Input Signal Destination	Description
		OUT	OUT_EN		
24	ACT_12	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 12 - See device datasheet for specific pin connectivity
25	ACT_13	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 13 - See device datasheet for specific pin connectivity
26	ACT_14	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 14 - See device datasheet for specific pin connectivity
27	ACT_15	Active Source OUT	Active Source OUT_EN	Active Source IN	Active functionality 15 - See device datasheet for specific pin connectivity
28	DS_4	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 4 - See device datasheet for specific pin connectivity
29	DS_5	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 5 - See device datasheet for specific pin connectivity
30	DS_6	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 6 - See device datasheet for specific pin connectivity
31	DS_7	DeepSleep Source OUT	DeepSleep Source OUT_EN	DeepSleep IN	DeepSleep functionality 7 - See device datasheet for specific pin connectivity

Note. The Active and DeepSleep sources are pin dependent. See the Pinouts section of the device datasheet for more details on the features supported by each pin. If the JTAG input pin is configured to the SWJ_TRSTN mode upon reset (refer to the related device datasheet for the pin number), change the mode of the pin from SWJ_TRSTN to GPIO according to the following sequence:

1. HSIOM_PRTx_PORT_SEL = 0 (GPIO)
2. GPIO_PRTx_CFG = 0

enabled as SWD lines during power up. The DAP connection does not provide pull-up or pull-down resistors; therefore, if left floating some crowbar current is possible. The DAP connection can be disabled or reconfigured for general-purpose use through the HSIOM only after the device boots and starts executing code.

22.9 Behavior in Low-Power Modes

22.8 I/O State on Power Up

During power up, all the GPIOs are in high-impedance analog state and the input buffers are disabled. During runtime, GPIOs can be configured by writing to the associated registers. Note that the pins supporting debug access port (DAP) connections (SWD lines) are always

To allow for DeepSleep Interrupt and Hibernate wake up functionality, the GPIOs hold/freeze their configuration information when entering either DeepSleep or Hibernate power mode. As a result, the configuration signals can be routed in the Active power domain. [Table 22-7](#) shows the status of GPIOs in low-power modes.


Table 22-7. GPIO in Low-Power Modes

Low-Power Mode	Status
Sleep	■ Standard GPIO pins are active and can be driven by most peripherals such as CapSense, TCPWM, and SCB, which can operate in sleep mode. ■ Inputs buffers are active; thus an interrupt on any I/O can be used to wake the CPU.
DeepSleep	■ GPIO pins, connected to deep-sleep domain peripherals, are functional. All other pins are hold/frozen and will maintain the last output driver state and configuration. ■ Pin interrupts are functional on all I/Os and can be used to wake the device.
Hibernate	■ Pin output states and configuration are latched and remain in the hold/frozen state. ■ Pin interrupts are functional on only select I/Os and can be used to wake the device. See the device datasheet for specific Hibernate pin connectivity.

22.10 Interrupt

All port pins have the capability to generate interrupts. There are two routing possibilities for pin signals to generate interrupts, as shown in [Figure 22-6](#).

Figure 22-6. Interrupt Signal Routing




Pin signal through the 'GPIO Edge Detect' block with direct connection to the CPU interrupt controller. Interrupt generation is independent of HSIOM configuration, consider disabling it (GPIO_PRTx_INTR_CFG.EDGEy_SEL = 00) to avoid unwanted GPIO interrupt. [Figure 22-7](#) shows the block diagram of the GPIO Edge Detect block.

Each GPIO pin interrupt can be activated either on GPIO pin rising or GPIO pin falling edge changes based on GPIO_PRTx_INTR_CFG register selection. Each of the up to eight GPIO pads in an I/O port has interrupt cause fields that are set to '1' on a GPIO pin input signal rising and falling edge respectively (GPIO_PRTx_INTR register). The GPIO_PRTx_INTR_MASK register of the I/O port specifies which interrupt is propagated to its interrupt output. This propagated GPIO pin interrupt is a DeepSleep functionality

interrupt and this allows for a GPIO input signal change to wake up the CPU from DeepSleep power mode after combined with other given GPIO port "i" pin interrupts as shown in [Figure 22-7](#) to form a port interrupt interrupts_gpio[i].

In addition, an register field specifies one specific I/O input signal that is routed to a 50-ns glitch filter (GPIO_PRTx_INTR_CFG.FLT_SEL register). The glitch filter output has a dedicated detection circuitry and has dedicated detection control fields (like GPIO_PRTx_INTR.FLT_EDGE). Each I/O port has a dedicated interrupt associated to it (interrupts_gpio[i] for I/O port i). [Figure 22-6](#) illustrates the interrupt functionality. [Figure 22-7](#) shows the GPIO Edge Detect block architecture.

Figure 22-7. GPIO Edge Detect Block Architecture



The software ISR can read the 8+1 interrupt cause fields to determine the I/O or glitch filter signal(s) that caused the interrupt activation. The ISR needs to clear the interrupt cause fields to deactivate the interrupt.

An edge detector is present at each pin. It is capable of detecting rising edge, falling edge, and both edges without any reconfiguration. The edge detector is configured by

writing into the GPIO_PRTx_INTR_CFG.EDGEy_SEL field, as shown in [Table 22-8](#).

Table 22-8. Edge Detector Configuration

EDGE_SEL	Configuration
00	Interrupt is disabled
01	Interrupt on Rising Edge
10	Interrupt on Falling Edge
11	Interrupt on Both Edges

Writing '1' to the corresponding status bit clears the pin edge state. It is important to clear the edge state status bit; otherwise, an interrupt can occur repeatedly for a single trigger or respond only once for multiple triggers, which is explained later in this section. When the Port Interrupt Control Status register is read at the same time an edge is occurring on the corresponding port, it can result in the edge not being properly detected. Therefore, when using GPIO interrupts, it is recommended to read the status register only inside the corresponding interrupt service routine and not in any other part of the code.

Firmware and the debug interface are able to trigger a hardware interrupt from any pin by setting the corresponding bit in the GPIO_PRTx_INTR_SET register.

In addition to the pins, each port provides a glitch filter connected to its own edge detector. This filter can be driven by one of the pins of a port. The selection of the driving pin is done by writing to the GPIO_PRTx_INTR_CFG.FLT_SEL field as shown in [Table 22-9](#).

Table 22-9. Glitch Filter Input Selection

FLT_SEL	Selected Pin
000	Pin 0 is selected
001	Pin 1 is selected
010	Pin 2 is selected
011	Pin 3 is selected
100	Pin 4 is selected
101	Pin 5 is selected
110	Pin 6 is selected
111	Pin 7 is selected

When a port pin edge occurs, it is required to know which pin caused the edge. This is done by reading the Port Interrupt Status register, GPIO_PRTx_INTR. This register includes both the latched information on which pin detected an edge and the current pin status. This allows the CPU to read both information in a single read operation. This register has an additional use - to clear the latched edge state.

The GPIO_PRTx_INTR_MASK register enables forwarding of the GPIO_PRTx_INTR edge detect signal to the interrupt controller when a '1' is written to a pin's corresponding bitfield. The GPIO_PRTx_INTR_MASKED register can then be read to determine the specific pin that generated the interrupt signal forwarded to the interrupt controller. The masked edge detector outputs of a port are then ORed together and routed to the interrupt controller (NVIC in the CPU subsystem). Thus, there is only one interrupt vector per port.

The masked and ORed edge detector block output is routed to the Interrupt Source Multiplexer (see the [Interrupts chapter on page 144](#) for details), which gives an option of Level and Rising Edge detection. If the Level

option is selected, an interrupt is triggered repeatedly as long as the Port Interrupt Status register bit is set. If the Rising Edge detect option is selected, an interrupt is triggered only once if the Port Interrupt Status register is not cleared. Thus, it is important to clear the interrupt status bit if the Edge Detect block is used.

There is a dedicated interrupt vector for each port when the interrupt signal is routed through the fixed-function route.

All the port interrupt vectors are also ORed together into a single interrupt vector for use on devices with more ports than there are interrupt vectors available. To determine the port that triggered the interrupt, the GPIO_INTR_CAUSEx registers can be read. A '1' present in a bit location indicates that the corresponding port has a pending interrupt. The indicated GPIO_PRTx_INTR register can then be read to determine the pin source.

The GPIO_VDD_ACTIVE register provides the capability to read the state of the external power supplies. It indicates the absence or presence of VDDIO supplies, VDDA and VDDD. In addition, power supply interrupts can be configured to generate interrupts on supply state change. The GPIO_VDD_INTR_MASK register is used to mask/enable the forwarding of interrupt to the CPUs. The status of interrupt can be checked with the GPIO_VDD_INTR register.

22.11 Peripheral Connections

22.11.1 Firmware-Controlled GPIO

For standard firmware-controlled GPIO using registers, the GPIO mode must be selected in the HSIOM_PRTx_PORT_SEL register.

The GPIO_PRTx_OUT register is used to read and write the output buffer state for GPIOs. A write operation to this register changes the GPIO's output driver state to the written value. A read operation reflects the output data written to this register and the resulting output driver state. It does not return the current logic level present on GPIO pins, which may be different. Using the GPIO_PRTx_OUT register, read-modify-write sequences can be safely performed on a port that has both input and output GPIOs.

In addition to the data register, three other registers – GPIO_PRTx_OUT_SET, GPIO_PRTx_OUT_CLR, and GPIO_PRTx_OUT_INV – are provided to set, clear, and invert the output data respectively on specific pins in a port without affecting other pins. This avoids the need for read-modify-write operations in most use cases. Writing '1' to these register bitfields will set, clear, or invert the respective pin; writing '0' will have no effect on the pin state.

GPIO_PRTx_IN is the port I/O pad register that provides the actual logic level present on the GPIO pin when read. Writes to this register have no effect.

22.11.2 Analog I/O

Analog resources, such as SAR ADC, which require low-impedance routing paths have dedicated pins. Dedicated analog pins provide direct connections to specific analog blocks. They help improve performance and should be given priority over other pins when using these analog resources. See the device datasheet for details on these dedicated pins of TVII-B-H.

To configure a GPIO as a dedicated analog I/O, it should be configured in high-impedance analog mode (see [Table 22-1](#)) with input buffer disabled and the respective connection should be enabled via registers in the specific analog resource.

While it is preferred that analog pins disable the input buffer, it is acceptable to enable the input buffer if simultaneous analog and digital input features are required.

22.11.3 Serial Communication Block (SCB)

SCB can be configured for UART, I²C, and SPI communication protocols. The SCB has dedicated connections to pins through the HSIOM. See the device datasheet for details on the dedicated pin connections. When the SCB I²C, UART, or SPI modes are used, the SCB controls the interface pins digital output state and output enable. In principle, all the peripherals with output connected to HSIOM also control the output enable for the pin. For details on the recommended interface pin drive

modes and controlling signals, refer to the [Serial Communications Block \(SCB\) chapter on page 274](#).

22.12 Smart I/O


The Smart I/O block adds programmable logic to an I/O port. This programmable logic integrates board-level Boolean logic functionality such as AND, OR, and XOR into the port. The Smart I/O block has these features:

- Integrate board-level Boolean logic functionality into a port
- Ability to pre-process HSIOM input signals from the GPIO port pins
- Ability to post-process HSIOM output signals to the GPIO port pins
- Support in all device power modes
- Integrate closely to the I/O pads, providing shortest signal paths with programmability

22.12.1 Overview

The Smart I/O block is positioned in the signal path between the HSIOM and the I/O port. The HSIOM multiplexes the output signals from fixed-function peripherals and CPU to a specific port pin and vice-versa. The Smart I/O block is placed on this signal path, acting as a bridge that can process signals between port pins and HSIOM, as shown in [Figure 22-8](#).

Figure 22-8. Smart I/O Interface



The signal paths supported through the Smart I/O block as shown in [Figure 22-8](#) are as follows:

1. Implement self-contained logic functions that directly operate on port I/O signals
2. Implement self-contained logic functions that operate on HSIOM signals
3. Operate on and modify HSIOM output signals and route the modified signals to port I/O signals
4. Operate on and modify port I/O signals and route the modified signals to HSIOM input signals

The following sections discuss the Smart I/O block components, routing, and configuration in detail. In these sections, the GPIO signal (io_data_in) refers to the input signal from the I/O port; device or chip (chip_data) signals refer to the output signal from HSIOM. smartio_data is the output of Smart I/O interface depending on the configuration of the blocks,

22.12.2 Block Components

The internal logic of the Smart I/O includes these components:

- Clock/reset
- Synchronizers
- Three-input lookup table (LUT3)
- Data unit

22.12.2.1 Clock and Reset

The clock and reset component selects the Smart I/O block's clock (`clk_block`) and reset signal (`rst_block_n`). A single clock and reset signal is used for all components in the block. The clock and reset sources are determined by the `SMARTIO_PRTx_CTL.CLOCK_SRC` field. The selected clock is used for the synchronous logic in the block components, which includes the I/O input synchronizers, LUT, and data unit components. The selected reset is used to asynchronously reset the synchronous logic in the LUT and data unit components.

Note that the selected clock (`clk_block`) for the block's synchronous logic is not phase-aligned with other synchronous logic in the device, operating on the same clock. Therefore, communication between Smart I/O and other synchronous logic should be treated as asynchronous.

The following clock sources are available for selection:

- GPIO input signals `io_data_in[7:0]`. These clock sources have no associated reset.
- HSIOM output signals `chip_data[7:0]`. These clock sources have no associated reset.
- The Smart I/O clock (`PCLK_SMARTIOx_CLOCK`). This is derived from the system clock (`CLK_SYS/CLK_HF`) using a peripheral clock divider. See the [Clocking System chapter on page 197](#) for details on peripheral clock dividers. This clock is available only in Active and Sleep power modes. The clock can have one out of two associated resets: `rst_sys_act_n` and `rst_sys_dpslp_n`. These resets determine in which system power modes the block synchronous state is reset; for example, `rst_sys_act_n` is intended for Smart I/O synchronous functionality in the Active power mode and reset is activated in the DeepSleep power mode.
- The low-frequency (40 kHz) system clock (`clk_lf`). This clock is available in DeepSleep power mode. This clock has an associated reset, `rst_lf_dpslp_n`.

When the block is enabled, the selected clock (`clk_block`) and associated reset (`rst_block_n`) are provided to the internal logic components. When the block is disabled, no clock is released to the internal logic components and the reset is activated (the LUT and data unit components are set to the reset value of '0').

The I/O input synchronizers introduce a delay of two `clk_block` cycles (when synchronizers are enabled). As a result, in the first two cycles, the block may be exposed to stale data from the synchronizer output. Hence, during the first two clock cycles, the reset is activated and the block is in bypass mode.


Table 22-10. Clock and Reset Register Control

Register[BIT_POS]	Bit Name	Description
SMARTIO_PRTx- _CTL[12:8]	CLK_SRC[4:0]	Clock (<code>clk_block</code>)/reset (<code>rst_block_n</code>) source selection: 0: <code>io_data_in[0]</code>'1' ... 7: <code>io_data_in[7]</code>'1' 8: <code>chip_data[0]</code>'1' ... 15: <code>chip_data[7]</code>'1' 16: <code>clk_smartio/rst_sys_act_n</code>; asserts reset in any power mode other than Active; that is, Smart I/O is active only in Active power mode with clock from the peripheral divider. 17: <code>clk_smartio/rst_sys_dpslp_n</code>. Smart I/O is active in all power modes with clock from the peripheral divider. However, the clock will not be active in DeepSleep power mode. 19: <code>clk_lf/rst_lf_dpslp_n</code>. Smart I/O is active in all power modes with clock from ILO. 20-30: Clock source is a constant '0'. 31: <code>CLK_SYS</code>'1'. This selection is not intended for <code>clk_sys</code> operation.

22.12.2.2 Synchronizer

Each GPIO input signal and device input signal (HSIOM input) can be used either asynchronously or synchronously. To use the signals synchronously, a double flip-flop synchronizer, as shown in [Figure 22-9](#), is placed on both these signal paths to synchronize the signal to the Smart I/O clock (clk_block). The synchronization for each pin/input is enabled or disabled by setting or clearing the IO_SYNC_EN[i] bit field for GPIO input signal and CHIP_SYNC_EN[i] for HSIOM signal in the SMARTIO_PRT0_SYNC_CTL register, where 'i' is the pin number.

Figure 22-9. Smart I/O Clock Synchronizer



22.12.2.3 LUT3

Each Smart I/O block contains eight lookup table (LUT3) components. The LUT3 component consists of a three-input LUT and a flip-flop. Each LUT3 block takes three input signals and generates an output based on the configuration set in the SMARTIO_PRTx_LUT_CTLy register ('y' denotes the LUT3 number). For each LUT3, the configuration is determined by an 8-bit lookup vector LUT[7:0] and a 2-bit opcode OPC[1:0] in the SMARTIO_PRTx_LUT_CTLy register. The 8-bit vector is used as a lookup table for the three input signals. The 2-bit opcode determines the usage of the flip-flop. The LUT3 configuration for different opcode is shown in [Figure 22-10](#).

SMARTIO_PRTx_LUT_SELy registers select the three input signals (tr0_in, tr1_in, and tr2_in) going into each LUT3. The input can come from the following sources:


- Data unit output
- Other LUT3 output signals (tr_out)
- HSIOM output signals (chip_data[7:0])
- GPIO input signals (io_data_in[7:0])

SMARTIO_PRTx_LUT_SELy.LUT_TR0_SEL register selects the tr0_in signal for the yth LUT3. Similarly, SMARTIO_PRTx_LUT_SELy.LUT_TR1_SEL bits and SMARTIO_PRTx_LUT_SELy.LUT_TR2_SEL bits select the tr1_in and tr2_in signals respectively. See [Table 22-11](#) for details.

Table 22-11. LUT3 Register Control

Register[BIT_POS]	Bit Name	Description
SMARTIO_PRTx_LUT_CTLy[7:0]	LUT[7:0]	LUT configuration. Depending on the LUT opcode (LUT_OPC), internal state, and LUT input signals tr0_in, tr1_in, and tr2_in, the LUT configuration is used to determine the LUT output signal and the next sequential state.
SMARTIO_PRTx_LUT_CTLy[9:8]	LUT_OPC[1:0]	LUT opcode specifies the LUT operation as illustrated in Figure 22-10 .
SMARTIO_PRTx_LUT_SELy[3:0]	LUT_TR0_SEL[3:0]	LUT input signal tr0_in source selection: 0: Data unit output 1: LUT 1 output 2: LUT 2 output 3: LUT 3 output 4: LUT 4 output 5: LUT 5 output 6: LUT 6 output 7: LUT 7 output 8: chip_data[0] (for LUTs 0, 1, 2, 3); chip_data[4] (for LUTs 4, 5, 6, 7) 9: chip_data[1] (for LUTs 0, 1, 2, 3); chip_data[5] (for LUTs 4, 5, 6, 7) 10: chip_data[2] (for LUTs 0, 1, 2, 3); chip_data[6] (for LUTs 4, 5, 6, 7) 11: chip_data[3] (for LUTs 0, 1, 2, 3); chip_data[7] (for LUTs 4, 5, 6, 7) 12: io_data_in[0] (for LUTs 0, 1, 2, 3); io_data_in[4] (for LUTs 4, 5, 6, 7) 13: io_data_in[1] (for LUTs 0, 1, 2, 3); io_data_in[5] (for LUTs 4, 5, 6, 7) 14: io_data_in[2] (for LUTs 0, 1, 2, 3); io_data_in[6] (for LUTs 4, 5, 6, 7) 15: io_data_in[3] (for LUTs 0, 1, 2, 3); io_data_in[7] (for LUTs 4, 5, 6, 7)
SMARTIO_PRTx_LUT_SELy[11:8]	LUT_TR1_SEL[3:0]	LUT input signal tr1_in source selection: 0: LUT 0 output 1: LUT 1 output 2: LUT 2 output 3: LUT 3 output 4: LUT 4 output 5: LUT 5 output 6: LUT 6 output 7: LUT 7 output 8: chip_data[0] (for LUTs 0, 1, 2, 3); chip_data[4] (for LUTs 4, 5, 6, 7) 9: chip_data[1] (for LUTs 0, 1, 2, 3); chip_data[5] (for LUTs 4, 5, 6, 7) 10: chip_data[2] (for LUTs 0, 1, 2, 3); chip_data[6] (for LUTs 4, 5, 6, 7) 11: chip_data[3] (for LUTs 0, 1, 2, 3); chip_data[7] (for LUTs 4, 5, 6, 7) 12: io_data_in[0] (for LUTs 0, 1, 2, 3); io_data_in[4] (for LUTs 4, 5, 6, 7) 13: io_data_in[1] (for LUTs 0, 1, 2, 3); io_data_in[5] (for LUTs 4, 5, 6, 7) 14: io_data_in[2] (for LUTs 0, 1, 2, 3); io_data_in[6] (for LUTs 4, 5, 6, 7) 15: io_data_in[3] (for LUTs 0, 1, 2, 3); io_data_in[7] (for LUTs 4, 5, 6, 7)
SMARTIO_PRTx_LUT_SELy[19:16]	LUT_TR2_SEL[3:0]	LUT input signal tr2_in source selection. Encoding is the same as for LUT_TR1_SEL.

Figure 22-10. Smart I/O LUT3 Configuration



22.12.2.4 Data Unit

Each Smart I/O block includes a data unit (DU) component. The data unit consists of an 8-bit datapath. It is capable of performing simple increment, decrement, increment/decrement, shift, and AND/OR operations. The operation performed by the DU is selected using a 4-bit opcode SMARTIO_PRTx_DU_CTL.DU_OPC field.

The data unit component supports up to three input trigger signals (tr0_in, tr1_in, tr2_in) similar to the LUT3 component. These signals are used to initiate an operation defined by the DU opcode. In addition, the data unit also includes two 8-bit data inputs (data0_in[7:0] and data1_in[7:0]) that are used to initialize the 8-bit internal state (data[7:0]) or to provide a reference. The 8-bit data input source is configured as:

- Constant '0x00'
- io_data_in[7:0]
- chip_data[7:0]
- SMARTIO_PRTx_DATA.DATA field

The trigger signals are selected using the SMARTIO_PRTx_DU_SEL.DU_TRy_SEL field. The SMARTIO_PRTx_DU_SEL.DU_DATAy_SEL field select the 8-bit input data source. The size of the DU (number of bits used by the datapath) is defined by the SMARTIO_PRTx_DU_CTL.DU_SIZE field. See [Table 22-12](#) for register control details.

Table 22-12. Data Unit Register Control

Register[BIT_POS]	Bit Name	Description
SMARTIO_PRTx_DU_CTL[2:0]	DU_SIZE[2:0]	Size/width of the data unit (in bits) is DU_SIZE+1. For example, if DU_SIZE is 7, the width is 8 bits.
SMARTIO_PRTx_DU_CTL[11:8]	DU_OPC[3:0]	Data unit opcode specifies the data unit operation: 1: INCR 2: DECR 3: INCR_WRAP 4: DECR_WRAP 5: INCR_DECR 6: INCR_DECR_WRAP 7: ROR 8: SHR 9: AND_OR 10: SHR_MAJ3 11: SHR_EQL Otherwise: Undefined.
SMARTIO_PRTx_DU_SEL[3:0]	DU_TR0_SEL[3:0]	Data unit input signal tr0_in source selection: 0: Constant '0'. 1: Constant '1'. 2: Data unit output. 10-3: LUT 7 - 0 outputs. Otherwise: Undefined.
SMARTIO_PRTx_DU_SEL[11:8]	DU_TR1_SEL[3:0]	Data unit input signal tr1_in source selection. Encoding same as DU_TR0_SEL
SMARTIO_PRTx_DU_SEL[19:16]	DU_TR2_SEL[3:0]	Data unit input signal tr2_in source selection. Encoding same as DU_TR0_SEL
SMARTIO_PRTx_DU_SEL[25:24]	DU_DATA0_SEL[1:0]	Data unit input data data0_in source selection: 0: 0x00 1: chip_data[7:0]. 2: io_data_in[7:0]. 3: SMARTIO_PRTx_DATA.DATA[7:0] register field.
SMARTIO_PRTx_DU_SEL[29:28]	DU_DATA1_SEL[1:0]	Data unit input data data1_in source selection. Encoding same as DU_DATA0_SEL.
SMARTIO_PRTx_DATA[7:0]	DATA[7:0]	Data unit input data source.

The data unit generates a single output trigger signal (tr_out). The internal state (du_data[7:0]) is captured in flip-flops and requires clk_block.

The following pseudo code describes the various datapath operations supported by the DU opcode. Note that “Comb” describes the combinatorial functionality – that is, functionalities that operate independent of previous output states. “Reg” describes the registered functionality – that is, functionalities that operate on inputs and previous output states (registered using flip-flops).

```
// The following is shared by all operations.

data_eql_data1_in = (data & mask) == (data1_in & mask));
data_eql_0        = (data & mask) == 0);
data_incr         = (data + 1) & mask;
data_decr         = (data - 1) & mask;
data0_masked      = data_in0 & mask;

// INCR operation:
Comb: tr_out = data_eql_data1_in;
Reg:  data ≤ data;
      if (tr0_in)      data ≤ data0_masked;
      else if (tr1_in) data ≤ data_eql_data1_in ? data : data_incr;

// INCR_WRAP operation:
Comb: tr_out = data_eql_data1_in;
Reg:  data ≤ data;
      if (tr0_in)      data ≤ data0_masked;
      else if (tr1_in) data ≤ data_eql_data1_in ? data0_masked : data_incr;

// DECR operation:
Comb: tr_out = data_eql_0;
Reg:  data ≤ data;
      if (tr0_in)      data ≤ data0_masked;
      else if (tr1_in) data ≤ data_eql_0          ? data : data_decr;

// DECR_WRAP operation:
Comb: tr_out = data_eql_0;
Reg:  data ≤ data;
      if (tr0_in)      data ≤ data0_masked;
      else if (tr1_in) data ≤ data_eql_0          ? data0_masked : data_decr;

// INCR_DECR operation:
Comb: tr_out = data_eql_data1_in | data_eql_0;
Reg:  data ≤ data;
      if (tr0_in)      data ≤ data0_masked;
      else if (tr1_in) data ≤ data_eql_data1_in ? data : data_incr;
      else if (tr2_in) data ≤ data_eql_0          ? data : data_decr;

// INCR_DECR_WRAP operation:
Comb: tr_out = data_eql_data1_in | data_eql_0;
Reg:  data ≤ data;
      if (tr0_in)      data ≤ data0_masked;
      else if (tr1_in) data ≤ data_eql_data1_in ? data0_masked : data_incr;
      else if (tr2_in) data ≤ data_eql_0          ? data0_masked : data_decr;

// ROR operation:
Comb: tr_out = data[0];
Reg:  data ≤ data;
      if (tr0_in)      data ≤ data0_masked;
      else if (tr1_in) {
          data ≤ {0, data[7:1]} & mask;
          data[du_size] ≤ data[0];
      }
}
```

```

// SHR operation:
Comb: tr_out = data[0];
Reg:  data ≤ data;
      if (tr0_in)      data          ≤ data0_masked;
      else if (tr1_in) {
                          data          ≤ {0, data[7:1]} & mask;
                          data[du_size] ≤ tr2_in;
      }

// SHR_MAJ3 operation:
Comb: tr_out = (data == 0x03)
           | (data == 0x05)
           | (data == 0x06)
           | (data == 0x07);
Reg:  data ≤ data;
      if (tr0_in)      data          ≤ data0_masked;
      else if (tr1_in) {
                          data          ≤ {0, data[7:1]} & mask;
                          data[du_size] ≤ tr2_in;
      }

// SHR_EQL operation:
Comb: tr_out = data_eql_data1_in;
Reg:  data ≤ data;
      if (tr0_in) data          ≤ data0_masked;
      else if (tr1_in) {
                          data          ≤ {0, data[7:1]} & mask;
                          data[du_size] ≤ tr2_in;
      }

// AND_OR operation:
Comb: tr_out = | (data & data1_in & mask);
Reg:  data ≤ data;
      if (tr0_in) data ≤ data0_masked;


```

The SHR_MAJ3 operation is useful to implement a digital filter. The filter selects the majority value of three signal values.

22.12.3 Routing

The Smart I/O block includes many switches that are used to route the signals in and out of the block and also between various components present in the block. The routing switches are handled through the SMARTIO_PRTx_LUT_SELy and SMARTIO_PRTx_DU_SEL registers. Refer to the *Traveo II Body Controller High Registers TRM* for details. The Smart I/O internal routing is shown in Figure 22-11. In the figure, note that LUT7 to LUT4 operate on io_data/chip_data[7] to io_data/chip_data[4] whereas LUT3 to LUT0 operate on io_data/chip_data[3] to io_data/chip_data[0].

Figure 22-11. Smart I/O Routing



22.12.4 Operation

The Smart I/O block should be configured and operated as follows:

1. Before enabling the block, all the components and routing should be configured.
2. In addition to configuring the components and routing, some block level settings need to be configured correctly for desired operation.
 - a. Bypass control: The Smart I/O path can be bypassed for a particular GPIO signal by setting the SMARTIO_PRTx_CTL.BYPASS bit register. When bit 'i' is set in the SMARTIO_PRTx_CTL.BYPASS bit field, the ith GPIO signal is bypassed to the HSIOM signal path directly - Smart I/O logic will not be present in that signal path. This is useful when the Smart I/O functionality is required only on select I/Os.
 - b. Pipelined trigger mode: The LUT3 input multiplexers and the LUT3 component itself do not include any combinatorial loops. Similarly, the data unit also does not include any combinatorial loops. However, when one LUT3 interacts with the other or to the data unit, inadvertent combinatorial loops are possible. To overcome this limitation, the SMARTIO_PRTx_CTL.PIPELINE_EN bit is used. When set, all the outputs (LUT3 and data unit) are registered before branching out to other components.
3. After the Smart I/O block is configured for the desired functionality, the block can be enabled by setting the SMARTIO_PRTx_CTL.ENABLE bit. If disabled, the Smart I/O block is put in bypass mode, where the GPIO signals are directly controlled by the HSIOM signals and vice-versa. The Smart I/O block must be configured; that is, all register settings must be updated before enabling the block to prevent glitches during register updates.

Table 22-13. Smart I/O Block Controls

Register [BIT_POS]	Bit Name	Description
SMARTIO_PRTx_CTL[25]	PIPELINE_EN	Enable for pipeline register: 0: Disabled (register is bypassed). 1: Enabled
SMARTIO_PRTx_CTL[31]	ENABLED	Enable Smart I/O. Should only be set to '1' when the Smart I/O is completely configured: 0: Disabled (signals are bypassed; behavior as if BYPASS[7:0] is 0xFF). When disabled, the block (data unit and LUTs) reset is activated. If the block is disabled: - The PIPELINE_EN register field should be set to '1', to ensure low power consumption. - The CLOCK_SRC register field should be set to 20 to 30 (clock is constant '0'), to ensure low power consumption. 1: Enabled. When enabled, it takes three clk_block clock cycles until the block reset is deactivated and the block becomes fully functional. This action ensures that the I/O pins' input synchronizer states are flushed when the block is fully functional.
SMARTIO_PRTx_CTL[7:0]	BYPASS[7:0]	Bypass of the Smart I/O, one bit for each I/O pin: BYPASS[i] is for I/O pin i. When ENABLED is '1', this field is used. When ENABLED is '0', this field is not used and Smart I/O is always bypassed. 0: No bypass (Smart I/O is present in the signal path) 1: Bypass (Smart I/O is absent in the signal path)

22.12.5 Example Application

Smart I/O can be useful when the application involves simple logic operations and routing of the signal coming from or going to the I/O pin. No CPU is required for these operations. Some applications of Smart I/O are:

- Change routing to/from pins (within the port)
- Invert the polarity of signal
- Clock or signal buffer
- Detect a pattern on pins

The following are detailed implementation of a few examples to better understand the Smart I/O operation.

Example 1

Change routing to and from pins (within the port)

Consider an example where HSIOM is sending data to Pin 1 at PORT 0 and you want to route it to a different pin (Pin 7, for example) at PORT0.

Figure 22-12 shows how the Smart I/O routing will look after configuration. Table 22-14 and Table 22-15 show LUT1 and LUT7 after configuration through the SMARTIO_PRT0_LUT_CTLx register.

Figure 22-12. Smart I/O Routing Connections after Configuration




Table 22-14. Lookup Table LUT1

tr0_in	tr1_in	tr2_in	Out
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	0
1	0	0	0
1	0	1	0
1	1	0	0
1	1	1	1

Table 22-15. Lookup Table LUT7

tr0_in	tr1_in	tr2_in	Out
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	0
1	0	0	0
1	0	1	0
1	1	0	0
1	1	1	1

Figure 22-13. Register Settings for Example 1

SMARTIO_PRT0_LUT_CTL1:

LUT = 0x80
 LUT_OPC= 0x0 (Combinatorial)
 LUT_TR0_SEL= 0x9
 LUT_TR1_SEL= 0x9
 LUT_TR2_SEL= 0x9

SMARTIO_PRT0_LUT_CTL7:

LUT = 0x80
 LUT_OPC= 0x0 (Combinatorial)
 LUT_TR0_SEL= 0x1
 LUT_TR1_SEL= 0x1
 LUT_TR2_SEL= 0x1

SMARTIO_PRT0_CTL:

CLK_SRC=0x10
 BYPASS =0x7F
 ENABLED=0x1 (Enable after all configuration is done)

Example 2. Breathing LED using a constant PWM signal.

Consider a TCPWM that is sourced by a 1-MHz clock and has a period of 65535 with a compare value of 32768. This generates a 50-percent duty cycle square wave with a period of approximately 65.5 ms. The CPU is used only to initialize the TCPWM.

Assume that the Smart I/O is clocked at 30 Hz using the divided clock sourced from CLK_HF and implements several logic functions using the LUTs. [Figure 22-14](#) shows

the Smart I/O LUT configuration. The Smart I/O can implement a divide-by-two circuit from the 30-Hz clock. Therefore, LUT2 will produce a signal with a period of approximately 66.6 ms. The signal is then XORed with the PWM output (coming for Port 1 Pin 4) using LUT4 to generate a signal whose duty cycle gradually increases and decreases over time as shown in [Figure 22-14](#) and inverted by LUT6. The example also shows how only one signal is routed to two output pins. [Figure 22-15](#) shows the register settings required for the configuration given here.

Figure 22-14. LUT Configuration and Timing Diagram




Table 22-16. LUT2

tr0_in	tr1_in	tr2_in	Out
0	0	0	1
0	0	1	0
0	1	0	0
0	1	1	0
1	0	0	0
1	0	1	0
1	1	0	0
1	1	1	0

Table 22-17. LUT4

tr0_in	tr1_in	tr2_in	Out
0	0	0	0
0	0	1	1
0	1	0	0
0	1	1	0
1	0	0	0
1	0	1	0
1	1	0	1
1	1	1	0

Table 22-18. LUT6

tr0_in	tr1_in	tr2_in	Out
0	0	0	1
0	0	1	0
0	1	0	0
0	1	1	0
1	0	0	0
1	0	1	0
1	1	0	0
1	1	1	0

Figure 22-15. Register Settings for Example 2

SMARTIO_PRT1_LUT_CTL2:

LUT = 0x01
 LUT_OPC= 0x2 (Gated Output)
 LUT_TR0_SEL= 0x2
 LUT_TR1_SEL= 0x2
 LUT_TR2_SEL= 0x2

SMARTIO_PRT1_LUT_CTL4:

LUT = 0x42
 LUT_OPC= 0x0 (Combinatorial)
 LUT_TR0_SEL= 0x2
 LUT_TR1_SEL= 0x2
 LUT_TR2_SEL= 0x8

SMARTIO_PRT1_LUT_CTL6:

LUT = 0x01
 LUT_OPC= 0x0 (Combinatorial)
 LUT_TR0_SEL= 0x4
 LUT_TR1_SEL= 0x4
 LUT_TR2_SEL= 0x4

SMARTIO_PRT1_CTL:

CLK_SRC=0x10
 BYPASS =0xAF
 ENABLED=0x1 (Enable after all configuration is done)

22.13 Registers

Table 22-19. I/O Registers

Register	Name	Description
GPIO_PRTx_OUT	Port output register	Port output data register reads and writes the output driver data for I/O pins in the port.
GPIO_PRTx_OUT_CLR	Port output clear register	Port output data clear register clears output data of specific I/O pins in the port.
GPIO_PRTx_OUT_SET	Port output set register	Port output data set register sets output data of specific I/O pins in the port.
GPIO_PRTx_OUT_INV	Port output invert register	Port output data invert register inverts output data of specific I/O pins in the port.
GPIO_PRTx_IN	Port input register	Port input state register reads the current pin state present on I/O pin inputs.
GPIO_PRTx_INTR	Port interrupt status register	Port interrupt status register reads the current pin interrupt state.
GPIO_PRTx_INTR_MASK	Port interrupt mask register	Port interrupt mask register configures the mask that forwards pin interrupts to the CPU's interrupt controller. This register only masks forwarding of interrupts to the CPU's interrupt controller; it does not enable/disable the logging of interrupts into the INTR register.
GPIO_PRTx_INTR_MASKED	Port interrupt masked status register	This register contains the AND-ed values of INTR and INTR_MASK registers forwarded to the CPU interrupt controller.
GPIO_PRTx_INTR_SET	Port interrupt set register	Port interrupt set register allows firmware to set pin interrupts.
GPIO_PRTx_INTR_CFG	Port interrupt configuration register	Port interrupt configuration register selects the edge detection type for each pin interrupt.
GPIO_PRTx_CFG	Port configuration register	Port configuration register selects the drive mode and input buffer enable for each pin.
GPIO_PRTx_CFG_IN	Port input configuration register	Port input buffer configuration register configures the input buffer mode (CMOS or TTL) for each pin. VTRIP_SEL[7:0]_0.
GPIO_PRTx_CFG_IN_AUTOLVL	Port Input configuration register	Configures the input buffer upper bit i.e. VTRIP_SEL for each pin. Lower bit is still selected by CFG_IN.VTRIP_SEL[7:0]_1 field.
GPIO_PRTx_CFG_OUT	Port output configuration register	Port output buffer configuration register selects the output driver slew rate for each pin.
GPIO_PRTx_CFG_OUT2	Port output configuration register	Port output buffer configuration register 2 selects the output drive select trim for each I/O pin.
HSIOM_PRTx_PORT_SELy	HSIOM port select register	High-speed I/O mux (HSIOM) port selection register selects the hardware peripheral connection to I/O pins.
GPIO_INTR_CAUSEz	Interrupt port cause register	This register provides interrupt status corresponding to ports (0 + z × 32) to (31 + z × 32). "z" can be from 0 to 3.
GPIO_VDD_ACTIVE	External power supply detection register	This register provides external power supply status.
GPIO_VDD_INTR	Supply detection interrupt register	This register is set whenever a supply ramp up or ramp down is detected. Some bits may be set after system power-up, depending on power supply sequencing.
GPIO_VDD_INTR_MASK	Supply detection interrupt mask register	This register configures the supply detection interrupts for all supplies. It only masks the forwarding of interrupts to the CPUs and does not enable/disable the logging of interrupts into the VDD_INTR register.
GPIO_VDD_INTR_MASKED	Supply detection interrupt masked register	This register contains the AND-ed values of VDD_INTR and VDD_INTR_MASK registers.
GPIO_VDD_INTR_SET	Supply detection interrupt set register	This register allows firmware or debugger to set interrupt bits in the VDD_INTR register by writing a '1' to the corresponding bit field. When read, it returns the same value as the VDD_INTR register.
SMARTIO_PRTx_CTL	SMARTIO control register	This is the control register for SMARTIO on the specific port. It controls Enable, Clock Source, Bypass, and so on

Table 22-19. I/O Registers

Register	Name	Description
SMARTIO_PRTx_SYNC_CTL	Synchronization control register	SMARTIO synchronization control
SMARTIO_PRTx_LUT_SELy	LUT component input selection register	LUT input selection register. LUT_TR0_SEL, LUT_TR1_SEL, and LUT_TR2_SEL
SMARTIO_PRTx_LUT_CTLy	LUT component control register	The LUT control register provides opcode for LUT
SMARTIO_PRTx_DU_SEL	Data unit component input selection register	Data unit input selection register
SMARTIO_PRTx_DU_CTL	Data unit component control register	Data unit control register.
SMARTIO_PRTx_DATA	Data register	Data unit input data source.

Note: The 'x' in GPIO_PRTx/HSIOM_PRTx/SMARTIO_PRTx denotes the port number. For example, GPIO_PTR1_OUT is the Port 1 output data register. 'y'/'z' can be 0 or 1.


Section E: Digital Subsystem



This section encompasses the following chapters:

- [Serial Communications Block \(SCB\) chapter on page 274](#)
- [CAN FD Controller chapter on page 328](#)
- [Timer, Counter, and PWM chapter on page 389](#)
- [Local Interconnect Network \(LIN\) chapter on page 455](#)
- [Cryptography Block chapter on page 480](#)
- [Event Generator \(EVTGEN\) chapter on page 483](#)
- [Trigger Multiplexer chapter on page 493](#)
- [FlexRay Controller chapter on page 499](#)
- [Ethernet MAC chapter on page 545](#)
- [Serial Memory Interface chapter on page 581](#)
- [SDHC Host Controller chapter on page 626](#)
- [Audio Subsystem chapter on page 639](#)

Top Level Architecture



23. Serial Communications Block (SCB)



The Serial Communications Block (SCB) supports three serial communication protocols: serial peripheral interface (SPI), universal asynchronous receiver transmitter (UART), and inter-integrated circuit (I²C or IIC). Only one of the protocols is supported by an SCB at any given time. Traveo II MCUs have several SCBs. One of them supports only I²C slave mode and SPI slave mode. This is the only SCB that is available in the DeepSleep power mode.


23.1 Features

The SCB supports the following features:

- Standard SPI master and slave functionality with Motorola, Texas Instruments, and National Semiconductor protocols
- Standard UART functionality with SmartCard reader, local interconnect network (LIN), and IrDA protocols
 - Standard LIN slave functionality with LIN v1.3 and LIN v2.1/2.2 specification complianceThe SCB has only standard LIN slave functionality.
- Standard I²C master and slave functionality
- EZ mode for SPI and I²C slaves; allows operation without CPU intervention
- CMD_RESP mode for SPI and I²C slaves; allows operation without CPU intervention and is available only on DeepSleep-capable SCB
- Low-power (DeepSleep) mode of operation for SPI and I²C slaves (using external clocking), available only on DeepSleep-capable SCB
- DeepSleep wakeup on I²C slave address match or SPI slave selection; available only on DeepSleep-capable SCB
- Trigger outputs for connection to DMA
- Multiple interrupt sources to indicate status of FIFOs and transfers
- Local loop-back control

23.2 Block Diagram

Figure 23-1. SCB Block Diagram



23.2.1 AHB-Lite Bus Interface

The SCB block is connected to the bus infrastructure through an AHB-Lite interface. This interface provides bus masters (such as the CPU) with access to the SCB block's registers. The registers control the block's operation and provide status information. The register map provides the details on the registers and register fields.

The AHB-Lite interface handles all accesses to the SCB block's 64-Kbyte memory aperture.

- The AHB-Lite interface generates an AHB-Lite bus error for non 32-bit accesses.
- The AHB-Lite interface generates an AHB-Lite bus error for non-aligned 32-bit accesses.
- Read accesses to memory aperture locations that are not populated by a register, return '0'.
- Write accesses to memory aperture locations that are not populated by a register, are ignored (no AHB-Lite bus error is generated).

23.2.2 Trigger Interface

23.2.2.1 DMA/DW Trigger Signals

The trigger interface provides status information on the TX FIFO (**tr_tx_req**) and RX FIFO (**tr_rx_req**):

- **tr_tx_req** indicates that the TX FIFO can accept a data element (to be transmitted), controlled by **SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL**.
- **tr_rx_req** indicates that the RX FIFO can provide a (received) data element, controlled by **SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL**.

These two signals are level-sensitive.

The trigger interface is typically connected (directly or indirectly) to a DW/DMA controller.

For a RX FIFO read case, it takes "2 CLK_AHB, 1 CLK_PERI" cycles, from AHB read RX FIFO to **tr_rx_req** being cleared. If **CLK_AHB = CLK_PERI**, it takes three CLK_PERI cycles.

For TX FIFO write case, it takes “3 CLK_AHB, 1 CLK_PERI” cycles, from AHB write TX FIFO to tr_rx_req being cleared. If CLK_AHB = CLK_PERI, it takes four CLK_PERI cycles.

23.2.2.2 tr_i2c_scl_filtered Signal

The tr_i2c_scl_filtered signal is added for timeout detection on the I²C SCL line. It is connected to SCL input analog filter output, so glitches are removed.

Along with the TCPWM block, it can be used for SMBus timeout, which is required as per the SMBus specification. It can also be used to detect SCL stretching.

23.2.3 Serial Protocol Interfaces

These are the SPI, UART, and I²C signal interfaces.

The interface signals connect to the High-speed I/O Matrix (HSIOM) in the I/O Subsystem (IOSS). The HSIOM multiplexes between on-chip signals and I/O pads. The multiplexing flexibility is chip-specific and is controlled by the IOSS/HSIOM registers. The HSIOM can expose a single serial interface at different pad locations to provide system-level flexibility.

If SMARTIO is available, it can be configured to short UART_TX and UART_RX to support single-line half-duplex UART, or short MOSI and MISO to support single-line half-duplex SPI.

23.2.4 Clock and Reset Interface

The SCB block receives the following clock and reset related signals:

- A high-frequency clock, CLK_PERI. This clock is used to derive an AHB-Lite interface clock (CLK_SYS) and an SCB functionality clock (CLK_SCB).
- A system clock enable (clock_sys_en) from SRSS, which is used to derive a system clock CLK_SYS from CLK_PERI.
- An SCB clock enable (clock_scb_en) from the PCLK component in PERI. This clock enable is used to derive CLK_SCB from CLK_PERI.
- CLK_SCB can be divided from integer or fractional divider in the PERI block.
 The fractional divider causes varying cycle times in generated CLK_SCB.
 The integer divider must be used for I²C and SPI (synchronous interface).
 Both integer and fractional dividers can be used for UART (asynchronous interface).
- CLK_SCB is used for internally-clocked mode. Note that CLK_SCB is available only in Active/Sleep power modes. As a result, internally-clocked mode is not available in DeepSleep power mode. The serial interface protocols (UART TX/RX functionality and I²C/SPI master

functionality) are implemented using CLK_SCB as an “oversampled multiple” of the desired interface clock. For example, to implement a 100-kHz UART, CLK_SCB can be set to 1 MHz and the oversample factor set to 10 (SCBx_CTRL.OVS = 10 – 1).

- A low/0’ active system reset for Active functionality rst_sys_act_n.
- A low/0’ active system reset for DeepSleep functionality rst_sys_dpslp_n.

In externally-clocked slave mode, serial interface input signals are used as clock (I2C: “SCL”, SPI: “SCLK”). These clocks are asynchronous to CLK_SYS and CLK_AHB, which are derived from CLK_PERI. In externally-clocked slave mode, reset signals are derived from rst_sys_dpslp_n that have a synchronous de-assertion with reference to the serial interface clock.

23.2.5 Block Enable

More details about initializing is given in the description field of the SCBx_CTRL.ENABLED register.

Note: All registers flagged with “NonRetention” will also be reset to the default state when the block is disabled. This includes the SCBx_INTR_XXX and SCBx_INTR_XXX_SET registers.

23.2.6 Interrupt Interface

The SCB block has six types of interrupts. Each interrupt has dedicated registers. For details, see [SCB Interrupts on page 322](#) and the *Traveo II Body Controller High Registers TRM*.

Table 23-1. Interrupt Interface Signals and Registers


Interrupt	Functionality	Active/DeepSleep	Sync/Async	Registers
interrupt_master	I2C master and SPI master functionality	Active	Sync	SCBx_INTR_M, SCBx_INTR_M_SET, SCBx_INTR_M_MASK, SCBx_INTR_M_MASKED
interrupt_slave	I2C slave and SPI slave functionality	Active	Sync	SCBx_INTR_S, SCBx_INTR_S_SET, SCBx_INTR_S_MASK, SCBx_INTR_S_MASKED
interrupt_tx	UART transmitter and TX FIFO functionality	Active	Sync	SCBx_INTR_TX, SCBx_INTR_TX_SET, SCBx_INTR_TX_MASK, SCBx_INTR_TX_MASKED
interrupt_rx	UART receiver and RX FIFO functionality	Active	Sync	SCBx_INTR_RX, SCBx_INTR_RX_SET, SCBx_INTR_RX_MASK, SCBx_INTR_RX_MASKED
interrupt_i2c_ec	Externally clocked I2C slave functionality	DeepSleep	Async	SCBx_INTR_I2C_EC, SCBx_INTR_I2C_EC_MASK, SCBx_INTR_I2C_EC_MASKED
interrupt_spi_ec	Externally clocked SPI slave functionality	DeepSleep	Async	SCBx_INTR_SPI_EC, SCBx_INTR_SPI_EC_MASK, SCBx_INTR_SPI_EC_MASKED

The Active functionality interrupts are generated synchronously to CLK_PERI. The DeepSleep functionality interrupts are generated asynchronously to CLK_PERI and need synchronization in the CPUSS interrupt multiplexer.

For chips with a limited number of available interrupt lines, the SCB block also provides “combined functionality” interrupts as follows:

- interrupt_ec is the OR of interrupt_i2c_ec and interrupt_spi_ec.
- interrupt_ic is the OR of interrupt_master, interrupt_slave, interrupt_tx, and interrupt_rx.
- “interrupt” is the OR of all six individual interrupts.

Figure 23-2. Interrupt Lines



SCBx_INTR_M, SCBx_INTR_S, SCBx_INTR_TX, and SCBx_INTR_RX are interrupts from internal-clocked logic; SCBx_INTR_I2C_EC and SCBx_INTR_SPI_EC are interrupts from external-clocked logic.

Figure 23-3. SCBx_INTR_M Generation





Figure 23-4. SCBx_INTR_SPI_EC and interrupt_spi_ec Generation



Note: Interrupt request registers such as SCBx_INTR_M can be set only by hardware (HW RW1S) and cleared only by software (SW RW1C).

To avoid being triggered by events from previous transactions, whenever the firmware enables an interrupt mask register bit (for example, SCBx_INTR_M_MASK.I2C_STOP), it should clear the interrupt request register (for example, SCBx_INTR_M.I2C_STOP) in advance.

23.3 Operation Modes

23.3.1 Buffer Modes

Each SCB has 256 bytes of dedicated RAM for transmit and receive operation. This RAM can be configured in three different modes (FIFO, EZ, or CMD_RESP). The following sections give a high-level overview of each mode. The sections on each protocol will provide more details.

- Masters can only use FIFO mode
- Slaves can use all three modes.


Note: CMD Response mode is available only on DeepSleep-capable SCB

- UART only uses FIFO mode

Figure 23-5 shows the buffer modes using a dedicated SRAM.

- Two 32 deep FIFOs for up to 32-bit data elements (SCBx_CTRL.MEM_WIDTH register is '2')
- Two 64 deep FIFOs for up to 16-bit data elements (SCBx_CTRL.MEM_WIDTH register is '1')
- Two 128 deep FIFOs for up to 8-bit data elements (SCBx_CTRL.MEM_WIDTH register is '0')
- One 256 Byte EZ memory buffer
- One 256 Byte CMD_RESP memory buffer

Figure 23-5. Buffer Modes Using a Dedicated SRAM



23.3.1.1 FIFO Mode

In this mode the RAM is split into two 128-byte FIFOs, one for transmit (TX) and one for receive (RX). The FIFOs can be configured to be 8 bits x 128 elements or 16 bits x 64 elements.

FIFO mode of operation is available only in Active and Sleep power modes. However, the I²C address or SPI slave select can be used to wake the device from DeepSleep on the DeepSleep-capable SCB.

A write access to the TX FIFO uses the SCBx_TX_FIFO_WR.DATA register. A read access from the RX FIFO uses the SCBx_RX_FIFO_RD.DATA register.

Furthermore, it is possible that reading a data frame will not remove the data frame from the FIFO using the SCBx_RX_FIFO_RD_SILENT.DATA register.

Status is provided for both the RX and TX buffers. Multiple interrupt sources that indicate the status of the FIFOs are available, such as full or empty; see [SCB Interrupts on page 322](#).

23.3.1.2 EZ Mode

In easy (EZ) mode the RAM is used as a single 256-byte buffer. The external master sets a base address and reads and writes start from that base address.

EZ mode is available only for SPI slave and I²C slave. It is available only on the DeepSleep-capable SCB.

EZ mode is available in Active, Sleep, and DeepSleep power modes.

23.3.1.3 CMD_RESP Mode

Command Response (CMD_RESP) mode is similar to EZ mode except that the base address is provided by the CPU not the external master.

CMD_RESP mode is available only for SPI slave and I²C slave. It is available only on the DeepSleep-capable SCB. CMD_RESP mode operation is available in Active, Sleep, and DeepSleep power modes.

23.3.2 Clocking Modes

The SCB can be clocked either by an internal clock provided by the peripheral clock dividers or by the external master.

- UART, SPI Master, and I²C Master modes must use the internal clock, called CLK_SCB in the rest of this chapter.
- Only SPI slave and I²C slave can use the clock from an external master, and only the DeepSleep-capable SCB supports this.

Internally- and externally-clocked slave functionality is determined by two register fields of the SCBx_CTRL register:

- SCBx_CTRL.EC_AM_MODE indicates whether SPI slave selection or I²C address matching is clocked internally ('0') or externally ('1').
- SCBx_CTRL.EC_OP_MODE indicates whether the rest of the protocol operation (besides SPI slave selection and I²C address matching) is clocked internally ('0') or externally ('1').

When using externally-clocked slave functionality, it is important to realize that:

- FIFO mode is not supported with externally-clocked operation (SCBx_CTRL.EC_OP_MODE is '1')
- EZ and CMD_RESP modes are supported with externally-clocked operation (SCBx_CTRL.EC_OP_MODE is '1')
- Before going to DeepSleep mode, the SCBx_CTRL.EC_ACCESS register should be set to '1'. When waking up from DeepSleep mode and PLL is locked (CLK_SCB is at the expected frequency), the SCBx_CTRL.EC_ACCESS should be set to '0'.

The following table provides an overview of which clocking modes and which buffer modes are supported for each communication mode.

Table 23-2. Clock Mode Compatibility

	Internally-clocked ("IC")			Externally-clocked ("EC")		
	FIFO	EZ	CMD_RESP	FIFO	EZ	CMD_RESP
I ² C master	Yes	No	No	No	No	No
I ² C slave	Yes	Yes	No	Yes ^a	Yes	Yes
I ² C master-slave	Yes	No	No	No	No	No
SPI master	Yes	No	No	No	No	No
SPI slave	Yes	Yes	No	Yes ^b	Yes	Yes
UART transmitter	Yes	No	No	No	No	No
UART receiver	Yes	No	No	No	No	No

a. In DeepSleep mode the externally-clocked logic can handle slave address matching; it then triggers an interrupt to wake up the CPU. The slave can be programmed to stretch the clock, or NACK until internal logic takes over. This only applies to the DeepSleep-capable SCB.

b. In DeepSleep mode the externally-clocked logic can handle slave selection detection; it then triggers an interrupt to wake up the CPU. Writes will be ignored and reads will return 0xFF until internal logic takes over. This only applies to the DeepSleep-capable SCB.

23.4 Serial Peripheral Interface (SPI)

The SPI protocol is a synchronous serial interface protocol. Devices operate in either master or slave mode. The only master can initiate the data transfer. The SCB supports single-master-multiple-slaves topology for SPI. Multiple slaves up to four are supported with individual slave select lines.

In the Traveo II MCU, all SCB blocks support full SPI and only one SCB (SCB[0]) is available in DeepSleep power mode; this allows externally-clocked operations.

23.4.1 Features


- Supports master and slave functionality
- Supports three types of SPI protocols:
 - Motorola SPI - modes 0, 1, 2, and 3
 - Texas Instruments SPI, with coinciding and preceding data frame indicator - mode 1 only
 - National Semiconductor (MicroWire) SPI - mode 0 only
- Master supports up to four slave select lines
 - Each slave select has configurable active polarity (high or low)

- Slave select can be programmed to stay active for a whole transfer, or just for each byte
- Master supports late sampling for better timing margin
- Master supports continuous SPI clock
- Data frame size programmable from 4 bits to 32 bits
- Variable SELECT output signal timing (SPI master):
 - SELECT setup time (select active to SPI clock)
 - SELECT hold time (SPI clock to select inactive)
 - Inter-data frame deselect time (select inactive to select active)
- Parity support (odd and even parity)
- Interrupts or polling CPU interface
- Programmable oversampling
- MSB or LSB first
- Median filter available for inputs
- Supports FIFO mode, EZ mode (slave only), and CMD_RESP mode (slave only).
- Wake-up interrupt cause activated on slave selection (SCB[0] only)
- Local loop-back control

23.4.2 General Description

Figure 23-6 illustrates an example of SPI master with four slaves.

Figure 23-6. SPI Example



A standard SPI interface consists of four signals as follows.

- SCLK: Serial clock (clock output from the master, input to the slave).
- MOSI: Master-out-slave-in (data output from the master, input to the slave).
- MISO: Master-in-slave-out (data input to the master, output from the slave).
- SELECT: Typically an active low signal (output from the master, input to the slave).

A simple SPI data transfer involves the following: the master selects a slave by driving its SELECT line, then it drives data on the MOSI line and a clock on the SCLK line. The slave uses either of the edges of SCLK depending on the configuration to capture the data on the MOSI line; it also drives data on the MISO line, which is captured by the master.

By default, the SPI interface supports a data frame size of eight bits (1 byte). The data frame size can be configured to any value in the range 4 to 32 bits. The serial data can be transmitted either most significant bit (MSb) first or least significant bit (LSb) first.

Three different variants of the SPI protocol are supported by the SCB:

- Motorola SPI: This is the original SPI protocol.
- Texas Instruments SPI: A variation of the original SPI protocol, in which data frames are identified by a pulse on the SELECT line.
- National Semiconductors SPI: A half-duplex variation of the original SPI protocol.

Notes about duplex control:

- Motorola and Texas Instruments modes are full-duplex; National Semiconductors mode is half-duplex.
- Full-duplex modes also work similar to half-duplex, controlled by SCBx_TX_FIFO_CTRL.FREEZE or SCBx_RX_FIFO_CTRL.FREEZE, to transmit dummy data words or ignore received data words.
- The MOSI can be set to Hi-Z state using IOSS/GPIO configuration.

23.4.3 SPI Modes of Operation

23.4.3.1 Motorola SPI

The original SPI protocol was defined by Motorola. It is a full duplex protocol. Multiple data transfers may happen with the SELECT line held at '0'. As a result, slave devices must keep track of the progress of data transfers to separate individual data frames. When not transmitting data, the SELECT line is held at '1' and SCLK is typically pulled low.

Clock Modes of Motorola SPI

The Motorola SPI protocol has four different clock modes based on how data is driven and captured on the MOSI and MISO lines. These modes are determined by clock polarity (SCBx_SPI_CTRL.CPOL) and clock phase (SCBx_SPI_CTRL.CPHA).

Clock polarity determines the value of the SCLK line when not transmitting data. SCBx_SPI_CTRL.CPOL = 0 indicates that SCLK is '0' when not transmitting data. SCBx_SPI_CTRL.CPOL = 1 indicates that SCLK is '1' when not transmitting data.

Clock phase determines when data is driven and captured. SCBx_SPI_CTRL.CPHA = 0 means sample (capture data) on the leading (first) clock edge, while SCBx_SPI_CTRL.CPHA = 1 means sample on the trailing (second) clock edge, regardless of whether that clock edge is rising or falling. With SCBx_SPI_CTRL.CPHA = 0, the data must be stable for setup time before the first clock cycle.

- Mode 0: SCBx_SPI_CTRL.CPOL is '0', SCBx_SPI_CTRL.CPHA is '0': Data is driven on a falling edge of SCLK. Data is captured on a rising edge of SCLK. SCLK idle state is '0'.
- Mode 1; SCBx_SPI_CTRL.CPOL is '0', SCBx_SPI_CTRL.CPHA is '1': Data is driven on a rising edge of SCLK. Data is captured on a falling edge of SCLK. SCLK idle state is '0'.
- Mode 2: SCBx_SPI_CTRL.CPOL is '1', SCBx_SPI_CTRL.CPHA is '0': Data is driven on a rising edge of SCLK. Data is captured on a falling edge of SCLK. SCLK idle state is '1'.
- Mode 3: SCBx_SPI_CTRL.CPOL is '1', SCBx_SPI_CTRL.CPHA is '1': Data is driven on a falling edge of SCLK. Data is captured on a rising edge of SCLK. SCLK idle state is '1'.

Figure 23-7 illustrates driving and capturing of MOSI/MISO data as a function of SCBx_SPI_CTRL.CPOL and SCBx_SPI_CTRL.CPHA.

Figure 23-7. SPI Motorola, 4 Modes





Figure 23-8 shows a single 8-bit and two successive 8-bit data transfers in mode 0 (SCBx_SPI_CTRL.CPOL is '0', SCBx_SPI_CTRL.CPHA is '0').

Figure 23-8. SPI Motorola Data Transfer Example



Configuring SCB for SPI Motorola Mode

To configure the SCB for SPI Motorola mode, set various register bits in the following order:

1. Select SPI by writing '01' to the SCBx_CTRL.MODE register.
2. Select SPI Motorola mode by writing '00' to the SCBx_CTRL.MODE register.
3. Select the mode of operation in Motorola by writing to the SCBx_SPI_CTRL.CPHA and SCBx_SPI_CTRL.CPOL register.
4. Follow steps 2 to 4 mentioned in [Enabling and Initializing SPI on page 294](#).


For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

23.4.3.2 Texas Instruments SPI

The Texas Instruments' SPI protocol redefines the use of the SELECT signal. It uses the signal to indicate the start of a data transfer, rather than a low active slave select signal, as in the Motorola SPI. As a result, slave devices need not keep track of the progress of data transfers to separate individual data frames. The start of a transfer is indicated by a high active pulse of a single-bit transfer period. This pulse may occur one cycle before the transmission of the first data bit, or may coincide with the transmission of the first data bit. The TI SPI protocol supports only mode 1 (SCBx_SPI_CTRL.CPOL is '0' and SCBx_SPI_CTRL.CPHA is '1'): data is driven on a rising edge of SCLK and captured on a falling edge of SCLK.


[Figure 23-9](#) illustrates a single 8-bit data transfer and two successive 8-bit data transfers. The SELECT pulse precedes the first data bit. Note how the SELECT pulse of the second data transfer coincides with the last data bit of the first data transfer.

Figure 23-9. SPI TI Data Transfer Example



[Figure 23-10](#) illustrates a single 8-bit data transfer and two successive 8-bit data transfers. The SELECT pulse coincides with the first data bit of a frame.

Figure 23-10. SPI TI Data Transfer Example



Configuring SCB for SPI TI Mode

To configure the SCB for SPI TI mode, set various register bits in the following order:

1. Select SPI by writing '01' to the SCBx_CTRL.MODE register.
2. Select SPI TI mode by writing '01' to the SCBx_CTRL.MODE register.
3. Select the mode of operation in TI by writing to the SCBx_SPI_CTRL.SELECT_PRECEDE register ('1' configures the SELECT pulse to precede the first bit of next frame and '0' otherwise).
4. Follow steps 2 to 4 mentioned in [Enabling and Initializing SPI on page 294](#).

For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.


23.4.3.3 National Semiconductors SPI

The National Semiconductors' SPI protocol is a half-duplex protocol. Rather than transmission and reception occurring at the same time, they take turns. The transmission and reception data sizes may differ. A single idle (= '0') bit transfer period separates transmission from reception. However, the successive data transfers are not separated by an idle bit transfer period.

The National Semiconductors SPI protocol only supports mode 0.

[Figure 23-11](#) illustrates a single data transfer and two successive data transfers. In both cases the transmission data transfer size is eight bits and the reception data transfer size is four bits.

Figure 23-11. SPI NS Data Transfer Example



Configuring SCB for SPI NS Mode

To configure the SCB for SPI NS mode, set various register bits in the following order:

1. Select SPI by writing '01' to the SCBx_CTRL.MODE register.
2. Select SPI NS mode by writing '10' to the SCBx_CTRL.MODE register.
3. Follow steps 2 to 4 mentioned in [Enabling and Initializing SPI on page 294](#).

For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

23.4.4 SPI Buffer Modes

SPI can operate in three different buffer modes – FIFO, EZ, and CMD_RESP modes. The buffer is used in different ways in each of these modes. The following subsections explain each of these buffer modes in detail.

23.4.4.1 FIFO Mode

The FIFO mode has a TX FIFO for the data being transmitted and an RX FIFO for the data received. Each FIFO is constructed out of the SRAM buffer. The FIFOs are either 32 elements deep with 32-bit data elements or 64 elements deep with 16-bit data elements or 128 elements deep with 8-bit data elements. The width of a FIFO is configured using the SCBx_CTRL.MEM_WIDTH register.

FIFO mode is available only in Active and Sleep power modes, and not in the DeepSleep mode.

Transmit and receive FIFOs allow write and read accesses. A write access to the transmit FIFO uses the SCBx_TX_FIFO_WR register. A read access from the receive FIFO uses the SCBx_RX_FIFO_RD register.

Transmit and receive FIFO status information is available through status registers, SCBx_TX_FIFO_STATUS and SCBx_RX_FIFO_STATUS. It is possible to define a programmable threshold that indicates a number of FIFO entries, a trigger/event is generated when the following conditions are met:

- The transmit FIFO has a SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL. A trigger/event is generated when the number of entries in the transmit FIFO is less than SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL.
- The receive FIFO has an SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL. A trigger/event is generated when the number of receive FIFO entries is greater than the SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL.

These triggers can be connected to a DMA channel.

Furthermore, numerous interrupt status bits are provided for both the RX and TX FIFOs. These can be found looking at SCBx_INTR_TX and SCBx_INTR_RX.

DeepSleep to Active Transition

SCBx_CTRL.EC_AM_MODE = 1,
 SCBx_CTRL.EC_OP_MODE = 0, FIFO Mode.

MISO transmits 0xFF until internally-clocked logic takes over and CPU writes to TX FIFO. Data on MOSI is ignored until internally-clocked logic takes over. When the internally-clocked logic takes over, there is no guarantee that the internal clock will be at the correct frequency due to PLL/FLL locking times. This may lead to corrupted data in the RX FIFO. Therefore, it is recommended to clear the RX FIFO before writing new data into the TX FIFO after the transition from DeepSleep to Active. Another option is to disable CLK_SCB before going to DeepSleep, and then wait to enable it until the PLL and FLL have stabilized. The external master needs to be aware that when it reads 0xFF on MISO the device is not ready yet.

23.4.4.2 EZSPI Mode

The easy SPI (EZSPI) protocol is based on the Motorola SPI operating in any mode (0, 1, 2, or 3). It allows communication between master and slave without the need for CPU intervention. In Traveo II MCU, only one SCB block supports EZSPI mode; the DeepSleep-capable SCB.

The EZSPI protocol defines a single memory buffer with an 8-bit EZ address that indexes the buffer (256-entry array of eight bit per entry) located on the slave device. The EZ address is used to address these 256 locations. All EZSPI data transfers have 8-bit data frames.

The CPU writes and reads to the memory buffer through the SCBx_EZ_DATA registers. These accesses are word accesses, but only the least significant byte of the word is used.

EZSPI has three types of transfers: a write of the EZ address from the master to the slave, a write of data from the master to an addressed slave memory location, and a read by the master from an addressed slave memory location.

Note: When multiple bytes are read or written the master must keep SELECT low during the entire transfer.

EZ Address Write

A write of the EZ address starts with a command byte (0x00) on the MOSI line indicating the master's intent to write the EZ address. The slave then drives a reply byte on the MISO line to indicate that the command is acknowledged (0xFE) or not (0xFF). The second byte on the MOSI line is the EZ address.

Memory Array Write


A write to a memory array index starts with a command byte (0x01) on the MOSI line indicating the master's intent to write to the memory array. The slave then drives a reply byte on the MISO line to indicate that the command was registered (0xFE) or not (0xFF). Any additional write data bytes on the MOSI line are written to the memory array at locations indicated by the communicated EZ address. The EZ address is automatically incremented by the slave as bytes are written into the memory array. When the EZ address exceeds the maximum number of memory entries (256), it remains there and does not wrap around to 0. The EZ base address is reset to the address written in the EZ Address Write phase on each slave selection.

Memory Array Read

A read from a memory array index starts with a command byte (0x02) on the MOSI line indicating the master's intent to read from the memory array. The slave then drives a reply byte on the MISO line to indicate that the command was registered (0xFE) or not (0xFF). Any additional read data bytes on the MISO line are read from the memory array at locations indicated by the communicated EZ address. The EZ address is automatically incremented by the slave as bytes are read from the memory array. When the EZ address exceeds the maximum number of memory entries (256), it remains there and does not wrap around to 0. The EZ base address is reset to the address written in the EZ Address Write phase on each slave selection.

Figure 23-12 illustrates the write of EZ address, write to a memory array and read from a memory array operations in the EZSPI protocol.

Figure 23-12. EZSPI Example



Configuring SCB for EZSPI Mode

By default, the SCB is configured for non-EZ mode of operation. To configure the SCB for EZSPI mode, set the register bits in the following order:

1. Select EZ mode by writing '1' to the SCBx_CTRL.EZ_MODE register.
2. Follow the steps in [“Configuring SCB for SPI Motorola Mode on page 284.”](#)
3. Follow steps 2 to 4 mentioned in [Enabling and Initializing SPI on page 294.](#)

For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

DeepSleep to Active Transition

- **SCBx_CTRL.EC_AM_MODE = 1, SCBx_CTRL.EC_OP_MODE = 0, EZ Mode.**

MISO transmits 0xFF until the internally-clocked logic takes over. Data on MOSI is ignored until the internally-clocked logic takes over. When this happens, there is no guarantee that the internal clock will be at the correct frequency due to PLL/FLL locking times. This may lead to corrupted data on MISO and in the EZ memory. Therefore, it is recommended to disable CLK_SCB before going to DeepSleep, and then wait to enable it until the PLL/FLL have stabilized. The external master needs to be aware that when it reads 0xFF on MISO the device is not ready yet.

- **SCBx_CTRL.EC_AM_MODE = 1, SCBx_CTRL.EC_OP_MODE = 1, EZ Mode.**

When transitioning from DeepSleep to Active mode, there is no guarantee that the internal clock will be at the correct frequency due to PLL/FLL locking times. This situation limits the SPI SCLK frequency to 2 MHz. After the FLL/PLL outputs have stabilized the clock can run faster.

23.4.4.3 Command-Response Mode

The command-response mode is defined only for an SPI slave. In the Traveo II MCU, only one SCB (SCB[0]) supports the command-response mode. This mode has a single memory buffer, a base read address, a current read address, a base write address, and a current write address that are used to index the memory buffer. The base addresses are provided by the CPU. The current addresses are used by the slave to index the memory buffer for sequential accesses of the memory buffer. The memory buffer holds 256 8-bit data elements. The base and current addresses are in the range [0, 255].

The CPU writes and reads to the memory buffer through the SCBx_EZ_DATA registers. These accesses are word accesses, but only the least significant byte of the word is used.

The slave interface accesses the memory buffer using the current addresses. At the start of a write transfer (SPI slave

selection), the base write address is copied to the current write address. A data element write is to the current write address location. After the write access, the current address is incremented by '1'. At the start of a read transfer, the base read address is copied to the current read address. A data element read is to the current read address location. After the read data element is transmitted, the current read address is incremented by '1'.

If the current addresses equal the last memory buffer address (255), the current addresses are not incremented. Subsequent write accesses will overwrite any previously written value at the last buffer address. Subsequent read accesses will continue to provide the (same) read value at the last buffer address. The bus master should be aware of the memory buffer capacity in command-response mode.


The base addresses are provided through SCBx_CMD_RESP_CTRL.BASE_RD_ADDR and SCBx_CMD_RESP_CTRL.BASE_WR_ADDR. The current addresses are provided through SCBx_CMD_RESP_STATUS.CURR_RD_ADDR and SCBx_CMD_RESP_STATUS.CURR_WR_ADDR. At the end of a transfer (SPI slave de-selection), the difference between a base and current address indicates how many read/write accesses were performed. The block provides interrupt cause fields to identify the end of a transfer. Command-response mode operation is available in Active, Sleep, and DeepSleep power modes.

The command-response mode has two phases of operation:

- **Write phase** - The write phase begins with a selection byte, which has its last bit set to '0' indicating a write. The master writes 8-bit data elements to the slave's memory buffer following the selection byte. The slave's current write address is set to the slave's base write address. Received data elements are written to the current write address memory location. After each memory write, the current write address is incremented.
- **Read phase** - The read phase begins with a selection byte, which has its last bit set to '1' indicating a read. The master reads 8-bit data elements from the slave's memory buffer. The slave's current read address is set to the slave's base read address. Transmitted data elements are read from the current address memory location. After each read data element is transferred, the current read address is incremented.

During the reception of the first byte, the slave (MISO) transmits either 0x62 (ready) or a value different from 0x62 (busy). When disabled or reset, the slave transmits 0xFF (busy). The byte value can be used by the master to determine whether the slave is ready to accept the SPI request.

Figure 23-13. Command-Response Mode Example



Note that a slave's base addresses are updated by the CPU and not by the master.

DeepSleep to Active Transition

SCBx_CTRL.EC_AM_MODE = 1,
 SCBx_CTRL.EC_OP_MODE = 1, CMD_RESP Mode.

When transitioning from DeepSleep to Active mode there is no guarantee that the internal clock will be at the correct frequency due to PLL/FLL locking times. This situation limits the SPI SCLK frequency to 2 MHz. After the FLL/PLL outputs have stabilized the clock can run faster.

Configuring SCB for CMD_RESP Mode

By default, the SCB is configured for non-CMD_RESP mode of operation. To configure the SCB for CMD_RESP mode, set the register bits in the following order:

1. Select the CMD_RESP mode by writing '1' to the SCBx_CTRL.CMD_RESP_MODE register.
2. Follow the steps in [Configuring SCB for SPI Motorola Mode on page 284](#).
3. Follow steps 2 to 4 mentioned in [Enabling and Initializing SPI on page 294](#).

For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

23.4.5 Clocking and Oversampling

23.4.5.1 Clock Modes

The SCB SPI supports both internally- and externally-clocked operation modes. SCBx_CTRL.EC_AM_MODE and SCBx_CTRL.EC_OP_MODE register determine the SCB clock mode. SCBx_CTRL.EC_AM_MODE indicates whether SPI slave selection is clocked internally (0) or externally (1). SCBx_CTRL.EC_OP_MODE indicates whether the rest of the protocol operation (besides SPI slave selection) is clocked internally (0) or externally (1).

An externally-clocked operation uses a clock provided by the external master (SPI SCLK).

Note: In the Traveo II MCU only the DeepSleep-capable SCB supports externally-clocked mode of operation and only for SPI slave mode.

An internally-clocked operation uses the programmable clock dividers. For more information on system clocking, see the [Clocking System chapter on page 182](#).

The SCBx_CTRL.EC_AM_MODE and SCBx_CTRL.EC_OP_MODE can be configured in the following ways.

- SCBx_CTRL.EC_AM_MODE is '0' and SCBx_CTRL.EC_OP_MODE is '0': Use this configuration when only Active mode functionality is required.

- ❑ FIFO mode: Supported.
- ❑ EZ mode: Supported.
- ❑ Command-response mode: Not supported. The slave (MISO) transmits a value different from a ready (0x62) byte during reception of the first byte, if the command-response mode is attempted in this configuration.
- SCBx_CTRL.EC_AM_MODE is '1' and SCBx_CTRL.EC_OP_MODE is '0': Use this configuration when both Active and DeepSleep functionality are required. This configuration relies on the externally-clocked functionality to detect the slave selection and relies on the internally-clocked functionality to access the memory buffer.
 The “handover” from external to internal functionality relies on a busy/ready byte scheme. This scheme relies on the master to retry the current transfer when it receives a busy byte and requires the master to support busy/ready byte interpretation. When the slave is selected, SCBx_INTR_SPI_EC.WAKE_UP is set to '1'. The associated DeepSleep functionality interrupt brings the system into Active power mode.
- ❑ FIFO mode: Supported. The slave (MISO) transmits 0xFF until the CPU is awoken and the TX FIFO is populated. Any data on the MOSI line will be dropped until CLK_SCB is enabled see [DeepSleep to Active Transition on page 287](#) for more details
- ❑ EZ mode: Supported. In DeepSleep power mode, the slave (MISO) transmits a busy (0xFF) byte during the reception of the command byte. In Active power mode, the slave (MISO) transmits a ready (0xFE) byte during the reception of the command byte.
- ❑ CMD_RESP mode: Not supported. The slave transmits (MISO) a value different from a ready (0x62) byte during the reception of the first byte.
- SCBx_CTRL.EC_AM_MODE is '1' and SCBx_CTRL.EC_OP_MODE is '1'. Use this mode when both Active and DeepSleep functionality are required. When the slave is selected, SCBx_INTR_SPI_EC.WAKE_UP is set to '1'. The associated DeepSleep functionality interrupt brings the system into Active power mode. When the slave is deselected, SCBx_INTR_SPI_EC.EZ_STOP and/or SCBx_INTR_SPI_EC.EZ_WRITE_STOP are set to '1'.
- ❑ FIFO mode: Not supported.
- ❑ EZ mode: Supported.
- ❑ CMD_RESP mode: Supported.

Table 23-3. SPI Modes Compatibility

	Internally-clocked (IC)			Externally-clocked (EC)		
	FIFO	EZ	CMD_RESP	FIFO	EZ	CMD_RESP
SPI master	Yes	No	No	No	No	No
SPI slave	Yes	Yes	No	Yes ^a	Yes	Yes

a. In SPI slave FIFO mode, the externally-clocked logic does selection detection, then triggers an interrupt to wake up the CPU. Writes will be ignored and reads will return 0xFF until the CPU is ready and the FIFO is populated.

If SCBx_CTRL.EC_OP_MODE is '1', the external interface logic accesses the memory buffer on the external interface clock (SPI SCLK). This allows for EZ and CMD_RESP mode functionality in Active and DeepSleep power modes.

In Active system power mode, the memory buffer requires arbitration between external interface logic (on SPI SCLK) and the CPU interface logic (on system peripheral clock). This arbitration always gives the highest priority to the external interface logic (host accesses). The external interface logic takes two serial interface clock/bit periods for SPI. During this period, the internal logic is denied service to the memory buffer. The Traveo II MCU provides two programmable options to address this “denial of service”:

- If the SCBx_CTRL.BLOCK is '1': An internal logic access to the memory buffer is blocked until the memory buffer is granted and the external interface logic has completed access. This option provides normal SCB register functionality, but the blocking time introduces additional internal bus wait states.
- If the SCBx_CTRL.BLOCK is '0': An internal logic access to the memory buffer is not blocked, but fails

when it conflicts with an external interface logic access. A read access returns the value 0xFFFF:FFFF and a write access is ignored. This option does not introduce additional internal bus wait states, but an access to the memory buffer may not take effect. In this case, the following failures are detected:

- ❑ Read Failure: A read failure is easily detected because the returned value is 0xFFFF:FFFF. This value is unique as non-failing memory buffer read accesses return an unsigned byte value in the range 0x0000:0000-0x0000:00ff.
- ❑ Write Failure: A write failure is detected by reading back the written memory buffer location, and confirming that the read value is the same as the written value.

For both options, a conflicting internal logic access to the memory buffer sets SCBx_INTR_TX.BLOCKED field to '1' (for write accesses) and SCBx_INTR_RX.BLOCKED field to '1' (for read accesses). These fields can be used as either status fields or as interrupt cause fields (when their associated mask fields are enabled).

If a series of read or write accesses is performed and SCBx_CTRL.BLOCK is '0', a failure is detected by comparing the "logical OR" of all read values to 0xFFFF:FFFF and checking the SCBx_INTR_TX.BLOCKED and SCBx_INTR_RX.BLOCKED fields to determine whether a failure occurred for a series of write or read operations.

23.4.5.2 Using SPI Master to Clock Slave

In a normal SPI master mode transmission, the SCLK is generated only when the SCB is enabled and data is being transmitted. This can be changed to always generate a clock on the SCLK line while the SCB is enabled. This is used when the slave uses the SCLK for functional operations other than the SPI functionality. To enable this, write '1' to the SCBx_SPI_CTRL.SCLK_CONTINUOUS register.

23.4.5.3 Oversampling and Bit Rate

SPI Master Mode

The SPI master does not support externally-clocked mode. In internally-clocked mode, the logic operates under internal clock. The internal clock has a higher frequency than the interface clock (SCLK), such that the master can oversample its input signals (MISO).

The SCBx_CTRL.OVS register specify the oversampling. The oversampling rate is calculated as the value in SCBx_CTRL.OVS register + 1. In SPI master mode, the valid range for oversampling is 4 to 16, when MISO is used; if MISO is not used then the valid range is 2 to 16. The bit rate is calculated as follows.

$$\text{Bit Rate} = \text{Input Clock} / \text{SCBx\_CTRL.OVS}$$

Hence, with an input clock of 100 MHz, the maximum bit rate is 25 Mbps with MISO, or 50 Mbps without MISO.

The numbers above indicate how fast the SCB hardware can run SCLK. It does not indicate that the master will be able to correctly receive data from a slave at those speeds. To determine that, the path delay of MISO must be calculated. It can be calculated using the following equation:

$$\frac{1}{2} * t_{SCLK} \geq t_{SCLK\_PCB\_D} + t_{DSO} + t_{SCLK\_PCB\_D} + t_{DSI} \quad \text{Equation 23-1}$$

Where:

t_{SCLK} is the period of the SPI clock

$t_{SCLK\_PCB\_D}$ is the SCLK PCB delay from master to slave

t_{DSO} is the total internal slave delay, time from SCLK edge at slave pin to MISO edge at slave pin


$t_{SCLK\_PCB\_D}$ is the MISO PCB delay from slave to master

t_{DSI} is the master setup time

Most slave datasheets will list t_{DSO} , It may have a different name; look for MISO output valid after SCLK edge. Most master datasheets will also list t_{DSI} , or master setup time. $t_{SCLK\_PCB\_D}$ and $t_{SCLK\_PCB\_D}$ must be calculated based on specific PCB geometries.

After doing these calculations, if the desired speed cannot be achieved then consider using the MISO late sample feature of the SCB. MISO late sample tells the SCB to sample the incoming MISO signal on the next edge of SCLK, thus allowing for a one-half SCLK cycle more timing margin, see [Figure 23-14](#).

Figure 23-14. MISO Sampling Timing



This changes the equation to:

$$t_{SCLK} \geq t_{SCLK\_PCB\_D} + t_{DSO} + t_{SCLK\_PCB\_D} + t_{DSI} \quad \text{Equation 23-2}$$

Because late sample allows for better timing, it is recommended to leave it enabled all the time.

The t_{DSI} specification in the device datasheet assumes that the late sample is enabled.

Note: The SCBx_SPI_CTRL.LATE_MISO_SAMPLE is set to '0' by default.

SPI Slave Mode

In SPI slave mode, the SCBx_CTRL.OVS register is not used. The data rate is determined by Equation 24-1 and Equation 24-2. Late MISO sample is determined by the external master and not by SCBx_SPI_CTRL.LATE_MISO_SAMPLE.

For Traveo II MCUs, tDSO is given in the device datasheet. For internally-clocked mode, it is proportional to the frequency of the internal clock. For example, it may be 20 nsec + 3 × tCLK_SCB. Assuming 0 nsec PCB delays, and a 0 nsec external master tDSI Equation 24-1 can be rearranged to

$$tCLK\_SCB \leq ((tSCLK) - 40 \text{ nsec})/6.$$

23.4.6 SPI Master SELECT Output Timing Control

The SPI master SELECT output signal “spi_select” timing is made variable. This applies to:


- The SELECT setup time (select active to SPI clock)
- The SELECT hold time (SPI clock to select inactive)

- The inter-data frame deselect time (select inactive to select active)

The following options can be selected for these delays:

- SELECT setup time (SCBx_SPI_CTRL.SSEL_SETUP_DEL register):
 - When SCBx_SPI_CTRL.CPHA = 0: 0.75 or 1.75 SPI clock cycles
 - When SCBx_SPI_CTRL.CPHA = 1: 0.25 or 1.25 SPI clock cycles
- SELECT hold time (SCBx_SPI_CTRL.SSEL_HOLD_DEL register):
 - When SCBx_SPI_CTRL.CPHA = 0: 0.25 or 1.25 SPI clock cycles
 - When SCBx_SPI_CTRL.CPHA = 1: 0.75 or 1.75 SPI clock cycles


Figure 23-15. SELECT Setup/hold Delay



- () means the activation/deactivation timing of SELECT from the sampling edge of SCLK.

- INTER-FRAME deselect time (SCBx_SPI_CTRL.SSEL_INTER_FRAME_DEL register):
 - 1.5 SPI clock cycles or 2.5 SPI clock cycles

Figure 23-16. SELECT INTER-FRAME Deselect Time



23.4.7 SPI Parity Functionality

Parity functionality is added to SPI mode.

- This applies to the SPI master and SPI slave with internally-clocked operation.
- Parity functionality adds a parity bit to the data frame and is used to identify single-bit data frame errors. The parity bit directly follows the data frame bits.
- Parity functionality can be configured to be enabled or disabled using SCBx_SPI_TX_CTRL.PARITY_ENABLED and SCBx_SPI_RX_CTRL.PARITY_ENABLED individually.
- When transmitting, a parity bit can be inserted. When receiving, the parity bit can be checked. If parity fails, it is possible to select whether the received data is sent to the RX FIFO or is dropped and lost, using the SCBx_SPI_RX_CTRL.DROP_ON_PARITY_ERROR register.
- Even and odd parity is supported (SCBx_SPI_TX_CTRL.PARITY, SCBx_SPI_RX_CTRL.PARITY).

23.4.8 Loop-back

In SPI Master mode, SCB supports internal loop-back from an output signal for MOSI to an input signal for MISO without affecting the information on the pins. It is configured using the SCBx_SPI_CTRL.LOOPBACK register.

This loop-back is not supported in National Semiconductors mode.

23.4.9 Enabling and Initializing SPI

The SPI must be programmed in the following order:

1. Program protocol specific information using the SCBx_SPI_CTRL register. This includes selecting the sub-modes of the protocol (MODE), master-slave functionality (MASTER_MODE), one of four SELECT (SSEL), whether SELECT stays active for a whole transfer or just for each data frame width (SSEL_CONTINUOUS), and SELECT polarity (SSEL_POLARITY0-3). EZSPI and CMD_RESP can be used with slave mode only.
2. Program the generic transmitter and receiver information using the SCBx_TX_CTRL and SCBx_RX_CTRL registers:
 - a. Specify the data frame width. This should always be 8 for EZSPI and CMD_RESP.
 - b. Specify whether MSb or LSB is the first bit to be transmitted/received. This should always be MSb first for EZSPI and CMD_RESP.
3. Program the transmitter and receiver FIFOs using the SCBx_TX_FIFO_CTRL and SCBx_RX_FIFO_CTRL registers respectively, as shown in SCBx_TX_FIFO_CTRL/SCBx_RX_FIFO_CTRL registers. Only for FIFO mode:
 - a. Set the trigger level (TRIGGER_LEVEL).
 - b. Clear the transmitter and receiver FIFO and Shift registers (CLEAR).

4. Enable the block (write a '1' to the SCBx_CTRL.ENABLED register). After the block is enabled, control bits should not be changed. Changes should be made after disabling the block; for example, to modify the operation mode (from Motorola mode to TI mode) or to go from externally-clocked to internally-clocked operation. The change takes effect only after the block is re-enabled. Note that re-enabling the block causes reinitialization and the associated state is lost (for example, FIFO content).

23.4.10 I/O Pad Connection

23.4.10.1 SPI Master

Figure 23-17 and Figure 23-12 list the use of the I/O pads for SPI master.

Figure 23-17. SPI Master I/O Pad Connections




Table 23-4. SPI Master I/O Pad Connection Usage

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
SCLK	Strong (CMOS output)	<i>spi_clk_out_en</i> <i>spi_clk_out</i>	Transmit a clock signal
SELECT	Strong (CMOS output)	<i>spi_select_out_en</i> <i>spi_select_out</i>	Transmit a select signal
MOSI	Strong (CMOS output)	<i>spi_mosi_out_en</i> <i>spi_mosi_out</i>	Transmit a data element
MISO	Input only	<i>spi_miso_in</i>	Receive a data element

23.4.10.2 SPI Slave

Figure 23-18 and Figure 23-13 list the use of I/O pads for SPI slave.

Figure 23-18. SPI Slave I/O Pad Connections




Table 23-5. SPI Slave I/O Signal Description

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
SCLK	Input only	<code>spi_clk_in</code>	Receive a clock signal
SELECT	Input only	<code>spi_select_in</code>	Receive a select signal
MOSI	Input only	<code>spi_mosi_in</code>	Receive a data element
MISO	Strong (CMOS output), or open drain drives low	<code>spi_miso_out_en</code> <code>spi_miso_out</code>	Transmit a data element

23.4.11 SPI Registers

The SPI interface is controlled using a set of 32-bit control and status registers listed in Table 23-19. For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

23.5 UART

The Universal Asynchronous Receiver/Transmitter (UART) protocol is an asynchronous serial interface protocol. UART communication is typically point-to-point. The UART interface consists of two signals:

- TX: Transmitter output
- RX: Receiver input

Additionally, two side-band signals are used to implement flow control in UART. Note that the flow control only applies to TX functionality.

- Clear to Send (CTS): This is an input signal to the transmitter. When active, it indicates that the slave is ready for the master to transmit data.
- Ready to Send (RTS): This is an output signal from the receiver. When active, it indicates that the receiver is ready to receive data.


23.5.1 Features

- Supports UART protocol
 - Standard UART
 - Multi-processor mode
- SmartCard (ISO7816) reader
- IrDA
- Supports LIN
 - Break detection
 - Baud rate detection
 - Collision detection (ability to detect that a driven bit value is not reflected on the bus, indicating that another component is driving the same bus)
- Data frame size programmable from 4 to 16 bits
- Programmable number of STOP bits, which can be set in terms of half bit periods between 1 and 4
- Parity support (odd and even parity)
- Median filter on RX input
- Programmable oversampling
- Start skipping
- FIFO mode operation only
- Local loop-back control

23.5.2 General Description

Figure 23-19 illustrates a standard UART TX and RX.

Figure 23-19. UART Example



A typical UART transfer consists of a Start Bit followed by multiple Data Bits, optionally followed by a Parity Bit and finally completed by one or more Stop Bits. The Start and Stop bits indicate the start and end of data transmission. The Parity bit is sent by the transmitter and is used by the receiver to detect single-bit errors. Because the interface does not have a clock (asynchronous), the transmitter and receiver use their own clocks; thus, the transmitter and receiver need to agree on the baud rate.

By default, UART supports a data frame width of eight bits. However, this can be configured to any value in the range of 4 to 9. This does not include start, stop, and parity bits. The number of stop bits can be in the range of 1 to 7 (SCBx_UART_TX_CTRL.STOP_BITS, SCBx_UART_RX_CTRL.STOP_BITS). The parity bit can be either enabled or disabled. If enabled, the type of parity can be set to either even parity or odd parity. The option of using the parity bit is available only in the Standard UART and SmartCard UART modes. For IrDA UART mode, the parity bit is automatically disabled. Figure 23-25 depicts the default configuration of the UART interface of the SCB.

Note: The UART interface does not support external clocking operation. Hence, UART operates only in the Active and Sleep system power modes. UART also supports only the FIFO buffer mode.

Note: The behavior of UART when an error is detected in a start or stop period is determined by the SCBx_UART_RX_CTRL.DROP_ON_FRAME_ERROR register.

23.5.3 UART Modes of Operation

23.5.3.1 Standard Protocol

A typical UART transfer consists of a start bit followed by multiple data bits, optionally followed by a parity bit and finally completed by one or more stop bits. The start bit value is always '0', the data bits values are dependent on the data transferred, the parity bit value is set to a value guaranteeing an even or odd parity over the data bits, and the stop bit value is '1'. The parity bit is generated by the transmitter and can be used by the receiver to detect single-bit transmission errors. When not transmitting data, the TX line is '1' – the same value as the stop bits.

Because the interface does not have a clock, the transmitter and receiver need to agree upon the baud rate. The transmitter and receiver have their own internal clocks. The receiver clock runs at a higher frequency than the bit transfer frequency, such that the receiver may oversample the incoming signal.


The transition of a stop bit to a start bit is represented by a change from '1' to '0' on the TX line. This transition can be used by the receiver to synchronize with the transmitter clock. Synchronization at the start of each data transfer allows error-free transmission even in the presence of

frequency drift between transmitter and receiver clocks. The required clock accuracy is dependent on the data transfer size. The stop period or the amount of stop bits between successive data transfers is typically agreed upon between

transmitter and receiver, and is typically in the range of 1 to 3-bit transfer periods.


Figure 23-20 illustrates the UART protocol.

Figure 23-20. UART, Standard Protocol Example




The receiver oversamples the incoming signal; the value of the sample point in the middle of the bit transfer period (on the receiver's clock) is used. Figure 23-21 illustrates this.

Figure 23-21. UART, Standard Protocol Example (Single Sample)



Alternatively, three samples around the middle of the bit transfer period (on the receiver's clock) are used for a majority vote to increase accuracy; this is enabled by enabling the RX_CTRL.MEDIAN register. Figure 23-22 illustrates this.

Figure 23-22. UART, Standard Protocol (Multiple Samples)



Parity

This functionality adds a parity bit to the data frame and is used to identify single-bit data frame errors. The parity bit directly follows the data frame bits. The transmitter calculates the parity bit (when SCBx_UART_TX_CTRL.PARITY_ENABLED is 1) from the data frame bits, such that data frame bits and parity bit have an even (SCBx_UART_TX_CTRL.PARITY is 0) or odd (SCBx_UART_TX_CTRL.PARITY is 1) parity. The receiver checks the parity bit (when SCBx_UART_RX_CTRL.PARITY_ENABLED is 1) from the received data frame bits, such that data frame bits and parity bit have an even (SCBx_UART_RX_CTRL.PARITY is 0) or odd (SCBx_UART_RX_CTRL.PARITY is 1) parity.


Parity applies to both TX and RX functionality and dedicated control fields are available.

- Transmit functionality: SCBx_UART_TX_CTRL.PARITY and SCBx_UART_TX_CTRL.PARITY_ENABLED.
- Receive functionality: SCBx_UART_RX_CTRL.PARITY and SCBx_UART_RX_CTRL.PARITY_ENABLED.

When a receiver detects a parity error, the data frame is either put in RX FIFO (SCBx_UART_RX_CTRL.DROP_ON_PARITY_ERROR is 0) or dropped (SCBx_UART_RX_CTRL.DROP_ON_PARITY_ERROR is 1).

The following figure illustrates the parity functionality (8-bit data frame).

Figure 23-23. UART Parity Examples



Start Skipping


Start skipping only applies to receive functionality. The standard UART mode supports “start skipping”. Regular receive operation synchronizes on the START bit period (a 1-to-0 transition on the UART RX line), start skipping receive operation synchronizes on the first received data frame bit, which must be a ‘1’ (a 0-to-1 transition on UART RX).

Start skipping is used to allow for wake up from system DeepSleep mode using UART. The process is described as follows:

1. Before entering DeepSleep power mode, UART receive functionality is disabled and the GPIO is programmed to set an interrupt cause to ‘1’ when UART RX line has a ‘1’ to ‘0’ transition (START bit).
2. While in DeepSleep mode, the UART receive functionality is not functional.
3. The GPIO interrupt is activated on the START bit and the system transitions from DeepSleep to Active power mode.
4. The CPU enables UART receive functionality, with SCBx_UART_RX_CTRL.SKIP_START bitfield set to ‘1’.
5. The UART receiver synchronizes data frame receipt on the next ‘0’ to ‘1’ transition. If the UART receive functionality is enabled in time, this is the transition from the START bit to the first received data frame bit.
6. The UART receiver proceeds with normal operation; that is, synchronization of successive data frames is on the START bit period.

Figure 23-24 illustrates the process.

Figure 23-24. UART Start Skip and Wakeup from DeepSleep



Note that this process only works for lower baud rates. The DeepSleep to Active power mode transition and CPU enabling the UART receive functionality should take less than a 1-bit period to ensure that the UART receiver is active in time to detect the ‘0’ to ‘1’ transition.

In step 4 of the above process, it takes some time for the firmware to finish the wakeup interrupt routine and enable the UART receive functionality, before the block can detect the input rising edge on the UART RX line. If the above steps cannot be completed in less than 1 bit period, then it is recommended to first send a “dummy” byte to the device to wake it up before sending real UART data. In this case, the SCBx_UART_RX_CTRL.SKIP_START bit can be left as 0.

Break Detection

Break detection is supported in the standard UART mode. This functionality detects when UART RX line is low (0) for more than SCBx_UART_RX_CTRL.BREAK_WIDTH bit periods. The break width should be larger than the maximum number of low (0) bit periods in a regular data transfer, plus an additional 1-bit period. The additional 1-bit period is a minimum requirement and preferably should be larger. The additional bit periods account for clock inaccuracies between transmitter and receiver.


For example, in an 8-bit data frame with parity support, the maximum number of low (0) bit periods is 10 (START bit, 8 ‘0’ data frame bits, and one ‘0’ parity bit). Therefore, the

break width should be larger than $10 + 1 = 11$ (SCBx_UART_CTRL.BREAK_WIDTH can be set to 11).

Note that the break detection only applies to receive functionality. A UART transmitter can generate a break by temporarily increasing SCBx_TX_CTRL.DATA_WIDTH and transmitting an all “zeroes data” frame. A break is used by the transmitter to signal a special condition to the receiver. This condition may result in a reset, shut down, or initialization sequence at the receiver.

Break detection is part of the LIN protocol. When a break is detected, the SCBx_INTR_RX.BREAK_DETECT interrupt cause is set to ‘1’. Figure 23-25 illustrates a regular data frame and break frame (8-bit data frame, parity support, and a break width of 12-bit periods). When SCBx_UART_RX_CTRL.BREAK_LEVEL is set to ‘1’, idle line detection is possible. For example, after successive transfer of several UART data frames, an idle (high) level longer than normal data frame length (start+8data+1parity +1stop) indicates the end of this successive transfer.

Figure 23-25. UART - Regular Frame and Data Frame



Flow Control

The standard UART mode supports flow control. This modem flow control controls the pace at which the transmitter transfers data to the receiver. Modem flow control is enabled through the `SCBx_UART_FLOW_CTRL.CTS_ENABLED` register field. When this field is '0', the transmitter transfers data when its TX FIFO is not empty. When '1', the transmitter transfers data when UART CTS line is active and its TX FIFO is not empty.

Note that the flow control only applies to TX functionality. Two UART side-band signals are used to implement flow control:


- **UART RTS (uart_rts_out):** This is an output signal from the receiver. When active, it indicates that the receiver is ready to receive data (RTS: Ready to Send).

- **UART CTS (uart_cts_in):** This is an input signal to the transmitter. When active, it indicates that the transmitter can transfer data (CTS: Clear to Send).

The receiver's `uart_rts_out` signal is connected to the transmitter's `uart_cts_in` signal. The receiver's `uart_rts_out` signal is derived by comparing the number of used receive FIFO entries with the `SCBx_UART_FLOW_CTRL.TRIGGER_LEVEL` field. If the number of used receive FIFO entries are less than `SCBx_UART_FLOW_CTRL.TRIGGER_LEVEL`, `uart_rts_out` is activated.

Typically, the UART side-band signals are active low. However, sometimes active high signaling is used. Therefore, the polarity of the side-band signals can be controlled using `SCBx_UART_FLOW_CTRL.RTS_POLARITY` and `SCBx_UART_FLOW_CTRL.CTS_POLARITY` bitfields. [Figure 23-26](#) gives an overview of the flow control functionality.


Figure 23-26. UART Flow Control Connection



23.5.3.2 UART Multi-Processor Mode

The UART_MP (multi-processor) mode is defined with single-master-multi-slave topology, as [Figure 23-27](#) shows. This mode is also known as UART 9-bit protocol because the data field is nine bits wide. UART_MP is part of standard UART mode.

Figure 23-27. UART MP Mode Bus Connections



The main properties of UART_MP mode are:

- Single master with multiple slave concept (multi-drop network).
- Each slave is identified by a unique address.
- Using 9-bit data field, with the ninth bit as address/data flag (MP bit). When set high, it indicates an address byte; when set low it indicates a data byte. A data frame is illustrated in [Figure 23-28](#).
- Parity bit is disabled.

Figure 23-28. UART MP Address and Data Frame




The SCB can be used either as a master or slave device in UART_MP mode. Both SCBx_TX_CTRL and SCBx_RX_CTRL registers should be set to 9-bit data frame size. When the SCB works as UART_MP master device, the firmware changes the MP flag for every address or data frame. When it works as UART_MP slave device, the SCBx_UART_RX_CTRL.MP_SCB register should be set to '1'. The SCBx_RX_MATCH register should be set for the slave address and address mask. The matched address is written in the RX FIFO when SCBx_CTRL.ADDR_ACCEPT register is set to '1'. If received address does not match its own address, then the interface ignores the following data, until the next address is received for compare.

single-master-multi-slave topology. There is one master node and multiple slave nodes on the LIN bus. The SCB UART supports only the LIN slave functionality. The LIN specification defines both physical layer (layer 1) and data link layer (layer 2). [Figure 23-29](#) illustrates the UART_LIN and LIN transceiver.

23.5.3.3 UART Local Interconnect Network (LIN) Mode

The LIN protocol is supported by the SCB as part of the standard UART. LIN is designed with

Figure 23-29. UART_LIN and LIN Transceiver




LIN protocol defines two tasks:

- Master task: This task involves sending a header packet to initiate a LIN transfer.
- Slave task: This task involves transmitting or receiving a response.

The master node supports master task and slave task; the slave node supports only slave task, as shown in [Figure 23-30](#).

Figure 23-30. LIN Bus Nodes and Tasks




LIN Frame Structure

LIN is based on the transmission of frames at pre-determined moments of time. A frame is divided into header and response fields, as shown in [Figure 23-31](#).

- The header field consists of:
 - Break field (at least 13 bit periods with the value '0').
 - Sync field (a 0x55 byte frame). A sync field can be used to synchronize the clock of the slave task with that of the master task.
 - Identifier field (a frame specifying a specific slave).
- The response field consists of data and checksum.

Figure 23-31. LIN Frame Structure



In LIN protocol communication, the least significant bit (LSb) of the data is sent first and the most significant bit (MSb) last. The start bit is encoded as zero and the stop bit is encoded as one. The following sections describe all the byte fields in the LIN frame.

Break Field


Every new frame starts with a break field, which is always generated by the master. The break field has logical zero with a minimum of 13 bit times and followed by a break delimiter. The break field structure is as shown in Figure 23-32.

Figure 23-32. LIN Break Field



Sync Field. This is the second field transmitted by the master in the header field; its value is 0x55. A sync field can be used to synchronize the clock of the slave task with that of the master task for automatic baud rate detection. Figure 23-33 shows the LIN sync field structure.

Figure 23-33. LIN Sync Field



Protected identifier (PID) Field. A PID field consists of two sub-fields: the frame identifier (bits 0-5) and the parity (bit 6 and bit 7). The PID field structure is shown in Figure 23-34.

- Frame identifier: The frame identifiers are split into three categories
 - Values 0 to 59 (0x3B) are used for signal carrying frames
 - 60 (0x3C) and 61 (0x3D) are used to carry diagnostic and configuration data
 - 62 (0x3E) and 63 (0x3F) are reserved for future protocol enhancements
- Parity: Frame identifier bits are used to calculate the parity

Figure 23-34 shows the PID field structure.

Figure 23-34. PID Field



Data. In LIN, every frame can carry a minimum of one byte and maximum of 8 bytes of data. Here, the LSB of the data byte is sent first and the MSB of the data byte is sent last.

Checksum. The checksum is the last byte field in the LIN frame. It is calculated by inverting the 8-bit sum along with carryover of all data bytes only or the 8-bit sum with the carryover of all data bytes and the PID field. There are two types of checksums in LIN frames. They are:

- Classic checksum: the checksum calculated over all the data bytes only (used in LIN 1.x slaves).
- Enhanced checksum: the checksum calculated over all the data bytes along with the protected identifier (used in LIN 2.x slaves).

LIN Frame Types

The type of frame refers to the conditions that need to be valid to transmit the frame. According to the LIN

specification, there are five different types of LIN frames. A node or cluster does not have to support all frame types.

Unconditional Frame. These frames carry the signals and their frame identifiers (of 0x00 to 0x3B range). The subscriber will receive the frames and make it available to the application; the publisher of the frame will provide the response to the header.

Event-Triggered Frame. The purpose of an event-triggered frame is to increase the responsiveness of the LIN cluster without assigning too much of the bus bandwidth to polling of multiple slave nodes with seldom occurring events. Event-triggered frames carry the response of one or more unconditional frames. The unconditional frames associated with an event-triggered frame should:

- Have equal length
- Use the same checksum model (either classic or enhanced)
- Reserve the first data field to its protected identifier

- Be published by different slave nodes
- Not be included directly in the same schedule table as the event-triggered frame

Sporadic Frame. The purpose of sporadic frames is to merge some dynamic behavior into the schedule table without affecting the rest of the schedule table. These frames have a group of unconditional frames that share the frame slot. When the sporadic frame is due for transmission, the unconditional frames are checked if they have any updated signals. If no signals are updated, no frame will be transmitted and the frame slot will be empty.

Diagnostic Frames. Diagnostic frames always carry transport layer, and contains eight data bytes.

The frame identifier for diagnostic frame is:

- Master request frame (0x3C), or
- Slave response frame (0x3D)

Before transmitting a master request frame, the master task queries its diagnostic module to see if it will be transmitted or if the bus will be silent. A slave response frame header will be sent unconditionally. The slave tasks publish and subscribe to the response according to their diagnostic modules.

Reserved Frames. These frames are reserved for future use; their frame identifiers are 0x3E and 0x3F.

LIN Go-To-Sleep and Wake-Up

The LIN protocol has the feature of keeping the LIN bus in Sleep mode if the master sends the go-to-sleep command. The go-to-sleep command is a master request frame (ID = 0x3C) with the first byte field equal to 0x00 and the remaining fields set to 0xFF. The slave node application may still be active after the go-to-sleep command is received. This behavior is application specific. The LIN slave nodes automatically enter Sleep mode if the LIN bus inactivity is more than four seconds.

Wake-up can be initiated by any node connected to the LIN bus – either LIN master or any of the LIN slaves by forcing the bus to be dominant for 250 μ s to 5 ms. Each slave should detect the wakeup request and be ready to process headers within 100 ms. The master should also detect the wakeup request and start sending headers when the slave nodes are active.

To support LIN, a dedicated (off-chip) line driver/receiver is required. Supply voltage range on the LIN bus is 7 V to 18 V. Typically, LIN line drivers will drive the LIN line with the value provided on the SCB TX line and present the value on the LIN line to the SCB RX line. By comparing TX and RX lines in the SCB, bus collisions can be detected (indicated by the SCBx_INTR_TX.UART_ARB_LOST register).

Configuring the SCB as Standard UART Interface

To configure the SCB as a standard UART interface, set various register bits in the following order:

1. Configure the SCB as UART interface by writing '10' to the SCBx_CTRL.MODE register.
2. Configure the UART interface to operate as a standard protocol by writing '00' to the SCBx_UART_CTRL.MODE register.
3. To enable the UART MP or UART LIN mode, write '1' to the SCBx_UART_RX_CTRL.MP_MODE or SCBx_UART_RX_CTRL.LIN_MODE register.
4. Follow steps 2 to 4 described in [Enabling and Initializing UART on page 307](#).

For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

23.5.3.4 SmartCard (ISO7816)


ISO7816 is an asynchronous serial interface, defined with single-master-single slave topology. ISO7816 defines both Reader (master) and Card (slave) functionality. For more information, refer to the [ISO7816 Specification](#). Only master (reader) function is supported by the SCB. This block provides the basic physical layer support with asynchronous character transmission. UART_TX line is connected to SmartCard I/O line, by internally multiplexing between UART_TX and UART_RX control modules.

The SmartCard transfer is similar to a UART transfer, with the addition of a negative acknowledgment (NACK) that may be sent from the receiver to the transmitter. A NACK is always '0'. Both master and slave may drive the same line, although never at the same time.

A SmartCard transfer has the transmitter drive the start bit and data bits (and optionally a parity bit). After these bits, it enters its stop period by releasing the bus. Releasing results in the line being '1' (the value of a stop bit). After one bit transfer period into the stop period, the receiver may drive a NACK on the line (a value of '0') for one bit transfer period. This NACK is observed by the transmitter, which reacts by extending its stop period by one bit transfer period (when SCBx_UART_TX_CTRL.RETRY_ON_NACK = 1). For this protocol to work, the stop period should be longer than one bit transfer period. Note that a data transfer with a NACK takes one bit transfer period longer than a data transfer without a NACK. Typically, implementations use a tristate driver with a pull-up resistor, such that when the line is not transmitting data or transmitting the Stop bit, its value is '1'.

[Figure 23-35](#) illustrates the SmartCard protocol.

Figure 23-35. SmartCard Example



The communication baud rate for ISO7816 is given as:

$$\text{Baud rate} = f_{7816} \times (D/F)$$

Where f_{7816} is the clock frequency, F is the clock rate conversion integer, and D is the baud rate adjustment integer.

By default, F = 372, D = 1, and maximum clock frequency is 5 MHz. Thus, maximum baud rate is 13.4 kbps. Typically, a 3.57-MHz clock is selected; the baud rate will then be 9.6 kbps.

Configuring SCB as UART SmartCard Interface

To configure the SCB as a UART SmartCard interface, set various register bits in the following order. For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

1. Configure the SCB as UART interface by writing '10' to the SCBx_CTRL.MODE register.
2. Configure the UART interface to operate as a Smart-Card protocol by writing '01' to the SCBx_UART_CTRL.MODE register.
3. Follow steps 2 to 4 described in [Enabling and Initializing UART on page 307](#).


23.5.3.5 IrDA

The SCB supports the Infrared Data Association (IrDA) protocol for data rates of up to 115.2 kbps using the UART interface. It supports only the basic physical layer of IrDA protocol with rates less than 115.2 kbps. Hence, the system instantiating this block must consider how to implement a complete IrDA communication system with other available system resources.

The IrDA protocol adds a modulation scheme to the UART signaling. At the transmitter, bits are modulated. At the receiver, bits are demodulated. The modulation scheme uses a Return-to-Zero-Inverted (RZI) format. A bit value of '0' is signaled by a short '1' pulse on the line and a bit value of '1' is signaled by holding the line to '0'. For these data rates (≤ 115.2 kbps), the RZI modulation scheme is used and the pulse duration is 3/16 of the bit period. The sampling clock frequency should be set 16 times the selected baud rate, by configuring the SCBx_CTRL.OVS register. The SCBx_UART_RX_CTRL.POLARITY register can invert the incoming UART_RX line signal. In addition, the Traveo II MCU SCB supports a low-power IrDA receiver mode, which allows it to detect pulses with a minimum width of 1.41 μ s.

Different communication speeds under 115.2 kbps can be achieved by configuring the corresponding block clock frequency. Additional allowable rates are 2.4 kbps, 9.6 kbps, 19.2 kbps, 38.4 kbps, and 57.6 kbps. [Figure 23-36](#) shows how a UART transfer is IrDA modulated.

Figure 23-36. IrDA Example



Configuring the SCB as UART IrDA Interface

To configure the SCB as a UART IrDA interface, set various register bits in the following order. For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

1. Configure the SCB as UART interface by writing '10' to the SCBx_CTRL.MODE register.
2. Configure the UART interface to operate as IrDA protocol by writing '10' to the SCBx_UART_CTRL.MODE register.
3. Enable the median filter on the input interface line by writing '1' to SCBx_RX_CTRL.MEDIAN register.
4. Configure the SCB as described in [Enabling and Initializing UART on page 307](#).

23.5.4 Clocking and Oversampling

The UART protocol is implemented using the SCB input clock as an oversampled multiple of the baud rate. For example, to implement a 100-kHz UART, SCB input clock should be set to 1 MHz and the oversample factor set to '10'. The oversampling is set using the SCBx_CTRL.OVS register field. The oversampling value is SCBx_CTRL.OVS + 1. In the UART standard sub-mode (including LIN) and the SmartCard sub-mode, the valid range for the SCBx_CTRL.OVS field is [7, 15].

In the UART transmit IrDA sub-mode, this field indirectly specifies the oversampling. Oversampling determines the interface clock per bit cycle and the width of the pulse. This sub-mode has only one valid SCBx_CTRL.OVS value-16; the pulse width is roughly 3/16 of the bit period (for all bit rates).

In UART receive IrDA sub-mode (1.2, 2.4, 9.6, 19.2, 38.4, 57.6, and 115.2 kbps), this field indirectly specifies oversampling. In normal transmission mode, this pulse is approximately 3/16 of the bit period (for all bit rates). In low-power transmission mode, this pulse is potentially smaller (down to 1.62 μ s typical and 1.41 μ s minimal) than 3/16 of the bit period (for less than 115.2 kbps bit rates).

Pulse widths greater or equal than two SCB input clock cycles are guaranteed to be detected by the receiver. Pulse widths less than two input clock cycles and greater or equal than one SCB input clock cycle may be detected by the receiver. Pulse widths less than one SCB input clock cycle will not be detected by the receiver. Note that the SCBx_RX_CTRL.MEDIAN should be set to '1' for IrDA receiver functionality.

The SCB input clock and the oversampling together determine the IrDA bit rate. Refer to the *Traveo II Body Controller High Registers TRM* for more details on the SCBx_CTRL.OVS values for different baud rates.

23.5.5 Loop-back

SCB supports internal loop-back from an output signal for UART_TX and UART_RTS to an input signal for UART_RX and UART_CTS without affecting the information on the pins. It is configured using the SCBx_UART_CTRL.LOOPBACK register.

23.5.6 Enabling and Initializing UART

The UART must be programmed in the following order:

1. Program protocol specific information using the SCBx_UART_TX_CTRL, SCBx_UART_RX_CTRL, and SCBx_UART_FLOW_CTRL registers. This includes selecting the submodes of the protocol, transmitter-receiver functionality, and so on.
2. Program the generic transmitter and receiver information using the SCBx_TX_CTRL and SCBx_RX_CTRL registers.
 - a. Specify the data frame width.
 - b. Specify whether MSb or LSb is the first bit to be transmitted or received.
3. Program the transmitter and receiver FIFOs using the SCBx_TX_FIFO_CTRL and SCBx_RX_FIFO_CTRL registers respectively.
 - a. Set the trigger level (TRIGGER_LEVEL).
 - b. Clear the transmitter and receiver FIFO and Shift registers (CLEAR).
4. Enable the block (write a '1' to the SCBx_CTRL.ENABLED register). After the block is enabled, control bits should not be changed. Changes should be made after disabling the block; for example, to modify the operation mode (from SmartCard to IrDA). The change takes effect only after the block is re-enabled. Note that re-enabling the block causes re-initialization and the associated state is lost (such as FIFO content).

23.5.7 I/O Pad Connection

23.5.7.1 Standard UART Mode

Figure 23-37, Figure 23-38, Figure 23-39 and Table 23-6 list the use of the I/O pads for the standard UART mode.

Figure 23-37. Standard UART Mode, I/O Pad Connections




Figure 23-38. Standard UART Mode, Flow Control I/O Pad Connection




Figure 23-39. Standard UART Mode, CTS Reused as TX_EN for RS485




Table 23-6. UART I/O Pad Connection Usage

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Strong (CMOS output)	uart_tx_out_en uart_tx_out	Transmit a data element
uart_rx	Input only	uart_rx_in	Receive a data element
uart_cts	Input only	Uart_cts_in	Indicate peer part readiness to receive data
uart_rts	Strong (CMOS output)	Uart_rts_out_en Uart_rts_out	Indicate DUT readiness to receive data
TX_EN (uart_cts)	Strong (CMOS output)	Uart_cts_out_en Uart_cts_out	Indicate DUT is transmitting data

23.5.7.2 SmartCard Mode

Figure 23-40 and Table 23-7 list the use of the I/O pads for the SmartCard mode.

Figure 23-40. SmartCard Mode I/O Pad Connections




Table 23-7. SmartCard Mode I/O Pad Connections

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Open drain drives low	uart_tx_in	Used to receive a data element. Receive a negative acknowledgment of a transmitted data element
		uart_tx_out_en uart_tx_out	Transmit a data element. Transmit a negative acknowledgment to a received data element.

23.5.7.3 LIN Mode

Figure 23-41 and Table 23-8 list the use of the I/O pads for LIN mode.

Figure 23-41. LIN Mode I/O Pad Connections




Table 23-8. LIN Mode I/O Pad Connections

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Strong (CMOS output)	uart_tx_out_en uart_tx_out	Transmit a data element.
uart_rx	Input only	uart_rx_in	Receive a data element.

23.5.7.4 IrDA Mode

Figure 23-42 and Table 23-9 list the use of the I/O pads for IrDA mode.

Figure 23-42. IrDA Mode I/O Pad Connections




Table 23-9. IrDA Mode I/O Pad Connections

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Strong (CMOS output)	uart_tx_out_en uart_tx_out	Transmit a data element.
uart_rx	Input only	uart_rx_in	Receive a data element.

23.5.8 UART Registers

The UART interface is controlled using a set of 32-bit registers listed in Table 23-20. For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

23.6 Inter Integrated Circuit (I²C)

This section explains the I²C implementation in the Traveo II MCUs. For more information on the I²C protocol specification, refer to the I²C-bus specification available on the [NXP website](#). In the Traveo II MCU, all SCB blocks support both I²C master and slave mode; only one SCB (SCB[0]) is available in DeepSleep power mode and allows externally-clocked operations.

23.6.1 Features

This block supports the following features:


- Master, slave, and master/slave mode
- Standard-mode (100 kbps), fast-mode (400 kbps), and fast-mode plus (1000 kbps) data-rates
- 7-bit slave addressing

- Clock stretching
- Collision detection
- Programmable oversampling of I²C clock signal (SCL)
- Auto ACK when RX FIFO not full, including address
- General address detection
- FIFO Mode
- EZ and CMD_RESP modes
- Interrupts or polling CPU interface
- Analog glitch filter
- Local loop-back control

23.6.2 General Description

Figure 23-43 illustrates an example of an I²C communication network.

Figure 23-43. I²C Interface Block Diagram



The standard I²C bus is a two-wire interface with the following lines:

- Serial Data (SDA)
- Serial Clock (SCL)

I²C devices are connected to these lines using open collector or open-drain output stages, with pull-up resistors (Rp). A simple master/slave relationship exists between

devices. Masters and slaves can operate as either transmitter or receiver. Each slave device connected to the bus is software addressable by a unique 7-bit address.

23.6.3 Terms and Definitions

Table 23-10 explains the commonly used terms in an I²C communication network.

Table 23-10. Definition of I²C Bus Terminology

Term	Description
Transmitter	The device that sends data to the bus
Receiver	The device that receives data from the bus
Master	The device that initiates a transfer, generates clock signals, and terminates a transfer
Slave	The device addressed by a master
Multi-master	More than one master can attempt to control the bus at the same time
Arbitration	Procedure to ensure that, if more than one master simultaneously tries to control the bus, only one is allowed to do so and the winning message is not corrupted
Synchronization	Procedure to synchronize the clock signals of two or more devices

23.6.3.1 Clock Stretching

When a slave device is not yet ready to process data, it may drive a '0' on the SCL line to hold it down. Due to the implementation of the I/O signal interface, the SCL line value will be '0', independent of the values that any other master or slave may be driving on the SCL line. This is known as clock stretching and is the only situation in which a slave drives the SCL line. The master device monitors the SCL line and detects it when it cannot generate a positive clock pulse ('1') on the SCL line. It then reacts by delaying the generation of a positive edge on the SCL line, effectively synchronizing with the slave device that is stretching the clock. The SCB on the Traveo II MCU can and will stretch the clock.

23.6.3.2 Bus Arbitration

The I²C protocol is a multi-master, multi-slave interface. Bus arbitration is implemented on master devices by monitoring the SDA line. Bus collisions are detected when the master observes an SDA line value that is not the same as the value it is driving on the SDA line. For example, when master 1 is driving the value '1' on the SDA line and master 2 is driving the value '0' on the SDA line, the actual line value will be '0' due to the implementation of the I/O signal interface. Master 1 detects the inconsistency and loses control of the bus. Master 2 does not detect any inconsistency and keeps control of the bus.

23.6.4 I²C Modes of Operation

I²C is a synchronous single master, multi-master, multi-slave serial interface. Devices operate in either master mode, slave mode, or master/slave mode. In master/slave mode, the device switches from master to slave mode when it is addressed. Only a single master may be active during a data transfer. The active master is responsible for driving the clock on the SCL line. [Table 23-11](#) illustrates the I²C modes of operation.

Table 23-11. I²C Modes

Mode	Description
Slave	Slave only operation (default)
Master	Master only operation
Multi-master	Supports more than one master on the bus

Data transfer through the I²C bus follows a specific format. [Table 23-12](#) lists some common bus events that are part of

an I²C data transfer. The [Write Transfer](#) and [Read Transfer](#) sections explain the I²C bus bit format during data transfer.

Table 23-12. I²C Bus Events Terminology

Bus Event	Description
START	A HIGH to LOW transition on the SDA line while SCL is HIGH
STOP	A LOW to HIGH transition on the SDA line while SCL is HIGH
ACK	The receiver pulls the SDA line LOW and it remains LOW during the HIGH period of the clock pulse, after the transmitter transmits each byte. This indicates to the transmitter that the receiver received the byte properly.
NACK	The receiver does not pull the SDA line LOW and it remains HIGH during the HIGH period of clock pulse after the transmitter transmits each byte. This indicates to the transmitter that the receiver received the byte unsuccessfully.
Repeated START	START condition generated by master at the end of a transfer instead of a STOP condition
DATA	SDA status change while SCL is low (data changing), and no change while SCL is high (data valid)

With all of these modes, there are two types of transfer-read and write. In write transfer, the master sends data to slave; in read transfer, the master receives data from slave.

Above START, STOP, ACK, NACK, and Repeated START is controlled by the following registers. For more information, see the *Traveo II Body Controller High Registers TRM*.


- SCBx_I2C_M_CMD.M_START
- SCBx_I2C_M_CMD.M_START_ON_IDLE
- SCBx_I2C_M_CMD.M_ACK
- SCBx_I2C_M_CMD.M_NACK
- SCBx_I2C_M_CMD.M_STOP
- SCBx_I2C_S_CMD.S_ACK
- SCBx_I2C_S_CMD.S_NACK

The behavior when received ACK or NACK can be configured by the following registers. For more information, see the *Traveo II Body Controller High Registers TRM*.

- SCBx_I2C_CTRL.M_READY_DATA_ACK
- SCBx_I2C_CTRL.M_NOT_READY_DATA_NACK
- SCBx_I2C_CTRL.S_GENERAL_IGNORE
- SCBx_I2C_CTRL.S_READY_ADDR_ACK
- SCBx_I2C_CTRL.S_READY_DATA_ACK
- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK
- SCBx_I2C_CTRL.S_NOT_READY_DATA_NACK

23.6.4.1 Write Transfer


Figure 23-44. Master Write Data Transfer



- A typical write transfer begins with the master generating a START condition on the I²C bus. The master then writes a 7-bit I²C slave address and a write indicator ('0') after the START condition. The addressed slave transmits an acknowledgment byte by pulling the data line low during the ninth bit time.
- If the slave address does not match any of the slave devices or if the addressed device does not want to acknowledge the request, it transmits a no acknowledgment (NACK) by not pulling the SDA line low. The absence of an acknowledgment, results in an SDA line value of '1' due to the pull-up resistor implementation.
- If no acknowledgment is transmitted by the slave, the master may end the write transfer with a STOP event.
- The master can also generate a repeated START condition for a retry attempt.
- The master may transmit data to the bus if it receives an acknowledgment. The addressed slave transmits an acknowledgment to confirm the receipt of every byte of data written. Upon receipt of this acknowledgment, the master may transmit another data byte.
- When the transfer is complete, the master generates a STOP condition.
- Individual data transfers (of one or more data elements) start with a START event and end with a STOP event. Combined data transfers consist of multiple individual transfers that are not separated by STOP events, but by repeated START events only.

23.6.4.2 Read Transfer

Figure 23-45. Master Read Data Transfer



- A typical read transfer begins with the master generating a START condition on the I²C bus. The master then writes a 7-bit I²C slave address and a read indicator ('1') after the START condition. The addressed slave transmits an acknowledgment by pulling the data line low during the ninth bit time.
- If the slave address does not match with that of the connected slave device or if the addressed device does not want to acknowledge the request, a no acknowledgment (NACK) is transmitted by not pulling the SDA line low. The absence of an acknowledgment, results in an SDA line value of '1' due to the pull-up resistor implementation.
- If no acknowledgment is transmitted by the slave, the master may end the read transfer with a STOP event. The master can also generate a repeated START condition for a retry attempt.
- If the slave acknowledges the address, it starts transmitting data after the acknowledgment signal. The master transmits an acknowledgment to confirm the receipt of each data byte sent by the slave. Upon receipt of this acknowledgment, the addressed slave may transmit another data byte.
- The master can send a NACK signal to the slave to stop the slave from sending data bytes. This completes the read transfer.

- When the transfer is complete, the master generates a STOP condition.
- Individual data transfers (of one or more data elements) start with a START event and end with a STOP event. Combined data transfers consist of multiple individual transfers that are not separated by STOP events, but by repeated START events only.

23.6.5 I²C Buffer Modes

I²C can operate in three different buffered modes - FIFO, EZ, and CMD_RESP modes. The buffer is used in different ways in each of the modes. The following subsections explain each of these buffered modes in detail.

23.6.5.1 FIFO Mode

The FIFO mode has a TX FIFO for the data being transmitted and an RX FIFO for the data being received. Each FIFO is constructed out of the SRAM buffer. The FIFOs are either 32 elements deep with 32-bit data elements or 64 elements deep with 16-bit data elements or 128 elements deep with 8-bit data elements. The width of the data elements are configured using the SCBx_CTRL.MEM_WIDTH. For I²C it is recommended to put the FIFO in BYTE mode because all transactions are a byte wide.

The FIFO mode operation is available only in Active and Sleep power modes, not in the DeepSleep power mode. However, on the DeepSleep-capable SCB the slave address can be used to wake the device from sleep.

Transmit and receive FIFOs allow write and read accesses. A write access to the transmit FIFO uses register SCBx_TX_FIFO_WR. A read access from the receive FIFO uses register SCBx_RX_FIFO_RD.

Transmit and receive FIFO status information is available through the SCBx_TX_FIFO_STATUS and SCBx_RX_FIFO_STATUS registers. It is possible to define a programmable threshold that indicates a number of FIFO entries, a trigger/event is generated when the following conditions are met:

- The transmit FIFO has a SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL. A trigger/event is generated when number of entries in the transmit FIFO is less than SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL.
- The receive FIFO has an SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL. A trigger/event is generated when number of receive FIFO entries is greater than the SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL.

Furthermore, several interrupt status bits are provided as well, which indicate if the FIFOs are full, empty, and so on.

DeepSleep to Active Transition

SCBx_CTRL.EC_AM_MODE = 1,
 SCBx_CTRL.EC_OP_MODE = 0, FIFO Mode.

Master Write:

- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK = 0, SCBx_I2C_CTRL.S_READY_ADDR_ACK = 1. The clock is stretched until the internally-clocked logic takes over, at which point the address is ACK'd and the master can start writing data. Before going to DeepSleep, CLK_SCB needs to be disabled. Upon wake up from DeepSleep CLK_SCB must be re-enabled; this is when the clock stretch will be released.
- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK = 0, SCBx_I2C_CTRL.S_READY_ADDR_ACK = 0. The clock is stretched until the internally-clocked logic takes over and the CPU writes either SCBx_I2C_S_CMD.S_ACK, or SCBx_I2C_S_CMD.S_NACK. Before going to DeepSleep CLK_SCB needs to be disabled. Upon wake up from DeepSleep CLK_SCB must be re-enabled, do this before setting SCBx_I2C_S_CMD.S_ACK or SCBx_I2C_S_CMD.S_NACK.
- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK = 1, SCBx_I2C_CTRL.S_READY_ADDR_ACK = x. The incoming address is NACK'd until the internally-clocked logic takes over. When the internally-clocked logic takes over, there is no guarantee that the internal clock will be at the correct frequency due to PLL/FLL locking times. This may lead to incorrect timing on the I²C bus for the ACK/NACK. To avoid this disable CLK_SCB before going to deep sleep, and then re-enable after the PLL/FLL have stabilized.

Master Read:

- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK = 0, SCBx_I2C_CTRL.S_READY_ADDR_ACK = x. The incoming address is stretched until the internally-clocked logic takes over and the CPU writes data into the TX FIFO. Before going to DeepSleep CLK_SCB needs to be disabled. Upon wake up from DeepSleep CLK_SCB must be re-enabled before writing data into the TX FIFO.
- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK = 1, SCBx_I2C_CTRL.S_READY_ADDR_ACK = x. The incoming address is NACK'd until the internally-clocked logic takes over. When this happens, there is no guarantee that the internal clock will be at the correct frequency due to PLL/FLL locking times. This may lead to incorrect timing on the I²C bus for the ACK/NACK. To avoid this, disable CLK_SCB before going to deep sleep, and then re-enable after the PLL/FLL have stabilized.

23.6.5.2 EZI2C Mode

The Easy I²C (EZI2C) protocol is a unique communication scheme built on top of the I²C protocol by Cypress. It uses a meta protocol around the standard I²C protocol to

communicate to an I²C slave using indexed memory transfers. This removes the need for CPU intervention.

The EZI²C protocol defines a single memory buffer with an 8-bit address that indexes the buffer (256-entry array of 8-bit per entry is supported) located on the slave device. The EZ address is used to address these 256 locations. The CPU writes and reads to the memory buffer through the EZ_DATA registers. These accesses are word accesses, but only the least significant byte of the word is used.

The slave interface accesses the memory buffer using the current address. At the start of a transfer (I²C START/RESTART), the base address is copied to the current address. A data element write or read operation is to the current address location. After the access, the current address is incremented by '1'.

If the current address equals the last memory buffer address (255), the current address is not incremented. Subsequent write accesses will overwrite any previously written value at the last buffer address. Subsequent read accesses will continue to provide the (same) read value at the last buffer address. The bus master should be aware of the memory buffer capacity in EZ mode.

The I²C base and current addresses are provided through I2C_STATUS. At the end of a transfer (I²C), the difference


between the base and current addresses indicates how many read or write accesses were performed. The block provides interrupt cause fields to identify the end of a transfer. EZ mode operation is available in Active, Sleep, and DeepSleep power modes. In Traveo II MCUs, only the DeepSleep-capable SCB block operate in EZI²C mode.

EZI²C distinguishes three operation phases:

- Address phase: The master transmits an 8-bit address to the slave. This address is used as the slave base and current address.
- Write phase: The master writes 8-bit data element(s) to the slave's memory buffer. The slave's current address is set to the slave's base address. Received data elements are written to the current address memory location. After each memory write, the current address is incremented.
- Read phase: The master reads 8-bit data elements from the slave's memory buffer. The slave's current address is set to the slave's base address. Transmitted data elements are read from the current address memory location. After each memory read, the current address is incremented.

Note that a slave's base address is updated by the master and not by the CPU.

Figure 23-46. EZI²C Write and Read Data Transfer



DeepSleep to Active Transition

SCBx_CTRL.EC_AM_MODE = 1,
 SCBx_CTRL.EC_OP_MODE = 0, EZ Mode.

- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK = 0, SCBx_I2C_CTRL.S_READY_ADDR_ACK = 1. The clock is stretched until the internally-clocked logic takes over at which point the address is ACK'd and master can start writing data. Before going to DeepSleep CLK_SCB needs to be disabled. Upon wake up from DeepSleep CLK_SCB must be re-enabled this is when the clock stretch will be released.
- SCBx_I2C_CTRL.S_NOT_READY_ADDR_NACK = 1, SCBx_I2C_CTRL.S_READY_ADDR_ACK = x. The incoming address is NACK'd until the internally-clocked logic takes over. When this happens, there is no guarantee that the internal clock will be at the correct frequency due to PLL/FLL locking times. To avoid this, disable CLK_SCB before going to deep sleep, and then re-enable after the PLL/FLL have stabilized.

23.6.5.3 Command-Response Mode

In the Traveo II MCU, only the DeepSleep-capable SCB supports the command-response mode. This mode has a single memory buffer, a base read address, a current read address, a base write address, and a current write address that are used to index the memory buffer. The base addresses are provided by the CPU. The current addresses are used by the slave to index the memory buffer for sequential accesses of the memory buffer. The memory buffer holds 256 8-bit data elements. The base and current addresses are in the range [0 to 255].

The CPU writes and reads to the memory buffer through the SCBx_EZ_DATA registers. These are word accesses, but only the least significant byte of the word is used.


The slave interface accesses the memory buffer using the current addresses. At the start of a write transfer (I2C START/RESTART), the base write address is copied to the current write address. A data element write is to the current write address location. After the write access, the current address is incremented by '1'. At the start of a read transfer, the base read address is copied to the current read address. A data element read is to the current read address location. After the read data element is transmitted, the current read address is incremented by '1'.

If the current addresses equal the last memory buffer address (255), the current addresses are not incremented. Subsequent write accesses will overwrite any previously written value at the last buffer address. Subsequent read accesses will continue to provide the (same) read value at the last buffer address. The bus master should be aware of the memory buffer capacity in command-response mode.

The base addresses are provided through SCBx_CMD_RESP_CTRL.BASE_RD_ADDR and SCBx_CMD_RESP_CTRL.BASE_WR_ADDR. The current

addresses are provided through SCBx_CMD_RESP_STATUS.CURR_RD_ADDR and SCBx_CMD_RESP_STATUS.CURR_WR_ADDR. At the end of a transfer (I²CSTOP), the difference between a base and current address indicates how many read/write accesses were performed. The block provides interrupt cause fields to identify the end of a transfer. Command-response mode operation is available in Active, Sleep, and DeepSleep power modes. The command-response mode has two phases of operation:

- Write phase - The write phase begins with a START/RESTART followed by the slave address with read/write bit set to '0' indicating a write. The slave's current write address is set to the slave's base write address. Received data elements are written to the current write address memory location. After each memory write, the current write address is incremented.
- Read phase - The read phase begins with a START/RESTART followed by the slave address with read/write bit set to '1' indicating a read. The slave's current read address is set to the slave's base read address. Transmitted data elements are read from the current address memory location. After each read data element is transferred, the current read address is incremented.

Figure 23-47. I²C Command-Response Mode


Note that a slave's base addresses are updated by the CPU and not by the master.

23.6.6 Clocking and Oversampling

The SCB I²C supports both internally and externally-clocked operation modes. SCBx_CTRL.EC_AM_MODE and SCBx_CTRL.EC_OP_MODE register determine the SCB clock mode. SCBx_CTRL.EC_AM_MODE indicates whether I²C address matching is clocked internally (0) or externally (1). I²C address matching comprises the first part of the I²C protocol. SCBx_CTRL.EC_OP_MODE indicates whether the rest of the protocol operation (besides I²C address matching) is clocked internally (0) or externally (1). The externally-clocked mode of operation is supported only in the I²C slave mode.

An internally-clocked operation uses the programmable clock dividers. For more information on system clocking, see the [Clocking System chapter on page 182](#). The internally-clocked mode does not support the command-response mode.

Note: In the Traveo II MCUs, only one SCB supports externally-clocked mode of operation.

The SCBx_CTRL.EC_AM_MODE and SCBx_CTRL.EC_OP_MODE can be configured in the following ways.

- SCBx_CTRL.EC_AM_MODE is '0' and SCBx_CTRL.EC_OP_MODE is '0': Use this configuration when only Active mode functionality is required.
 - FIFO mode: Supported.
 - EZ mode: Supported.
 - Command-response mode: Not supported. The slave NACKs every slave address.
- SCBx_CTRL.EC_AM_MODE is '1' and SCBx_CTRL.EC_OP_MODE is '0': Use this

configuration when both Active and DeepSleep functionality are required. This configuration relies on the externally-clocked functionality for the I²C address matching and relies on the internally-clocked functionality to access the memory buffer. The "handover" from external to internal functionality relies either on an ACK/NACK or clock stretching scheme. The former may result in termination of the current transfer and relies on a master retry. The latter stretches the current transfer after a matching address is received. This mode requires the master to support either NACK generation (and retry) or clock stretching. When the I²C address is matched, SCBx_INTR_I2C_EC.WAKE_UP is set to '1'. The associated DeepSleep functionality interrupt brings the system into Active power mode.

- FIFO mode: See [DeepSleep to Active Transition on page 314](#).
- EZ mode: See [DeepSleep to Active Transition on page 316](#).
- CMD_RESP mode: Not supported. The slave NACKs every slave address.

- SCBx_CTRL.EC_AM_MODE is '1' and SCBx_CTRL.EC_OP_MODE is '1'. Use this mode when both Active and DeepSleep functionality are required. This mode may cause a "denial of service" for memory buffer accesses made by the CPU. When the slave is selected, SCBx_INTR_I2C_EC.WAKE_UP is set to '1'. The associated DeepSleep functionality interrupt brings the system into Active power mode. When the slave is deselected, SCBx_INTR_I2C_EC.EZ_STOP and SCBx_INTR_I2C_EC.EZ_WRITE_STOP are set to '1'.
 - FIFO mode: Not supported.
 - EZ mode: Supported.
 - CMD_RESP mode: Supported.

Table 23-13. Clock Configuration and Mode support

Mode	SCBx_CTRL.EC_AM_MODE is '0'; SCBx_CTRL.EC_OP_MODE is '0'	SCBx_CTRL.EC_AM_MODE is '1'; SCBx_CTRL.EC_OP_MODE is '0'	SCBx_CTRL.EC_AM_MODE is '1'; SCBx_CTRL.EC_OP_MODE is '1'
FIFO mode	Yes	Yes	No
EZ mode	Yes	Yes	Yes
CMD_RESP mode	No	No	Yes

An externally-clocked operation uses a clock provided by the serial interface. The externally-clocked mode does not support FIFO mode. If SCBx_CTRL.EC_OP_MODE is '1', the external interface logic accesses the memory buffer on the external interface clock (I²C SCL). This allows for EZ and CMD_RESP mode functionality in Active and DeepSleep power modes.

In Active system power mode, the memory buffer requires arbitration between external interface logic (on I²C SCL) and the CPU interface logic (on system peripheral clock). This arbitration always gives the highest priority to the external interface logic (host accesses). The external interface logic takes one serial interface clock/bit periods for the I²C. During this period, the internal logic is denied service to the memory buffer. The Traveo II MCU provides two programmable options to address this “denial of service”:

- If the SCBx_CTRL.BLOCK is '1': An internal logic access to the memory buffer is blocked until the memory buffer is granted and the external interface logic has completed access. For a 100-kHz I²C interface, the maximum blocking period of one serial interface bit period measures 10 μ s (approximately 208 clock cycles on a 48 MHz SCB input clock). This option provides normal SCB register functionality, but the blocking time introduces additional internal bus wait states.
- If the SCBx_CTRL.BLOCK is '0': An internal logic access to the memory buffer is not blocked, but fails when it conflicts with an external interface logic access. A read access returns the value 0xFFFF:FFFF and a write access is ignored. This option does not introduce additional internal bus wait states, but an access to the memory buffer may not take effect. In this case, following failures are detected:
 - Read Failure: A read failure is easily detected, as the returned value is 0xFFFF:FFFF. This value is unique as non-failing memory buffer read accesses return an unsigned byte value in the range 0x0000:0000-0x0000:00FF.
 - Write Failure: A write failure is detected by reading back the written memory buffer location, and confirming that the read value is the same as the written value.

For both options, a conflicting internal logic access to the memory buffer sets SCBx_INTR_TX.BLOCKED field to '1' (for write accesses) and SCBx_INTR_RX.BLOCKED field to '1' (for read accesses). These fields can be used as either

status fields or as interrupt cause fields (when their associated mask fields are enabled).


If a series of read or write accesses is performed and SCBx_CTRL.BLOCK is '0', a failure is detected by comparing the “logical-or” of all read values to 0xFFFF:FFFF and checking the SCBx_INTR_TX.BLOCKED and SCBx_INTR_RX.BLOCKED fields to determine whether a failure occurred for a (series of) write or read operation(s).

Note: In Traveo II MCUs, only one SCB supports externally-clocked mode of operation.

23.6.6.1 Glitch Filtering

The Traveo II MCU SCB I²C has analog and digital glitch filters. Analog glitch filters are applied on the i2c_scl_in and i2c_sda_in input signals (AF_in) to filter glitches of up to 50 ns. An analog glitch filter is also applied on the i2c_sda_out output signal (AF_out). Analog glitch filters are enabled and disabled in the SCBx_I2C_CFG register. Do not change the _TRIM bitfields, only change the _SEL bitfields in this register.

Digital glitch filters are applied on the i2c_scl_in and i2c_sda_in input signals (DF_in). The digital glitch filter is enabled in the SCBx_RX_CTRL.MEDIAN.

Figure 23-48. I²C Glitch Filtering Connection


The following table lists the useful combinations of glitch filters.

Table 23-14. Glitch Filter Combinations

AF_in	AF_out	DF_in	Comments
0	0	1	Used when operating in internally-clocked mode and in Master in Fast-mode plus (1-MHz speed mode)
1	0	0	Used when operating in internally-clocked mode (SCBx_CTRL.EC_OP_MODE is '0')
1	1	0	Used when operating in externally-clocked mode (SCBx_CTRL.EC_OP_MODE is '1'). Only slave mode.

When operating in EC_OP_MODE = 1, the 100-kHz, 400-kHz, and 1000-kHz modes require the following settings for AF_out:

AF_in	AF_out	DF_in	
1	1	0	100-kHz mode: SCBx_I2C_CFG.SDA_OUT_FILT_SEL = 3 400-kHz mode: SCBx_I2C_CFG.SDA_OUT_FILT_SEL = 3 1000-kHz mode: SCBx_I2C_CFG.SDA_OUT_FILT_SEL = 1

23.6.6.2 Oversampling and Bit Rate

Internally-clocked Master

The Traveo II MCU implements the I²C clock as an oversampled multiple of the SCB input clock. In master mode, the block determines the I²C frequency. Routing delays on the PCB, on the device, and the block (including analog and digital glitch filters) all contribute to the signal interface timing. In master mode, the block operates off CLK_SCB and uses programmable oversampling factors for the SCL high SCBx_I2C_CTRL.HIGH_PHASE_OVS and low SCBx_I2C_CTRL.LOW_PHASE_OVS times.

 Table 23-15. I²C Frequency and Oversampling Requirements in I²C Master Mode

AF_in	AF_out	DF_in	Mode	Supported Frequency	SCBx_I2C_CTRL.LOW_PHASE_OVS	SCBx_I2C_CTRL.HIGH_PHASE_OVS	Input Clock Frequency
0	0	1	100 kHz	[62, 100] kHz	[9, 15]	[9, 15]	[1.98-3.2] MHz
			400 kHz	[264, 400] kHz	[13, 15]	[7, 15]	[8.45-10] MHz
			1000 kHz	[447, 1000] kHz	[8, 15]	[5, 15]	[14.32-25.8] MHz
1	0	0	100 kHz	[48, 100] kHz	[7, 15]	[7, 15]	[1.55-3.2] MHz
			400 kHz	[244, 400] kHz	[12, 15]	[7, 15]	[7.82-10] MHz
			1000 kHz	Not supported			

Table 23-15 assumes worst-case conditions on the I²C bus. The following equations can be used to determine the settings for your own system. This will involve measuring the rise and fall times on SCL and SDA lines in your system.

$$t_{CLK\_SCB}(\text{Min}) = (t_{LOW} + t_F) / SCBx\_I2C\_CTRL.LOW\_PHASE\_OVS$$

If CLK_SCB is any faster than this, the t_{LOW} of the I²C specification will be violated. t_F needs to be measured in your system.

$$t_{CLK\_SCB}(\text{Max}) = (t_{VD} - t_{RF} - 100 \text{ nsec}) / 3 \text{ (When analog filter is enabled and digital filter disabled)}$$

$$t_{CLK\_SCB}(\text{Max}) = (t_{VD} - t_{RF}) / 4 \text{ (When analog filter is disabled and digital filter is enabled)}$$

t_{RF} is the maximum of either the rise or fall time. If CLK_SCB is slower than this frequency, t_{VD} will be violated.

Internal-clocked Slave

In slave mode, the I²C frequency is determined by the incoming I²C SCL signal. To ensure proper operation, CLK_SCB must be significantly higher than the I²C bus frequency. Unlike master mode, this mode does not use programmable oversampling factors. Table 23-16 assumes worst-case conditions on the I²C bus including the chip internal delay.

Table 23-16. SCB Input Clock Requirements in I²C Slave Mode

AF_in	AF_out	DF_in	Mode	CLK_SCB Frequency Range
0	0	1	100 kHz	[1.98-12.8] MHz
			400 kHz	[8.45-17.14] MHz
			1000 kHz	[14.32-44.77] MHz
1	0	0	100 kHz	[1.55-12.8] MHz
			400 kHz	[7.82-15.38] MHz
			1000 kHz	[15.84-89.0] MHz

$$t_{CLK\_SCB}(\text{Max}) = (t_{VD} - t_{RF} - 100 \text{ nsec}) / 3 \text{ (When analog filter is enabled and digital filter disabled)}$$

$$t_{CLK\_SCB}(\text{Max}) = (t_{VD} - t_{RF}) / 4 \text{ (When analog filter is disabled and digital filter is enabled)}$$

t_{RF} is the maximum of either the rise or fall time. If CLK_SCB is slower than this frequency, t_{VD} will be violated.

The minimum period of CLK_SCB is determined by one of the following equations:

$$t_{CLK\_SCB}(\text{MIN}) = (t_{SU;DAT}(\text{min}) + t_{RF}) / 16$$

or

$$t_{CLK\_SCB}(\text{min}) = (0.6 \times t_F - 50 \text{ nsec}) / 2 \text{ (When analog filter is enabled and digital filter disabled)}$$

$$t_{CLK\_SCB}(\text{min}) = (0.6 \times t_F) / 3 \text{ (When analog filter is disabled and digital filter enabled)}$$

The result that yields the largest period from the two sets of equations above should be used to set the minimum period of CLK_SCB.

Master-Slave

In this mode, when the SCB is acting as a master device, the block determines the I²C frequency. When the SCB is acting as a slave device, the block does not determine the I²C frequency. Instead, the incoming I²C SCL signal does.

To guarantee operation in both master and slave modes, choose clock frequencies that work for both master and slave using the tables above.

23.6.7 Loop-back

In master-slave mode, SCB supports internal SCL and SDA lines are routed internally in the peripheral. As a result, it is unaffected by other I²C devices.

It is configured using the SCBx_I2C_CTRL.LOOPBACK register.

23.6.8 Enabling and Initializing the I²C

The following section describes the method to configure the I²C block for standard (non-EZ) mode and EZI²C mode.

23.6.8.1 Configuring for I²C FIFO Mode

The I²C interface must be programmed in the following order.

1. Program protocol specific information using the SCBx_I2C_CTRL register. This includes selecting master - slave functionality (MASTER_MODE, SLAVE_MODE).
2. Program the generic transmitter and receiver information using the SCBx_TX_CTRL and SCBx_RX_CTRL registers.
 - a. Specify the data frame width (DATA_WIDTH = 7).
 - b. Specify that MSb is the first bit to be transmitted/received (MSB_FIRST = 1).
3. Set the SCBx_CTRL.MEM_WIDTH to '1' to enable the byte mode.

4. Program the transmitter and receiver FIFOs using the SCBx_TX_FIFO_CTRL and SCBx_RX_FIFO_CTRL registers respectively.
 - a. Set the trigger level (TRIGGER_LEVEL).
 - b. Clear the transmitter and receiver FIFO and Shift registers (CLEAR).
5. Program the SCBx_CTRL register to enable the I²C block and select the I²C mode. For a complete description of the I²C registers, see the *Traveo II Body Controller High Registers TRM*.

23.6.8.2 Configuring for EZ and CMD_RESP Modes

To configure the I²C block for EZ and CMD_RESP modes, set the following I²C register bits

- 1a. Select the EZI2C mode by writing '1' to the SCBx_CTRL.EZ_MODE register.
- 1b. Select CMD_RESP mode by writing a 1 to the SCBx_CTRL.CMD_RESP register.
2. Set the S_READY_ADDR_ACK (bit 12) and SCBx_I2C_CTRL.S_READY_DATA_ACK register.

23.6.9 I/O Pad Connections

Figure 23-49. I²C I/O Pad Connections




Table 23-17. I²C I/O Pad Descriptions

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
i2c_scl	Open drain drives low	i2c_scl_in	Receive a clock
		i2c_scl_out	Transmit a clock
i2c_sda	Open drain drives low	i2c_sda_in	Receive data
		i2c_sda_out	Transmit data

23.6.10 I²C Registers

The I²C interface is controlled by reading and writing a set of configuration, control, and status registers, as listed in [Table 23-21](#).

Note: Detailed descriptions of the I²C register bits are available in the *Traveo II Body Controller High Registers TRM*.

23.7 SCB Interrupts

SCB supports interrupt generation on various events. The interrupts generated by the SCB block vary depending on the mode of operation.

Table 23-18. SCB Interrupts

Interrupt	Functionality	Active/ DeepSleep	Registers
interrupt_master	I ² C master and SPI master functionality	Active	SCBx_INTR_M, SCBx_INTR_M_SET, SCBx_INTR_M_MASK, SCBx_INTR_M_MASKED
interrupt_slave	I ² C slave and SPI slave functionality	Active	SCBx_INTR_S, SCBx_INTR_S_SET, SCBx_INTR_S_MASK, SCBx_INTR_S_MASKED
interrupt_tx	UART transmitter and TX FIFO functionality	Active	SCBx_INTR_TX, SCBx_INTR_TX_SET, SCBx_INTR_TX_MASK, SCBx_INTR_TX_MASKED
interrupt_rx	UART receiver and RX FIFO functionality	Active	SCBx_INTR_RX, SCBx_INTR_RX_SET, SCBx_INTR_RX_MASK, SCBx_INTR_RX_MASKED
interrupt_i2c_ec	Externally-clocked I ² C slave functionality	DeepSleep	SCBx_INTR_I2C_EC, SCBx_INTR_I2C_EC_MASK, SCBx_INTR_I2C_EC_MASKED
interrupt_spi_ec	Externally-clocked SPI slave functionality	DeepSleep	SCBx_INTR_SPI_EC, SCBx_INTR_SPI_EC_MASK, SCBx_INTR_SPI_EC_MASKED

The following sections explain the different interrupt sources for each mode of SCB operation.

Note: To avoid being triggered by events from previous transactions, whenever the firmware enables an interrupt mask register bit, it should clear the interrupt request register in advance.

23.7.1 SPI Interrupts

SPI interrupts can be classified as Master interrupts, Slave interrupts, TX interrupts, RX interrupts, and externally-clocked (EC) mode interrupts. Each interrupt output is the logical OR of the group of all possible interrupt sources classified under the section. For example, the TX interrupt output is the logical OR of the group of all possible TX interrupt sources. This signal goes high when any of the enabled TX interrupt sources are true. The SCB also provides an interrupt cause register (SCBx_INTR_CAUSE) that can be used to determine interrupt source. The interrupt registers are cleared by writing '1' to the corresponding bit field. Note that certain interrupt sources are triggered again as long as the condition is met even if the interrupt source was cleared. For example, the TX_FIFO_EMPTY is set as long as the transmit FIFO is empty even if the interrupt source is cleared. For more information on interrupt

registers, see the *Traveo II Body Controller High Registers TRM*. The SPI supports interrupts on the following events:

- SPI Master interrupts (SCBx_INTR_M)
 - SPI master transfer done (SPI_DONE)
- SPI Slave interrupts (SCBx_INTR_S)
 - SPI slave deselected after a write EZSPI transfer occurred (SPI_EZ_WRITE_STOP)
 - SPI slave deselected after any EZSPI transfer occurred (SPI_EZ_STOP)
 - SPI Bus Error – Slave deselected unexpectedly in the SPI transfer. The firmware may decide to clear the TX and RX FIFOs for this error. (SPI_BUS_ERROR)
- SPI TX (SCBx_INTR_TX)
 - TX FIFO has less entries than the value specified by SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL (TRIGGER)
 - TX FIFO is not full (NOT_FULL)
 - TX FIFO is empty (EMPTY)
 - TX FIFO overflow (OVERFLOW)
 - TX FIFO underflow (UNDERFLOW)

■ SPI RX (SCBx_INTR_RX)

- ❑ RX FIFO has more entries than the value specified by SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL (TRIGGER)
- ❑ RX FIFO is not empty (NOT_EMPTY)
- ❑ RX FIFO is full (FULL)
- ❑ RX FIFO overflow (OVERFLOW)
- ❑ RX FIFO underflow (UNDERFLOW)

■ SPI Externally-clocked (SCBx_INTR_SPI_EC)

- ❑ Wake up request on slave select (WAKE_UP)
- ❑ SPI STOP detection at the end of each transfer (EZ_STOP)
- ❑ SPI STOP detection at the end of a write transfer (EZ_WRITE_STOP)
- ❑ SPI STOP detection at the end of a read transfer (EZ_READ_STOP)

■ UART RX (INTR_RX)

- ❑ RX FIFO has more entries than the value specified by SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL (TRIGGER)
- ❑ RX FIFO is not empty (NOT_EMPTY)
- ❑ RX FIFO is full (FULL)
- ❑ RX FIFO overflow (OVERFLOW)
- ❑ RX FIFO underflow (UNDERFLOW)
- ❑ Frame error in received data frame (FRAME_ERROR)
- ❑ Parity error in received data frame (PARITY_ERROR)
- ❑ LIN baud rate detection is completed (BAUD_DETECT)
- ❑ LIN break detection is successful (BREAK_DETECT)

23.7.2 UART Interrupts

UART interrupts can be classified as TX interrupts and RX interrupts. Each interrupt output is the logical OR of the group of all possible interrupt sources classified under the section. For example, the TX interrupt output is the logical OR of the group of all possible TX interrupt sources. This signal goes high when any of the enabled TX interrupt sources are true. The SCB also provides an interrupt cause register (SCBx_INTR_CAUSE) that can be used to determine interrupt source. The interrupt registers are cleared by writing '1' to the corresponding bitfield. Note that certain interrupt sources are triggered again as long as the condition is met even if the interrupt source was cleared. For example, the TX_FIFO_EMPTY is set as long as the transmit FIFO is empty even if the interrupt source is cleared. For more information on interrupt registers, see the *Traveo II Body Controller High Registers TRM*. The UART block generates interrupts on the following events:

■ UART TX (SCBx_INTR_TX)

- ❑ TX FIFO has less entries than the value specified by SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL (TRIGGER)
- ❑ TX FIFO is not full (NOT_FULL)
- ❑ TX FIFO is empty (EMPTY)
- ❑ TX FIFO overflow (OVERFLOW)
- ❑ TX FIFO underflow (UNDERFLOW)
- ❑ TX received a NACK in SmartCard mode (UART_NACK)
- ❑ TX done. This happens when the UART completes transferring all data in the TX FIFO and the last stop field is transmitted (both TX FIFO and transmit shifter register are empty). (UART_DONE)
- ❑ Arbitration lost (in LIN or SmartCard modes) (UART_ARB_LOST)

23.7.3 I²C Interrupts

I²C interrupts can be classified as Master interrupts, Slave Interrupts, TX interrupts, RX interrupts, and Externally-clocked (EC) mode interrupts. Each interrupt output is the logical OR of the group of all possible interrupt sources classified under the section. For example, the TX interrupt output is the logical OR of the group of all possible TX interrupt sources. This signal goes high when any of the enabled TX interrupt sources are true. The SCB also provides an interrupt cause register (SCBx_INTR_CAUSE) that can be used to determine interrupt source. The interrupt registers are cleared by writing '1' to the corresponding bit field. Note that certain interrupt sources are triggered again as long as the condition is met even if the interrupt source was cleared. For example, the TX_FIFO_EMPTY is set as long as the transmit FIFO is empty even if the interrupt source is cleared. For more information on interrupt registers, see the *Traveo II Body Controller High Registers TRM*. The I²C block generates interrupts for the following conditions.

■ I²C Master (SCBx_INTR_M)

- ❑ I²C master lost arbitration (I2C_ARB_LOST)
- ❑ I²C master received NACK (I2C_NACK)
- ❑ I²C master received ACK (I2C_ACK)
- ❑ I²C master sent STOP (I2C_STOP)
- ❑ I²C bus error (unexpected stop/start condition detected) (I2C_BUS_ERROR)

■ I²C Slave (SCBx_INTR_S)

- ❑ I²C slave lost arbitration (I2C_ARB_LOST)
- ❑ I²C slave received NACK (I2C_NACK)
- ❑ I²C slave received ACK (I2C_ACK)

- ❑ I²C slave received Write STOP (I2C_WRITE_STOP)
- ❑ I²C slave received STOP (I2C_STOP)
- ❑ I²C slave received START (I2C_START)
- ❑ I²C slave address matched (I2C_ADDR_MATCH)
- ❑ I²C slave general call address received (I2C_GENERAL)
- ❑ I²C bus error – unexpected stop/start condition detected (I2C_BUS_ERROR)
- I²C TX (SCBx_INTR_TX)
 - ❑ TX FIFO has less entries than the value specified by SCBx_TX_FIFO_CTRL.TRIGGER_LEVEL (TRIGGER)
 - ❑ TX FIFO is not full (NOT_FULL)
 - ❑ TX FIFO is empty (EMPTY)
 - ❑ TX FIFO overflow (OVERFLOW)
 - ❑ TX FIFO underflow (UNDERFLOW)
- I²C RX (SCBx_INTR_RX)
 - ❑ RX FIFO has more entries than the value specified by SCBx_RX_FIFO_CTRL.TRIGGER_LEVEL (TRIGGER)
 - ❑ RX FIFO is not empty (NOT_EMPTY)
 - ❑ RX FIFO is full (FULL)
 - ❑ RX FIFO overflow (OVERFLOW)
 - ❑ RX FIFO underflow (UNDERFLOW)
- I²C Externally-clocked (SCBx_INTR_I2C_EC)
 - ❑ Wake up request on address match (WAKE_UP)
 - ❑ I²C STOP detection at the end of each transfer (EZ_STOP)
 - ❑ I²C STOP detection at the end of a write transfer (EZ_WRITE_STOP)
 - ❑ I²C STOP detection at the end of a read transfer (EZ_READ_STOP)

23.8 Registers

23.8.1 SPI Registers

Table 23-19. SPI Registers

Register	Name	Description
SCBx_CTRL	SCB Control Register	Enables the SCB, selects the type of serial interface (SPI, UART, I ² C), and selects internally and externally-clocked operation, EZ and non-EZ modes of operation.
SCBx_STATUS	SCB Status Register	In EZ mode, this register indicates whether the externally-clocked logic is potentially using the EZ memory.
SCBx_SPI_CTRL	SCB SPI Control Register	Configures the SPI as either a master or a slave, selects SPI protocols (Motorola, TI, National) and clock-based submodes in Motorola SPI (modes 0,1,2,3), selects the type of SELECT signal in TI SPI.
SCBx_SPI_STATUS	SCB SPI Status Register	Indicates whether the SPI bus is busy and sets the SPI slave EZ address in the internally-clocked mode.
SCBx_TX_CTRL	SCB TX Control Register	Specifies the data frame width and specifies whether MSb or LSb is the first bit in transmission.
SCBx_RX_CTRL	SCB RX Control Register	Performs the same function as that of the SCBx_TX_CTRL register, but for the receiver. Also decides whether a median filter is to be used on the input interface lines.
SCBx_TX_FIFO_CTRL	SCB TX FIFO Control Register	Specifies the trigger level, clears the transmitter FIFO and shift registers, and performs the FREEZE operation of the transmitter FIFO.
SCBx_RX_FIFO_CTRL	SCB RX FIFO Control Register	Performs the same function as that of the SCBx_TX_FIFO_CTRL register, but for the receiver.
SCBx_TX_FIFO_WR	SCB TX FIFO Write Register	Holds the data frame written into the transmitter FIFO. Behavior is similar to that of a PUSH operation.
SCBx_RX_FIFO_RD	SCB RX FIFO Read Register	Holds the data frame read from the receiver FIFO. Reading a data frame removes the data frame from the FIFO - behavior is similar to that of a POP operation. This register has a side effect when read by software: a data frame is removed from the FIFO.
SCBx_RX_FIFO_RD_SILENT	SCB RX FIFO Read Silent Register	Holds the data frame read from the receiver FIFO. Reading a data frame does not remove the data frame from the FIFO; behavior is similar to that of a PEEK operation.
SCBx_TX_FIFO_STATUS	SCB TX FIFO Status Register	Indicates the number of bytes stored in the transmitter FIFO, the location from which a data frame is read by the hardware (read pointer), the location from which a new data frame is written (write pointer), and decides if the transmitter FIFO holds the valid data.
SCBx_RX_FIFO_STATUS	SCB RX FIFO Status Register	Performs the same function as that of the SCBx_TX_FIFO_STATUS register, but for the receiver.
SCBx_EZ_DATA	SCB EZ Data Register	Holds the data in EZ memory location

23.8.2 UART Registers

Table 23-20. UART Registers

Register	Name	Description
SCBx_CTRL	SCB Control Register	Enables the SCB; selects the type of serial interface (SPI, UART, I ² C)
SCBx_UART_CTRL	SCB UART Control Register	Used to select the sub-modes of UART (standard UART, SmartCard, IrDA), also used for local loop back control.
SCBx_UART_RX_STATUS	SCB UART RX Status Register	Used to specify the BR_COUNTER value that determines the bit period. This is used to set the accuracy of the SCB clock. This value provides more granularity than the OVS bit in SCBx_CTRL register.
SCBx_UART_TX_CTRL	SCB UART TX Control Register	Used to specify the number of stop bits, enable parity, select the type of parity, and enable retransmission on NACK.
SCBx_UART_RX_CTRL	SCB UART RX Control Register	Performs same function as SCBx_UART_TX_CTRL but is also used for enabling multi processor mode, LIN mode drop on parity error, and drop on frame error.
SCBx_TX_CTRL	SCB TX Control Register	Used to specify the data frame width and to specify whether MSb or LSb is the first bit in transmission.
SCBx_RX_CTRL	SCB RX Control Register	Performs the same function as that of the SCBx_TX_CTRL register, but for the receiver. Also decides whether a median filter is to be used on the input interface lines.
SCBx_UART_FLOW_CONTROL	SCB UART Flow Control Register	Configures flow control for UART transmitter.

23.8.3 I²C Registers

Table 23-21. I²C Registers

Register	Name	Description
SCBx_CTRL	SCB Control Register	Enables the SCB block and selects the type of serial interface (SPI, UART, I ² C). Also used to select internally and externally-clocked operation and EZ and non-EZ modes of operation.
SCBx_I2C_CTRL	SCB I2C Control Register	Selects the mode (master, slave) and sends an ACK or NACK signal based on receiver FIFO status.
SCBx_I2C_STATUS	SCB I2C Status Register	Indicates bus busy status detection, read/write transfer status of the slave/master, and stores the EZ slave address.
SCBx_I2C_M_CMD	SCB I2C Master Command Register	Enables the master to generate START, STOP, and ACK/NACK signals.
SCBx_I2C_S_CMD	SCB I2C Slave Command Register	Enables the slave to generate ACK/NACK signals.
SCBx_STATUS	SCB Status Register	Indicates whether the externally-clocked logic is using the EZ memory. This bit can be used by software to determine whether it is safe to issue a software access to the EZ memory.
SCBx_I2C_CFG	SCB I2C Configuration Register	Configures filters, which remove glitches from the SDA and SCL lines.
SCBx_TX_CTRL	SCB TX Control Register	Specifies the data frame width; also used to specify whether MSb or LSb is the first bit in transmission.
SCBx_TX_FIFO_CTRL	SCB TX FIFO Control Register	Specifies the trigger level, clearing of the transmitter FIFO and shift registers, and FREEZE operation of the transmitter FIFO.
SCBx_TX_FIFO_STATUS	SCB TX FIFO Status Register	Indicates the number of bytes stored in the transmitter FIFO, the location from which a data frame is read by the hardware (read pointer), the location from which a new data frame is written (write pointer), and decides if the transmitter FIFO holds the valid data.

Table 23-21. I²C Registers

Register	Name	Description
SCBx_TX_FIFO_WR	SCB TX FIFO Write Register	Holds the data frame written into the transmitter FIFO. Behavior is similar to that of a PUSH operation.
SCBx_RX_CTRL	SCB RX Control Register	Performs the same function as that of the SCBx_TX_CTRL register, but for the receiver. Also decides whether a median filter is to be used on the input interface lines.
SCBx_RX_FIFO_CTRL	SCB RX FIFO Control Register	Performs the same function as that of the SCBx_TX_FIFO_CTRL register, but for the receiver.
SCBx_RX_FIFO_STATUS	SCB RX FIFO Status Register	Performs the same function as that of the SCBx_TX_FIFO_STATUS register, but for the receiver.
SCBx_RX_FIFO_RD	SCB RX FIFO Read Register	Holds the data read from the receiver FIFO. Reading a data frame removes the data frame from the FIFO; behavior is similar to that of a POP operation. This register has a side effect when read by software: a data frame is removed from the FIFO.
SCBx_RX_FIFO_RD_SILENT	SCB RX FIFO Read Silent Register	Holds the data read from the receiver FIFO. Reading a data frame does not remove the data frame from the FIFO; behavior is similar to that of a PEEK operation.
SCBx_RX_MATCH	SCB RX Match Register	Stores slave device address and is also used as slave device address MASK.
SCBx_EZ_DATA	SCB EZ Data Register	Holds the data in an EZ memory location.

24. CAN FD Controller



24.1 Overview

The CAN FD controller complies with the ISO11898-1 (CAN specification Rev. 2.0 parts A and B). In addition, it supports the Time-Triggered CAN (TTCAN) protocol defined in ISO 11898-4.

All message handling functions are implemented by the RX and TX handlers. The RX handler manages message acceptance filtering, transfer of received messages from the CAN core to a message RAM, and receive message status information. The TX handler transfers transmit messages from the message RAM to the CAN core and provides transmit status information.

Two separate clocks are provided to the CAN FD controller: CAN clock (PCLK_CANFD[x]_CLOCK_CAN[y]) for CAN operation and system clock (CLK_SYS/CLK_GR5) for internal block operation. Acceptance filtering is implemented by a combination of up to 192 filter elements, where each can be configured as a range, as a bit mask, or as a dedicated ID filter.

The CAN FD controller functions only in Active and Sleep power modes. In DeepSleep mode, it is not functional but is fully retained except the Shared Time Stamp (TS) counter. In Hibernate power mode, the controller is neither functional nor retained.

24.1.1 Features

The CAN FD controller has the following features:

- Flexible data-rate (FD) (ISO 11898-1: 2015)
 - Up to 64 data bytes per message
 - Maximum 8 Mbps supported
- Time-Triggered (TT) communication on CAN (ISO 11898-4: 2004)
 - TTCAN protocol level 1 and level 2 completely in hardware
- AUTOSAR support
- Acceptance filtering
- Two configurable receive FIFOs
- Up to 64 dedicated receive buffers
- Up to 32 dedicated transmit buffers
- Configurable transmit FIFO
- Configurable transmit queue
- Configurable transmit event FIFO
- Programmable loop-back test mode
- Power-down support
- Shared message RAM
- ECC protection for message RAM
- Global fault structure to handle ECC errors
- Receive FIFO top pointer logic
 - Enables DMA access on the FIFO
- DMA for debug message and received FIFOs
- Shared time stamp counter

Note: Refer to the device datasheet to find the supported number of M_TTCAN groups, M_TTCAN channels in each group, and total message RAM allocated to each group.


24.1.2 Features Not Supported

- Asynchronous serial communication (ASC)
- Interrupt of Bit Error Corrected (CANFDx_CHy_IR.BEC) in M_TTCAN
 - This bit is fixed at '0'. If this occurs, the fault structures report as correctable ECC error.

24.2 Configuration

24.2.1 Block Diagram

Figure 24-1. M_TTCAN Block Diagram



24.2.2 Dual Clock Sources

Each M_TTCAN channel has two clock inputs: PCLK_CANFD[x]_CLOCK_CAN[y] and CLK_GR5. The PCLK_CANFD[x]_CLOCK_CAN[y] (clk_cclk) is used for the CAN (or CAN FD) operation. The CLK_GR5 (CLK_SYS) is used for everything except CAN operation, such as register accesses and SRAM accesses. For the CAN FD operation, it is recommended to use 20 MHz, 40 MHz, or 80 MHz for PCLK_CANFD[x]_CLOCK_CAN[y] (clk_cclk). CLK_SYS must run equal or faster than clk_cclk.

Each M_TTCAN channel has its own PCLK_CANFD[x]_CLOCK_CAN[y]. This allows channels to have the flexibility to communicate at independent speeds. However, the host clock (CLK_SYS) will be the same for all M_TTCAN channels.

See the [Clocking System chapter on page 182](#) for more details about clock configuration.

24.2.3 Interrupt Lines


The two kind of interrupts from the M_TTCAN group are as follows:

- Interrupt0 and interrupt1 from each M_TTCAN channel within the M_TTCAN group
- Consolidated interrupt0 and consolidated interrupt1 for one M_TTCAN group

Each M_TTCAN channel provides two interrupt lines: interrupt0 and interrupt1. Interrupts from any source within the M_TTCAN channel can be routed either to interrupt0 or interrupt1 by using CANFDx_CHy_ILS and CANFDx_CHy_TTILS registers. By default, all interrupts are routed to interrupt0. By programming Enable Interrupt Line 0 (CANFDx_CHy_ILE.EINT0) and Enable Interrupt Line 1 (CANFDx_CHy_ILE.EINT1), the interrupt lines can be enabled or disabled separately for each interrupt source.

In Traveo II, one device may contain multiple M_TTCAN channels in one M_TTCAN instance. Therefore, Interrupt line 0 and Interrupt line 1 from each M_TTCAN channel are routed to a common interrupt0 and interrupt1. Common interrupt0 and interrupt1 are ORed of all interrupt0 and interrupt1 coming from all present channels within one M_TTCAN group. Interrupt cause registers CANFDx_INTR0_CAUSE and CANFDx_INTR1_CAUSE provide information about the active interrupt causing channel from a particular group.

Figure 24-2. Interrupts in M_TTCAN Group



24.3 Functional Description

24.3.1 Operation Modes

24.3.1.1 Software Initialization

This refers to setting or resetting the initialization bit (CANFDx_CHy_CCCR.INIT). The CANFDx_CHy_CCCR.INIT bit is set

- either by software or hardware reset
- when an uncorrected bit error is detected in message RAM
- by going Bus Off

While the CANFDx_CHy_CCCR.INIT is set:

- message transfer from and to the CAN bus is stopped

- the status of the CAN bus output CANx_y_TX is recessive (high)
- the protocol error counters are unchanged

Setting CANFDx_CHy_CCCR.INIT does not change any configuration register.

Resetting CANFDx_CHy_CCCR.INIT finishes the software initialization. The CAN FD controller then synchronizes itself to the data transfer on the CAN bus by waiting for the occurrence of a sequence of 11 consecutive recessive bits (Bus Idle) before it can take part in bus activities and start the message transfer.

Access/Set/Reset Properties of Registers Affected by Configuration Change Enable (CANFDx_CHy_CCCR.CCE)

Access to the configuration registers is only enabled when both bits CANFDx_CHy_CCCR.INIT and

CANFDx_CHy_CCCR.CCE are set (write-protected). CANFDx_CHy_CCCR.CCE can only be set/reset while CANFDx_CHy_CCCR.INIT is 1. CANFDx_CHy_CCCR.CCE is automatically reset when CANFDx_CHy_CCCR.INIT is reset.

The following registers are reset when CANFDx_CHy_CCCR.CCE is set

- CANFDx_CHy_HPMS - High Priority Message Status
- CANFDx_CHy_RXF0S - RX FIFO 0 Status
- CANFDx_CHy_RXF1S - RX FIFO 1 Status
- CANFDx_CHy_TXFQS - TX FIFO/Queue Status
- CANFDx_CHy_TXBRP - TX Buffer Request Pending
- CANFDx_CHy_TXBTO - TX Buffer Transmission Occurred
- CANFDx_CHy_TXBCF - TX Buffer Cancellation Finished
- CANFDx_CHy_TXEFS - TX Event FIFO Status
- CANFDx_CHy_TTOST - TT Operation Status
- CANFDx_CHy_TTLGT - TT Local and Global Time, only Global Time CANFDx_CHy_TTLGT.GT is reset
- CANFDx_CHy_TTCTC - TT Cycle Time and Count
- CANFDx_CHy_TTCSM - TT Cycle Sync Mark

In addition

- Timeout Counter value (CANFDx_CHy_TOCV.TOC[15:0]) is preset to the value configured by the Timeout Period (CANFDx_CHy_TOCC.TOP[15:0]) when CANFDx_CHy_CCCR.CCE is set.
- State machines of TX and RX handlers are held in idle state while CANFDx_CHy_CCCR.CCE is 1.

The following registers are only writable while CANFDx_CHy_CCCR.CCE is 0.

- CANFDx_CHy_TXBAR - TX Buffer Add Request
- CANFDx_CHy_TXBCR - TX Buffer Cancellation Request

Test Mode Enable (CANFDx_CHy_CCCR.TEST) and Bus Monitoring mode (CANFDx_CHy_CCCR.MON) can only be set by the CPU while CANFDx_CHy_CCCR.INIT is 1 and CANFDx_CHy_CCCR.CCE is 1. Both bits may be reset at any time.

Disable Automatic Retransmission (CANFDx_CHy_CCCR.DAR) can only be set/reset while CANFDx_CHy_CCCR.INIT is 1 and CANFDx_CHy_CCCR.CCE is 1.

Message RAM Initialization

Each message RAM word should be reset by writing 0x00000000 before configuration of the CAN FD controller. This prevents message RAM bit errors when reading uninitialized words, and also avoids unexpected filter element configurations in message RAM.

24.3.1.2 Normal Operation

The M_TTCAN's default operating mode after hardware reset is event-driven CAN communication without time triggers (CANFDx_CHy_TTOCF.OM = 00). Both CANFDx_CHy_CCCR.INIT and CANFDx_CHy_CCCR.CCE must be set before the TT operation mode is changed.

When M_TTCAN is initialized and CANFDx_CHy_CCCR.INIT is reset to zero, M_TTCAN synchronizes itself to the CAN bus and is ready for communication.

After passing the acceptance filtering, received messages including Message ID and DLC are stored into a dedicated RX buffer or into RX FIFO 0 or RX FIFO 1.

For messages to be transmitted, dedicated TX buffers and a TX FIFO/TX queue can be initialized or updated. Automated transmission on reception of remote frames is not implemented.

24.3.1.3 CAN FD Operation

The two variants in CAN FD frame transmission are:

- CAN FD frame without bit rate switching
- CAN FD frame where the control, data, and CRC fields are transmitted with a higher bit rate than the beginning and end of the frame

The previously reserved bit in CAN frames with 11-bit identifiers and 29-bit identifiers will now be decoded as FDF bit.

- FDF = recessive signifies a CAN FD frame
- FDF = dominant signifies a classic CAN frame

In a CAN FD frame, the two bits following FDF, reserved bits, and bit rate switch (BRS) decide whether the bit rate inside the CAN FD frame is switched. A CAN FD bit rate switch signified by res is dominant and BRS is recessive. The coding of res as recessive is reserved for future expansion of the protocol. If the M_TTCAN receives a frame with FDF and res as recessive, it will signal a Protocol Exception Event by setting the CANFDx_CHy_PSR.PXE bit. When Protocol Exception Handling is enabled (CANFDx_CHy_CCCR.PXHD = 0), it causes the operation state to change from Receiver (CANFDx_CHy_PSR.ACT = 10) to Integrating (CANFDx_CHy_PSR.ACT = 00) at the next sample point. If Protocol Exception Handling is disabled (CANFDx_CHy_CCCR.PXHD = 1), the M_TTCAN will treat a recessive res bit as a form error and respond with an error frame.

CAN FD operation is enabled by programming CANFDx_CHy_CCCR.FDOE. If CANFDx_CHy_CCCR.FDOE is '1', transmission and reception of CAN FD frames is enabled. Transmission and reception of classic CAN frames is always possible. Whether a CAN FD frame or a classic CAN frame is transmitted can be configured via the FDF bit in the

respective TX buffer element. With CANFDx_CHy_CCCR.FDOE as '0', received frames are interpreted as classic CAN frames, which leads to the transmission of an error frame when receiving a CAN FD frame. When CAN FD operation is disabled, no CAN FD frames are transmitted even if the FDF bit of a TX buffer element is set. CANFDx_CHy_CCCR.FDOE and CANFDx_CHy_CCCR.BRSE can only be changed while CANFDx_CHy_CCCR.INIT and CANFDx_CHy_CCCR.CCE are both set.

With CANFDx_CHy_CCCR.FDOE as '0', the setting of FDF and BRS is ignored and frames are transmitted in classic CAN format. When CANFDx_CHy_CCCR.FDOE = 1 and CANFDx_CHy_CCCR.BRSE = 0, only FDF of a TX buffer element is evaluated. When CANFDx_CHy_CCCR.FDOE = 1 and CANFDx_CHy_CCCR.BRSE = 1, transmission of CAN FD frames with bit rate switching is enabled. All TX buffer elements with FDF and BRS bits set are transmitted in CAN FD format with bit rate switching.

A mode change during CAN operation is only recommended under the following conditions:

- The failure rate in the CAN FD data phase is significantly higher than in the CAN FD arbitration phase. In this case disable the CAN FD bit rate switching option for transmissions.
- During system startup, all nodes transmit classic CAN messages until it is verified that they can communicate in CAN FD format. If this is true, all nodes switch to CAN FD operation.
- Wake-up messages in CAN partial networking must be transmitted in classic CAN format.
- End-of-line programming occurs in case all nodes are not CAN FD capable. Non-CAN FD nodes are held in Silent mode until programming is completed. Then all nodes switch back to classic CAN communication.

In the CAN FD format, the coding of the DLC differs from the standard CAN format. The DLC codes 0 to 8 have the same coding as in standard CAN, codes 9 to 15, which in standard CAN have a data field of 8 bytes, are coded according to Table 24-1.

Table 24-1. Coding of DLC in CAN FD

DLC	9	10	11	12	13	14	15
Number of Data Bytes	12	16	20	24	32	48	64

In CAN FD frames, the bit timing will be switched inside the frame after the (BRS) bit, if this bit is recessive. Before the BRS bit, in the CAN FD arbitration phase, the nominal CAN bit timing is used as defined by the Nominal Bit Timing and Prescaler Register (CANFDx_CHy_NBTP). In the following CAN FD data phase, the data phase bit timing is used as defined by the Data Bit Timing and Prescaler Register (CANFDx_CHy_DBTP). The bit timing is switched back from

the data phase timing at the CRC delimiter or when an error is detected, whichever occurs first.

The maximum configurable bit rate in the CAN FD data phase depends on the CAN clock frequency (clk_can). For example, with a CAN clock frequency of 20 MHz and the shortest configurable bit time of 4 tq, the bit rate in the data phase is 5 Mbit/s.

In both data frame formats, CAN FD and CAN FD with bit rate switching, the value of the bit ESI (Error Status Indicator) is determined by the transmitter's error state at the start of the transmission. If the transmitter is error passive, ESI is transmitted recessive; otherwise, it is transmitted dominant.

24.3.1.4 Transmitter Delay Compensation

During the data phase of a CAN FD transmission only one node is a transmitter; all others are receivers. The length of the bus line has no impact. When transmitting via pin CANx_y_TX, the M_TTCAN receives the transmitted data from its local CAN transceiver via pin CANx_y_RX. The received data is delayed by the transmitter delay. In case this delay is greater than TSEG1 (time segment before sample point), a bit error is detected. To enable a data phase bit time that is even shorter than the transmitter delay, the delay compensation is introduced. Without transmitter delay compensation, the bit rate in the data phase of a CAN FD frame is limited by the transmitter delay.

Description

The M_TTCAN's protocol unit has implemented a delay compensation mechanism to compensate the transmitter delay. This enables transmission with higher bit rates during the CAN FD data phase, independent of the delay of a specific CAN transceiver.

To check for bit errors during the data phase of transmitting nodes, the delayed transmit data is compared against the received data at the Secondary Sample Point (SSP). If a bit error is detected, the transmitter will react on this bit error at the next following regular sample point. During the arbitration phase the delay compensation is always disabled.

The transmitter delay compensation enables configurations where the data bit time is shorter than the transmitter delay, it is described in detail in the new ISO 11898-1:2015. It is enabled by setting bit CANFDx_CHy_DBTP.TDC.

The received bit is compared against the transmitted bit at the SSP. The SSP position is defined as the sum of the measured delay from the M_TTCAN's transmit output CANx_y_TX through the transceiver to the receive input RX plus the transmitter delay compensation offset as configured by CANFDx_CHy_TDCR.TDCO. The transmitter delay compensation offset is used to adjust the position of the SSP inside the received bit (for example, half of the bit time in the data phase). The position of the secondary sample

point is rounded down to the next integer number of mtq (PCLK_CANFD[x]_CLOCK_CAN[y] period).

CANFDx_CHy_PSR.TDCV shows the actual transmitter delay compensation value. CANFDx_CHy_PSR.TDCV is cleared when CANFDx_CHy_CCCR.INIT is set and is updated at each transmission of an FD frame while CANFDx_CHy_DBTP.TDC is set.

The following boundary conditions must be considered for the transmitter delay compensation implemented in the M_TTCAN:


- The sum of the measured delay from CANx_y_TX to CANx_y_RX and the configured transmitter delay compensation offset CANFDx_CHy_TDCR.TDCO must be less than 6 bit times in the data phase.

- The sum of the measured delay from CANx_y_TX to CANx_y_RX and the configured transmitter delay compensation offset CANFDx_CHy_TDCR.TDCO should be less than or equal 127 mtq. In case this sum exceeds 127 mtq, the maximum value of 127 mtq is used for transmitter delay compensation
- The data phase ends at the sample point of the CRC delimiter that stops checking of receive bits at the SSPs.

Transmitter Delay Compensation Measurement

If transmitter delay compensation is enabled by programming CANFDx_CHy_DBTP.TDC = 1, the measurement is started within each transmitted CAN FD frame at the falling edge of bit FDF to bit res. The measurement is stopped when this edge is seen at the receive input CANx_y_RX of the transmitter. The resolution of this measurement is one mtq (minimum time quanta).

Figure 24-3. Transmitter Delay Measurement



To avoid this, a dominant glitch inside the received FDF bit ends the delay compensation measurement before the falling edge of the received res bit, resulting in an early SSP position. The use of a transmitter delay compensation filter window can be enabled by programming CANFDx_CHy_TDCR.TDCF. This defines a minimum value for the SSP position. Dominant edges on CANx_y_RX, that results in an earlier SSP position are ignored for transmitter delay measurement. The measurement is stopped when the SSP position is at least CANFDx_CHy_TDCR.TDCF and CANx_y_RX is low.

24.3.1.5 Restricted Operation mode

In Restricted Operation mode, the node is able to receive data and remote frames and acknowledge valid frames, but it does not send data frames, remote frames, active error frames, or overload frames. In case of an error or overload condition, it does not send dominant bits; instead it waits for the occurrence of a bus idle condition to resynchronize itself to the CAN communication. The error counters (CANFDx_CHy_ECR.REC and CANFDx_CHy_ECR.TEC)

are frozen while Error Logging (CANFDx_CHy_ECR.CEL) is active.

The CPU can set the CAN FD controller into Restricted Operation mode by setting the Restricted Operation mode bit (CANFDx_CHy_CCCR.ASM). CANFDx_CHy_CCCR.ASM can only be set by the CPU when both CANFDx_CHy_CCCR.CCE and CANFDx_CHy_CCCR.INIT are set to '1'. CANFDx_CHy_CCCR.ASM can be reset by the CPU at any time.

The CAN FD controller enters Restricted Operation mode automatically when the TX handler is not able to read data from the message RAM in time. To leave this mode, the CPU should reset CANFDx_CHy_CCCR.ASM.

The Restricted Operation mode can be used in applications that adapt themselves to different CAN bit rates. In this case, the application tests different bit rates and leaves the mode after it has received a valid frame.


Note: The Restricted Operation mode must not be combined with the Loop Back mode (internal or external).

24.3.1.6 Bus Monitoring Mode

The M_TTCAN is set in Bus Monitoring mode by programming CANFDx_CHy_CCCR.MON to '1' or when error level S3 (CANFDx_CHy_TTOST.EL = 11) is entered. In Bus Monitoring mode, the M_TTCAN is able to receive valid data frames and valid remote frames, but cannot start a transmission. In this mode, it sends only recessive bits on the CAN bus, if the M_TTCAN is required to send a dominant bit (ACK bit, overload flag, or active error flag), the bit is rerouted internally so that the M_TTCAN monitors this dominant bit, although the CAN bus may remain in recessive state. In Bus Monitoring mode, the CANFDx_CHy_TXBRP register is held in reset state.

The Bus Monitoring mode can be used to analyze the traffic on a CAN bus without affecting it by the transmission of dominant bits. Figure 24-4 shows the connection of signals CANx_y_TX and CANx_y_RX to the M_TTCAN in Bus Monitoring mode.

Figure 24-4. Pin Control in Bus Monitoring Mode



Bus Monitoring Mode

24.3.1.7 Disable Automatic Retransmission

M_TTCAN supports automatic retransmission of frames that have lost arbitration or that have been disturbed by errors during transmission. By default, automatic retransmission is enabled. To support time-triggered communication (as described in ISO 11898-1:2015, chapter 9.2), the automatic retransmission may be disabled via CANFDx_CHy_CCCR.DAR.

In DAR mode, all transmissions are automatically canceled after they are started on the CAN bus. The TX Request Pending bit CANFDx_CHy_TXBRP.TRPx is reset after successful transmission, when a transmission has not yet started at the point of cancellation, is aborted due to lost arbitration, or when an error occurred during frame transmission.

- Successful transmission:
 - Corresponding TX Buffer Transmission Occurred bit CANFDx_CHy_TXBTO.TOx set
 - Corresponding TX Buffer Cancellation Finished bit CANFDx_CHy_TXBCF.CFx not set

- Successful transmission in spite of cancellation:
 - Corresponding TX Buffer Transmission Occurred bit CANFDx_CHy_TXBTO.TOx set
 - Corresponding TX Buffer Cancellation Finished bit CANFDx_CHy_TXBCF.CFx set
- Arbitration lost or frame transmission disturbed:
 - Corresponding TX Buffer Transmission Occurred bit CANFDx_CHy_TXBTO.TOx not set
 - Corresponding TX Buffer Cancellation Finished bit CANFDx_CHy_TXBCF.CFx set

In successful frame transmissions, and if storage of TX events is enabled, a TX Event FIFO element is written with Event Type ET = 10 (transmission despite cancellation).

24.3.1.8 Power Down (Sleep Mode)

The M_TTCAN channel can be set into power down mode via Clock Stop Request (CANFDx_CTL.STOP_REQ). As long as clock stop request is active, STOP_REQ bit is read as one.

When all pending transmission requests have completed, the M_TTCAN waits until bus idle state is detected. Then the M_TTCAN sets CANFDx_CHy_CCCR.INIT to one to prevent any further CAN transfers. Now the M_TTCAN acknowledges that it is ready for power down by setting Clock Stop Acknowledge (CANFDx_STATUS.STOP_ACK). Upon receiving acknowledgment from channel, hardware automatically switches off the clock to the respective channel.

To leave power down mode, the application must reset CANFDx_CTL.STOP_REQ. The M_TTCAN will acknowledge this by resetting CANFDx_STATUS.STOP_ACK. Afterwards, the application can restart CAN communication by resetting bit CANFDx_CHy_CCCR.INIT.

When the clock stop request is triggered through CANFDx_CTL.STOP_REQ, it must not be cleared before CANFDx_STATUS.STOP_ACK bit is set.

Note: Do not use the TTCAN CANFDx_CHy_CCCR.CSR register for the power down control, instead use CANFDx_CTL.STOP_REQ. Similarly, use of CANFDx_CHy_CCCR.CSA should be avoided, instead use CANFDx_STAUTS.STOP_ACK.

24.3.1.9 Test Mode

To enable write access to CANFDx_CHy_TEST register, Test Mode Enable bit (CANFDx_CHy_CCCR.TEST) must be set to one. This allows the configuration of the test modes and test functions.

Four output functions are available for the CAN transmit pin CANx_y_TX by programming CANFDx_CHy_TEST.TX. Apart from its default function of serial data output, it can drive the CAN Sample Point signal to monitor the

M_TTCAN's bit timing; it can also drive constant dominant or recessive values. The actual value at the CANx_y_RX pin can be read from CANFDx_CHy_TEST.RX. Both functions can be used to check the CAN bus physical layer.

Due to the synchronization mechanism between CAN clock and host clock domain, there may be a delay of several host clock periods between writing to CANFDx_CHy_TEST.TX until the new configuration is visible at output pin CANx_y_TX. This applies also when reading input pin CANx_y_RX via CANFDx_CHy_TEST.RX.

Note: Test modes should be used for production tests or self-test only. The software control for pin CANx_y_TX interferes with all CAN protocol functions. It is not recommended to use test modes for application.

External Loop Back Mode

The M_TTCAN can be set in External Loop Back mode by programming CANFDx_CHy_TEST.LBCK to one. In Loop Back mode, the M_TTCAN treats its own transmitted messages as received messages and stores them (if they pass acceptance filtering) into an RX buffer or an RX FIFO. [Figure 24-5](#) shows the connection of signals CANx_y_TX


and CANx_y_RX to the M_TTCAN in External Loop Back mode.

This mode is provided for hardware self-test. To be independent from external stimulation, the M_TTCAN ignores acknowledge errors (recessive bit sampled in the acknowledge slot of a data/remote frame) in Loop Back mode. In this mode the M_TTCAN performs an internal feedback from its TX output to its RX input. The actual value of the CANx_y_RX input pin is disregarded by the M_TTCAN. The transmitted messages can be monitored at the CANx_y_TX pin.

Internal Loop Back Mode

Internal Loop Back mode is entered by programming the CANFDx_CHy_TEST.LBCK and CANFDx_CHy_CCCR.MON bits to one. This mode can be used for a “Hot Selftest”, meaning the M_TTCAN can be tested without affecting a running CAN system connected to the CANx_y_TX and CANx_y_RX pins. In this mode, CANx_y_RX pin is disconnected from the M_TTCAN and CANx_y_TX pin is held recessive. [Figure 24-5](#) shows the connection of CANx_y_TX and CANx_y_RX to the M_TTCAN in Internal Loop Back mode.

Figure 24-5. Pin Control in Loop Back Modes



24.3.1.10 Application Watchdog

The application watchdog is served by reading the CANFDx_CHy_TTOST register. When the application watchdog is not served in time, CANFDx_CHy_TTOST.AWE bit is set, all TTCAN communication is stopped, and the M_TTCAN is set into Bus Monitoring mode.

The TT application watchdog can be disabled by programming the Application Watchdog Limit CANFDx_CHy_TTOCF.AWL to 0x00. The TT application watchdog should not be disabled in a TTCAN application program.


24.3.2 Timestamp Generation

The M_TTCAN channel uses a 16-bit counter to record when messages are sent or received. This allows the

application software to know the order in which events occurred.

To keep event ordering across multiple M_TTCAN channels, a global timestamp counter is implemented, which must be selected by setting '10' to CANFDx_CHy_TSCC.TSS[1:0]. This global timestamp counter is shared among all M_TTCAN groups present in the device. For instance, if the device contains two M_TTCAN groups, timestamp counter is shared among all the channels present in both the groups.

Figure 24-6. Timestamp Connection between Two M_TTCAN Group



The timestamp counter is configured through the CANFDx_TS_CTL register. The CANFDx_TS_CTL.ENABLED bit will enable the counter. Upon enabling, it will start incrementing according to the CANFDx_TS_CTL.PRESCALE [15:0]. The application can read the counter value through the CANFDx_TS_CNT register. Write access to the CANFDx_TS_CNT register will clear the CANFDx_TS_CNT.

When the timestamp counter is enabled, internal counter for prescaler counts with every cycle of CLK_SYS; when the counter value reaches the prescaler value, the timestamp counter increments by one and internal prescaler counter is cleared. When CANFDx_TS_CTL.PRESCALE changes, CANFDx_TS_CNT should be written to reset them. This can make the internal prescaler counter follow a new value of CANFDx_TS_CTL.PRESCALE immediately.

The shared timestamp counter is a wrap-around counter. When the counter wraps around, CANFDx_CHy_IR.TSW for all M_TTCAN channels will be raised.

On start of frame reception/transmission, the timestamp counter value is captured and stored into the timestamp section of an RX buffer/RX FIFO (RXTS [15:0]) or TX Event FIFO (TXTS [15:0]) element. **Note:** The counter value CANFDx_TS_CNT is not retained in DeepSleep mode whereas the CANFDx_TS_CTL is retained.

24.3.3 Timeout Counter

To signal timeout conditions for RX FIFO 0, RX FIFO 1, and the TX Event FIFO, the M_TTCAN supplies a 16-bit Timeout Counter. It operates as down-counter and uses the same prescaler controlled by CANFDx_CHy_TSCC.TCP as the timestamp Counter. A prescaler CANFDx_CHy_TSCC.TCP

should be configured to clock the timeout counter in multiples of CAN bit times (1...16). The timeout counter is configured via register CANFDx_CHy_TOCC. The actual counter value can be read from CANFDx_CHy_TOCV.TOC.

The timeout counter can only be started while CANFDx_CHy_CCCR.INIT = 0. It is stopped when CANFDx_CHy_CCCR.INIT = 1; for example, when the M_TTCAN enters Bus_Off state.

The operation mode is selected by CANFDx_CHy_TOCC.TOS. When operating in Continuous mode, the counter starts when CANFDx_CHy_CCCR.INIT is reset. A write to CANFDx_CHy_TOCV presets the counter to the value configured by CANFDx_CHy_TOCC.TOP and continues down-counting.

When the timeout counter is controlled by one of the FIFOs, an empty FIFO presets the counter to the value configured by CANFDx_CHy_TOCC.TOP. Down-counting is started when the first FIFO element is stored. Writing to CANFDx_CHy_TOCV has no effect.

When the counter reaches zero, interrupt flag CANFDx_CHy_IR.TOO is set. In Continuous mode, the counter is immediately restarted at CANFDx_CHy_TOCC.TOP.

Note: The clock signal for the timeout counter is derived from the CAN Core's sample point signal. Therefore, the time the Timeout Counter is decremented may vary due to the synchronization/resynchronization mechanism of the CAN Core. If the bit rate switch feature in CAN FD is used, the timeout counter is clocked differently in arbitration and data field.

24.3.4 RX Handling

The RX handler controls acceptance filtering, transfer of received messages to the RX buffers or to one of the two RX FIFOs, as well as RX FIFO's Put and Get Indices.

24.3.4.1 Acceptance Filtering

The M_TTCAN offers the possibility to configure two sets of acceptance filters, one for standard identifiers and one for extended identifiers. These filters can be assigned to an RX buffer or to RX FIFO 0,1. For acceptance filtering each list of filters is executed from element #0 until the first matching element. Acceptance filtering stops at the first matching element. The following filter elements are not evaluated for this message.

The main features are:

- Each filter element can be configured as
 - Range filter (from - to)
 - Filter for one or two dedicated IDs
 - Classic bit mask filter
- Each filter element is configurable for acceptance or rejection filtering
- Each filter element can be enabled/disabled individually
- Filters are checked sequentially; execution stops with the first matching filter element

Related configuration registers are:

- Global Filter Configuration (CANFDx_CHy_GFC)
- Standard ID Filter Configuration (CANFDx_CHy_SIDFC)
- Extended ID Filter Configuration (CANFDx_CHy_XIDFC)
- Extended ID AND Mask (CANFDx_CHy_XIDAM)

Depending on the configuration of the filter element (SFEC/EFEC) a match triggers one of the following actions:

- Store received frame in FIFO 0 or FIFO 1
- Store received frame in RX buffer
- Store received frame in RX buffer and generate pulse at filter event pin
- Reject received frame
- Set High-Priority Message interrupt flag CANFDx_CHy_IR.HPM
- Set High-Priority Message interrupt flag CANFDx_CHy_IR.HPM and store received frame in FIFO 0 or FIFO 1

Acceptance filtering is started after the complete identifier is received. After acceptance filtering has completed, if a matching RX buffer or RX FIFO is found, the message handler starts writing the received message data in portions of 32 bits to the matching RX buffer or RX FIFO. If the CAN protocol controller has detected an error condition (such as CRC error), this message is discarded with the following impact on the affected RX buffer or RX FIFO:

■ RX Buffer

New data flag of the matching RX buffer is not set, but RX buffer is (partly) overwritten with received data. For error type, see CANFDx_CHy_PSR.LEC and CANFDx_CHy_PSR.DLEC, respectively.

■ RX FIFO

Put index of matching RX FIFO is not updated, but related RX FIFO element is (partly) overwritten with received data. For error type, see CANFDx_CHy_PSR.LEC and CANFDx_CHy_PSR.DLEC, respectively. If the matching RX FIFO is operated in overwrite mode, the boundary conditions described in [RX FIFO Overwrite Mode](#) should be considered.

Note: When an accepted message is written to one of the two RX FIFOs, or into an RX buffer, the unmodified received identifier is stored independent of the filter(s) used. The result of the acceptance filter process depends on the sequence of configured filter elements.

Range Filter

The filter matches for all received frames with Message IDs in the range defined by SF1ID/SF2ID resp. EF1ID/EF2ID.

The two possibilities when range filtering is used with extended frames are:

- EFT = 00: The Message ID of received frames is ANDed with the CANFDx_CHy_XIDAM before the range filter is applied
- EFT = 11: The CANFDx_CHy_XIDAM is not used for range filtering

Filter for specific IDs

A filter element can be configured to filter one or two specific Message IDs. To filter a specific Message ID, the filter element should be configured with SF1ID = SF2ID and EF1ID = EF2ID, respectively.

Classic Bit Mask Filter

Classic bit mask filtering is intended to filter groups of Message IDs by masking single bits of a received Message ID. With classic bit mask filtering, SF1ID/EF1ID is used as Message ID filter, while SF2ID/EF2ID is used as filter mask.

A zero bit at the filter mask will mask the corresponding bit position of the configured ID filter; for example, the value of the received Message ID at that bit position is not relevant for acceptance filtering. Only those bits of the received Message ID where the corresponding mask bits are one are relevant for acceptance filtering.


In case all mask bits are one, a match occurs only when the received Message ID and the Message ID filter are identical. If all mask bits are zero, all Message IDs match.

Standard Message ID Filtering

Figure 24-7 shows the flow for standard Message ID (11-bit Identifier) filtering. The Standard Message ID Filter element is described in [Standard Message ID Filter Element](#) on page 354.

Controlled by the CANFDx_CHy_GFC and CANFDx_CHy_SIDFC Message IDs, the Remote Transmission Request bit (RTR) and the Identifier Extension bit (IDE) of received frames are compared against the list of configured filter elements.

Figure 24-7. Standard Message ID Filter




Extended Message ID Filtering

Figure 24-8 shows the flow for extended Message ID (29-bit Identifier) filtering. The Extended Message ID Filter element is described in [Extended Message ID Filter Element](#) on page 356.

Controlled by the CANFDx_CHy_GFC and CANFDx_CHy_XIDFC Message IDs, the RTR bit, and IDE bit of received frames are compared against the list of configured filter elements.

The Extended ID AND Mask XIDAM[28:0] is ANDed with the received identifier before the filter list is executed.

Figure 24-8. Extended Message ID Filter Path



24.3.4.2 RX FIFOs


RX FIFO 0 and RX FIFO 1 can be configured to hold up to 64 elements each. The two RX FIFOs are configured via the CANFDx_CHy_RXF0C and CANFDx_CHy_RXF1C registers.

Received messages that pass acceptance filtering are transferred to the RX FIFO as configured by the matching filter element. For a description of the filter mechanisms available for RX FIFO 0 and RX FIFO 1, see [Acceptance Filtering on page 337](#). The RX FIFO element is described in [RX Buffer and FIFO Element on page 350](#).

To avoid an RX FIFO overflow, the RX FIFO watermark can be used. When the RX FIFO fill level reaches the RX FIFO watermark configured by CANFDx_CHy_RXFnC.FnWM,

interrupt flag CANFDx_CHy_IR.RFnW is set. When the RX FIFO Put Index reaches the RX FIFO Get Index, an RX FIFO Full condition is signaled by CANFDx_CHy_RXFoS.FnF. In addition, interrupt flag CANFDx_CHy_IR.RFnF is set. The FIFO watermark interrupt flags can be used to trigger the DMA. DMA request for FIFO will remain set until the respective trigger is cleared by software. Software can clear the trigger by clearing the watermark flag.

Figure 24-9. RX FIFO Status



When reading from an RX FIFO, RX FIFO Get Index $\text{CANFDx\_CHy\_RXFnS.FnGI} \times \text{FIFO Element Size}$ has to be added to the corresponding RX FIFO start address $\text{CANFDx\_CHy\_RXFnC.FnSA}$. RX FIFO Top pointer logic is added to the CAN FD controller to make reading faster. See [RX FIFO Top Pointer on page 341](#).

Table 24-2. RX Buffer/FIFO Element size

$\text{CANFDx\_CHy\_RXE SC.RBDS}[2:0]$ $\text{CANFDx\_CHy\_RXE SC.FnDS}[2:0]$	Data Field [bytes]	FIFO Element Size [RAM words]
000	8	4
001	12	5
010	16	6
011	20	7
100	24	8
101	32	10
110	48	14
111	64	18

RX FIFO Blocking Mode

The RX FIFO blocking mode is configured by $\text{CANFDx\_CHy\_RXFnC.FnOM} = 0$. This is the default operation mode for RX FIFOs.

When an RX FIFO full condition is reached ($\text{CANFDx\_CHy\_RXFnS.FnPI} = \text{CANFDx\_CHy\_RXFnS.FnGI}$), no further messages are written to the corresponding RX FIFO until at least one message is read and the RX FIFO Get Index is incremented. An RX FIFO full condition is signaled by $\text{CANFDx\_CHy\_RXFnS.FnF} = 1$. In addition, the interrupt flag $\text{CANFDx\_CHy\_IR.RFnF}$ is set.

If a message is received while the corresponding RX FIFO is full, this message is discarded and the message lost condition is signaled by $\text{CANFDx\_CHy\_RXFnS.RFnL} = 1$. In addition, the interrupt flag $\text{CANFDx\_CHy\_IR.RFnL}$ is set.


RX FIFO Overwrite Mode

The RX FIFO overwrite mode is configured by $\text{CANFDx\_CHy\_RXFnC.FnOM} = 1$.

When an RX FIFO full condition ($\text{CANFDx\_CHy\_RXFnS.FnPI} = \text{CANFDx\_CHy\_RXFnS.FnGI}$) is signaled by $\text{CANFDx\_CHy\_RXFnS.FnF} = 1$, the next message accepted for the FIFO will overwrite the oldest FIFO message. Put and Get indices are both incremented by one.

When an RX FIFO is operated in overwrite mode and an RX FIFO full condition is signaled, reading of the RX FIFO elements should start at least at Get Index + 1. This is because a received message may be written to the message RAM (Put Index) while the CPU is reading from the message RAM (Get Index). In this case, inconsistent data may be read from the respective RX FIFO element. Adding an offset to the Get Index when reading from the RX FIFO avoids this problem. The offset depends on how fast the CPU accesses the RX FIFO. [Figure 24-10](#) shows an offset of two with respect to the Get Index when reading the RX FIFO. In this case, the two messages stored in element 1 and 2 are lost.

Figure 24-10. RX FIFO Overflow Handling



After reading from the RX FIFO, the number of the last element read must be written to the RX FIFO Acknowledge Index CANFDx_CHy_RXFnA.FnA. This increments the Get Index to that element number. If the Put Index is not incremented to this RX FIFO element, the RX FIFO full condition is reset (CANFDx_CHy_RxFnS.FnF = 0).

RX FIFO Top Pointer

M_TTCAN supports two receive FIFOs. Reading from these FIFOs requires application to go through following steps:

- Retrieve read pointer
- Calculate correct message RAM address
- Read the data from message RAM
- Update the read pointer

To avoid all these steps, RX FIFO Top Pointer logic has been integrated in the CAN FD controller. It provides a single MMIO location (CANFDx_CHy_RXFTOPn_DATA; n = 0,1) to read the data from. Using such hardware logic has the following benefits:

- Higher performance data access
- Less bus traffic
- Reduced CPU load
- Reduced power
- Enables DMA access to FIFO

This logic is enabled when CANFDx_CHy_RXFTOP_CTL.FnTPE is set. Setting this bit enables the logic to set the FIFO top address (FnTA) and internal message word counter. Receive FIFO in the top status register (CANFDx_CHy_RXFTOPn_STAT) shows the respective FIFO top address and CANFDx_CHy_RXFTOPn_DATA provides the data located at the top address. Refer to register definitions for more details on both registers.

If CANFDx_CHy_RXFTOPn_DATA is read, the top pointer logic also updates the RX FIFO Acknowledge Index (CANFDx_CHy_RXFnA.FnA) in TTCAN channel.

Note: Top pointer logic is disabled when the channel is being configured (CANFDx_CHy_CCCR.CCE = 1). Reading CANFDx_CHy_RXFTOPn_DATA while the logic is disabled will return the invalid data.

24.3.4.3 Dedicated RX Buffers

The M_TTCAN supports up to 64 dedicated RX buffers. The start address of the dedicated RX buffer section is configured via CANFDx_CHy_RXBC.RBSA.

For each RX buffer, a Standard or Extended Message ID Filter Element with SFEC/EFEC = 111 and SFID2/EFID2[10:9] = 00 must be configured (see 24.4.5 Standard Message ID Filter Element and 24.4.6 Extended Message ID Filter Element).

After a received message is accepted by a filter element, the message is stored into the RX buffer in the message RAM referenced by the filter element. The format is the same as for an RX FIFO element. In addition, the flag CANFDx_CHy_IR.DRX (message stored in a dedicated RX buffer) in the interrupt register is set.

Table 24-3. Example Filter Configuration for RX Buffers

Filter Element	SFID1[10:0] EFID1[28:0]	SFID2[10:9] EFID2[10:9]	SFID2[5:0] EFID2[5:0]
0	ID message 1	00	00 0000
1	ID message 2	00	00 0001
2	ID message 3	00	00 0010

After the last word of a matching received message is written to the message RAM, the respective New Data flag in CANFDx_CHy_NDAT1 and CANFDx_CHy_NDAT2 registers is set. As long as the New Data flag is set, the

respective RX buffer is locked against updates from received matching frames. The New Data flags should be reset by the host by writing a '1' to the respective bit position.

While an RX buffer's New Data flag is set, a Message ID Filter Element referencing the specific RX buffer will not match, causing the acceptance filtering to continue. The following Message ID Filter Elements may cause the received message to be stored into another RX buffer, or into an RX FIFO, or the message may be rejected, depending on filter configuration.

- Reset interrupt flag CANFDx_CHy_IR.DRX
- Read New Data registers
- Read messages from message RAM
- Reset New Data flags of processed messages

24.3.4.4 Debug on CAN Support

Debug messages are stored into RX buffers; three consecutive RX buffers (for example, #61, #62, and #63) should be used to store debug messages A, B, and C. The format is the same for RX buffer and RX FIFO elements.

To filter debug messages Standard/Extended Filter Elements with SFEC/EFEC = "111" should be set up. Messages that match these filter elements are stored into the RX buffers addressed by SFID2/EFID2[5:0].

After message C is stored, the DMA request is activated and the three messages can be read from the message RAM under DMA control. The RAM words holding the debug messages will not be changed by the M_TTCAN while DMA request is activated. The behavior is similar to that of an RX buffer with its New Data flag set.

After the DMA transfer is completed, an acknowledge from DMA resets the DMA request. Now the M_TTCAN is prepared to receive the next set of debug messages.

Filtering Debug Messages

Debug messages are filtered by configuring one Standard/Extended Message ID filter element for each of the three debug messages. To enable a filter element to filter debug messages, SFEC/EFEC should be programmed to "111". In this case the SFID1/SFID2 and EFID1/EFID2 fields have a different meaning (see [Standard Message ID Filter Element on page 354](#) and [Extended Message ID Filter Element on page 356](#)). While SFID2/EFID2[10:9] controls the debug message handling state machine, SFID2/EFID2[5:0] controls the storage location of a received debug message.

When a debug message is stored, neither the respective New Data flag nor CANFDx_CHy_IR.DRX are set. The reception of debug messages can be monitored via CANFDx_CHy_RXF1S.DMS.

Table 24-4. Example Filter Configuration for Debug Message


Filter Element	SFID1[10:0] EFID1[28:0]	SFID2[10:9] EFID2[10:9]	SFID2[5:0] EFID2[5:0]
0	ID debug message A	01	11 1101
1	ID debug message B	10	11 1110
2	ID debug message C	11	11 1111

Debug Message Handling

The debug message handling state machine assures that debug messages are stored to three consecutive RX buffers in the correct order. If there are missing messages, the process is restarted. The DMA request is activated only when all three debug messages A, B, and C are received in correct order.

The status of the debug message handling state machine is signaled via CANFDx_CHy_RXF1S.DMS.

Figure 24-11. Debug Message Handling State Machine



24.3.5 TX Handling

The TX handler handles transmission requests for the dedicated TX buffers, TX FIFO, and TX Queue. It controls the transfer of transmit messages to the CAN Core, the Put and Get Indices, and the TX Event FIFO. Up to 32 TX buffers can be set up for message transmission. The CAN mode for transmission (Classic CAN or CAN FD) can be configured separately for each TX buffer element. The TX buffer element is described in [TX Buffer Element on page 352](#). [Table 24-5](#) describes the possible configurations for frame transmission.

Table 24-5. Possible Configuration for Frame Transmission

CANFDx_CHy_C CCR		TX Buffer Element		Frame Transmission
BRSE	FDOE	FDF	BRS	
ignored	0	ignored	ignored	Classic CAN
0	1	0	ignored	Classic CAN
0	1	1	ignored	FD without bit rate switching
1	1	0	ignored	Classic CAN
1	1	1	0	FD without bit rate switching
1	1	1	1	FD with bit rate switching

The TX handler starts a TX scan to check for the highest priority pending TX request (TX buffer with lowest Message ID) when the TX Buffer Request Pending register (CANFDx_CHy_TXBRP) is updated, or when a transmission is started.

24.3.5.1 Transmit Pause

The transmit pause feature is intended for use in CAN systems where the CAN message identifiers are (permanently) assigned to specific values and cannot be changed easily. These message identifiers may have a higher CAN arbitration priority than other defined messages, while in a specific application their relative arbitration priority should be inverse. This may lead to a case where one Electronic Control Unit (ECU) sends a burst of CAN messages that cause another ECU's CAN messages to be delayed because the other messages have a lower CAN arbitration priority.

For example, if CAN ECU-1 has the transmit pause feature enabled and is requested by the application software to transmit four messages, it will, after the first successful message transmission, wait for two nominal bit times of bus idle before it is allowed to start the next requested message. If there are other ECUs with pending messages, those messages are started in the idle time, they will not need to arbitrate with the next message of ECU-1. After having received a message, ECU-1 is allowed to start its next transmission as soon as the received message releases the CAN bus.

The transmit pause feature is controlled by the Transmit Pause bit (CANFDx_CHy_CCCR.TXP). If the bit is set, the M_TTCAN controller will, each time it has successfully transmitted a message, pause for two nominal bit times before starting the next transmission. This enables other CAN nodes in the network to transmit messages even if their messages have lower prior identifiers. Default is transmit pause disabled (CANFDx_CHy_CCCR.TXP = 0).

This feature loses burst transmissions coming from a single node and protects against "babbling idiot" scenarios where

the application program erroneously requests too many transmissions.

24.3.5.2 Dedicated TX Buffers

Dedicated TX buffers are intended for message transmission under complete control of the CPU. Each dedicated TX buffer is configured with a specific Message ID. If multiple TX buffers are configured with the same Message ID, the TX buffer with the lowest buffer number is transmitted first.

If the data section is updated, a transmission is requested by an Add Request via CANFDx_CHy_TXBAR.ARn. The requested messages arbitrate internally with messages from an optional TX FIFO or TX Queue and externally with messages on the CAN bus, and are sent out according to their Message ID.

Addressing Dedicated TX Buffers

A dedicated TX buffer allocates an Element Size of 32-bit words in the message RAM as shown in the [Table 24-6](#). Therefore, the start address of a dedicated TX buffer in the message RAM is calculated by

(transmit buffer index (0 to 31) × Element Size) + TX Buffers Start Address (CANFDx_CHy_TXBC.TBSA[15:2]).

Table 24-6. TX Buffer/FIFO/Queue Element size

CANFDx_CHy_TXESC.TBDS[2:0]	Data Field [bytes]	Element Size [RAM words]
000	8	4
001	12	5
010	16	6
011	20	7
100	24	8
101	32	10
110	48	14
111	64	18

24.3.5.3 TX FIFO

TX FIFO operation is configured by programming CANFDx_CHy_TXBC.TFQM to '0'. Messages stored in the TX FIFO are transmitted starting with the message referenced by the Get Index CANFDx_CHy_TXFQS.TFGI. After each transmission, the Get Index is incremented cyclically until the TX FIFO is empty. The TX FIFO enables transmission of messages with the same Message ID from different TX buffers in the order these messages are written to the TX FIFO. The M_TTCAN calculates the TX FIFO Free Level CANFDx_CHy_TXFQS.TFFL as difference between Get and Put Index. It indicates the number of available (free) TX FIFO elements.

New transmit messages must be written to the TX FIFO starting with the TX buffer referenced by the Put Index CANFDx_CHy_TXFQS.TFQPI. An Add Request increments the Put Index to the next free TX FIFO element. When the Put Index reaches the Get Index, TX FIFO Full (CANFDx_CHy_TXFQS.TFQF = 1) is signaled. In this case no further messages should be written to the TX FIFO until the next message is transmitted and the Get Index is incremented.

When a single message is added to the TX FIFO, the transmission is requested by writing a '1' to the CANFDx_CHy_TXBAR bit related to the TX buffer referenced by the TX FIFO's Put Index.

When multiple (n) messages are added to the TX FIFO, they are written to n consecutive TX buffers starting with the Put Index. The transmissions are then requested via CANFDx_CHy_TXBAR. The Put Index is then cyclically incremented by n. The number of requested TX buffers should not exceed the number of free TX buffers as indicated by the TX FIFO Free Level.

When a transmission request for the TX buffer referenced by the Get Index is canceled, the Get Index is incremented to the next TX buffer with pending transmission request and the TX FIFO Free Level is recalculated. When transmission cancellation is applied to any other TX buffer, the Get Index and the FIFO Free Level remain unchanged.

A TX FIFO element allocates Element Size 32-bit words in the message RAM as shown in [Table 24-6](#). Therefore, the start address of the next available (free) TX FIFO buffer is calculated by adding TX FIFO/Queue Put Index CANFDx_CHy_TXFQS.TFQPI (0...31) • Element Size to the TX buffer Start Address CANFDx_CHy_TXBC.TBSA.

24.3.5.4 TX Queue

TX Queue operation is configured by programming CANFDx_CHy_TXBC.TFQM to '1'. Messages stored in the TX Queue are transmitted starting with the message with the lowest Message ID (highest priority). If multiple queue buffers are configured with the same Message ID, the queue buffer with the lowest buffer number is transmitted first.

New messages must be written to the TX buffer referenced by the Put Index CANFDx_CHy_TXFQS.TFQPI. An Add Request cyclically increments the Put Index to the next free TX buffer. If the TX Queue is full (CANFDx_CHy_TXFQS.TFQF = 1), the Put Index is not valid and no further message should be written to the TX Queue until at least one of the requested messages is sent out or a pending transmission request is canceled.

The application may use the CANFDx_CHy_TXBRP register instead of the Put Index and may place messages to any TX buffer without pending transmission request.

A TX Queue buffer allocates element size of 32-bit words in the message RAM as shown in [Table 24-6](#). Therefore, the


start address of the next available (free) TX Queue buffer is calculated by adding TX FIFO/Queue Put Index `CANFDx_CHy_TXFQS.TFQPI` (0...31) × Element Size to the TX buffer Start Address `CANFDx_CHy_TXBC.TBSA`.

24.3.5.5 Mixed Dedicated TX Buffers/TX FIFO

In this case, the TX Buffers section in the message RAM is subdivided into a set of dedicated TX buffers and a TX

FIFO. The number of dedicated TX buffers is configured by `CANFDx_CHy_TXBC.NDTB`. The number of TX buffers assigned to the TX FIFO is configured by `CANFDx_CHy_TXBC.TFQS`. If `CANFDx_CHy_TXBC.TFQS` is programmed to zero, only the dedicated TX buffers are used.

Figure 24-12. Example of Mixed Configuration Dedicated TX Buffers/TX FIFO




TX prioritization:

- Scan dedicated TX buffers and oldest pending TX FIFO buffer (referenced by `CANFDx_CHy_TXFS.TFGI`)
- Buffer with the lowest Message ID gets highest priority and is transmitted next

24.3.5.6 Mixed Dedicated TX Buffers/TX Queue

In this case the TX Buffers section in the message RAM is subdivided into a set of dedicated TX buffers and a TX Queue. The number of dedicated TX buffers is configured by `CANFDx_CHy_TXBC.NDTB`. The number of TX Queue buffers is configured by `CANFDx_CHy_TXBC.TFQS`. In case `CANFDx_CHy_TXBC.TFQS` is programmed to zero, only dedicated TX buffers are used.

Figure 24-13. Example of Mixed Configuration Dedicated TX Buffers/TX Queue



TX prioritization:

- Scan all TX buffers with activated transmission request
- TX buffer with the lowest Message ID gets highest priority and is transmitted next

24.3.5.7 Transmit Cancellation

The `M_TTCAN` supports transmit cancellation. This feature is especially intended for gateway applications and AUTOSAR-based applications. To cancel a requested transmission from a dedicated TX buffer or a TX Queue buffer the host must write a '1' to the corresponding bit position (number of TX buffers) of the

`CANFDx_CHy_TXBCR` register. Transmit cancellation is not intended for TX FIFO operation.

Successful cancellation is signaled by setting the corresponding bit of the `CANFDx_CHy_TXBCF` register to '1'.

In case a transmit cancellation is requested while a transmission from a TX buffer is ongoing, the corresponding `CANFDx_CHy_TXBRP` bit remains set as long as the transmission is in progress. If the transmission was successful, the corresponding `CANFDx_CHy_TXBTO` and `CANFDx_CHy_TXBCF` bits are set. If the transmission was not successful, it is not repeated and only the corresponding `CANFDx_CHy_TXBCF` bit is set.

Note: If a pending transmission is canceled immediately before this transmission is started, there follows a short time window where no transmission is started even if another message is also pending in this node. This may enable another node to transmit a message, which may have a lower priority than the second message in this node.

24.3.5.8 TX Event Handling

To support TX event handling, the M_TTCAN has implemented a TX Event FIFO. After the M_TTCAN has transmitted a message on the CAN bus, the Message ID and timestamp are stored in a TX Event FIFO element. To link a TX event to a TX Event FIFO element, the Message Marker from the transmitted TX buffer is copied into the TX Event FIFO element.

The TX Event FIFO can be configured to a maximum of 32 elements. The TX Event FIFO element is described in [TX Event FIFO Element on page 353](#).

The purpose of the TX Event FIFO is to decouple handling transmit status information from transmit message handling; that is, a TX buffer holds only the message to be transmitted, while the transmit status is stored separately in the TX Event FIFO. This has the advantage, especially when operating a dynamically managed transmit queue, that a TX buffer can be used for a new message immediately after successful transmission. There is no need to save transmit status information from a TX buffer before overwriting that TX buffer.

When a TX Event FIFO full condition is signaled by CANFDx_CHy_IR.TEFF, no further elements are written to the TX Event FIFO until at least one element is read out and the TX Event FIFO Get Index is incremented. In case a TX event occurs while the TX Event FIFO is full, this event is discarded and interrupt flag CANFDx_CHy_IR.TEFL is set.

To avoid a TX Event FIFO overflow, the TX Event FIFO watermark can be used. When the TX Event FIFO fill level reaches the TX Event FIFO watermark configured by CANFDx_CHy_TXEFC.EFWM, interrupt flag CANFDx_CHy_IR.TEFW is set.

When reading from the TX Event FIFO, the TX Event FIFO Get Index CANFDx_CHy_TXEFS.EFGI must be added twice to the TX Event FIFO start address CANFDx_CHy_TXEFC.EFSA.

24.3.6 FIFO Acknowledge Handling

The Get indices of RX FIFO 0, RX FIFO 1, and the TX Event FIFO are controlled by the corresponding FIFO Acknowledge Index.

When RX FIFO top pointer hardware logic is used, it updates the RX FIFO Acknowledge Index. After CANFDx_CHy_RXFTOPn_DATA is read, the Acknowledge Index (CANFDx_CHy_RXFnA.FnA) is updated automatically, which will eventually set the FIFO Get Index

to the FIFO Acknowledge Index plus one and thereby updates the FIFO Fill Level.

When the application does not use RX FIFO top pointer logic, the Acknowledge Index must be updated. This can be done using one of the following two use cases:

- When only a single element is read from the FIFO (the one being pointed to by the Get Index), the Get Index value is written to the FIFO Acknowledge Index.
- When a sequence of elements is read from the FIFO, it is sufficient to write the FIFO Acknowledge Index only once at the end of that read sequence (value is the index of the last element read), to update the FIFO's Get Index.

Because the CPU has free access to the M_TTCAN's message RAM, take care when reading FIFO elements in an arbitrary order (Get Index not considered). This may be useful when reading a high-priority message from one of the two RX FIFOs. In this case the FIFO Acknowledge Index should not be written because this will set the Get Index to a wrong position and alter the FIFO's Fill Level. Some older FIFO elements are lost.

Note: The application must ensure that a valid value is written to the FIFO Acknowledge Index. The M_TTCAN does not check for erroneous values.

24.3.7 Configuring the CAN Bit Timing

Each node in the CAN network has its own clock generator (usually a quartz oscillator). The time parameter of the bit time can be configured individually for each CAN node. Even if each CAN node's oscillator has a different period, a common bit rate can be generated.


The oscillator frequencies vary slightly because of changes in temperature or voltage, or deterioration of components. As long as the frequencies vary only within the tolerance range of the oscillators, the CAN nodes can compensate for the different bit rates by resynchronizing to the bit stream.

24.3.7.1 CAN Bit Timing

The CAN FD operation defines two bit times – nominal bit time and data bit time. The nominal bit time is for the arbitration phase. The data bit time has an equal or shorter length and can be used to accelerate the data phase (see [CAN FD Operation on page 331](#)).

The basic construction of a bit time is shared with both the nominal and data bit times. The bit time can be divided into four segments according to the CAN specifications (see [Figure 24-14](#): the synchronization segment (Sync_Seg), the propagation time segment (Prop_Seg), the phase buffer segment 1 (Phase_Seg1), and the phase buffer segment 2 (Phase_Seg2). The sample point at which the bus level is read and interpreted as the value of that respective bit, is located at the end of Phase_Seg1.

Figure 24-14. Bit Time Construction



Each segment consists of a programmable number of time quanta, which is a multiple of the time quantum that is defined by PCLK_CANFD[x]_CLOCK_CAN[y] and a prescaler. The values and prescalers used to define these parameters differ for the nominal and data bit times, and are configured by CANFDx_CHy_NBTP (Nominal Bit Timing and Prescaler Register) and CANFDx_CHy_DBTP (Data Bit Timing and Prescaler Register) as shown in [Table 24-7](#).

Table 24-7. Bit Time Parameters

Parameter	Description
Time quantum tq (nominal) and tqd (data)	Time quantum. Derived by multiplying the basic unit time quanta (the PCLK_CANFD[x]_CLOCK_CAN[y] period) with the respective prescaler. The time quantum is configured by the CAN FD controller as nominal: $tq = (CANFDx\_CHy\_NBTP.NBRP[8:0] + 1) \times PCLK\_CANFD[x]\_CLOCK\_CAN[y] \text{ period}$ data: $tqd = (CANFDx\_CHy\_DBTP.DBRP[4:0] + 1) \times PCLK\_CANFD[x]\_CLOCK\_CAN[y] \text{ period}$
Sync_Seg	Sync_Seg is fixed to one time quantum as defined by the CAN specifications and is not configurable (inherently built into the CAN FD controller). nominal: 1 tq data: 1 tqd
Prop_Seg	Prop_Seg is the part of the bit time that is used to compensate for the physical delay times within the network. The CAN FD controller configures the sum of Prop_Seg and Phase_Seg1 with a single parameter: nominal: $Prop\_Seg + Phase\_Seg1 = CANFDx\_CHy\_NBTP.NTSEG1[7:0] + 1$ data: $Prop\_Seg + Phase\_Seg1 = CANFDx\_CHy\_DBTP.DTSEG1[4:0] + 1$
Phase_Seg1	Phase_Seg1 is used to compensate for edge phase errors before the sampling point. Can be lengthened by the resynchronization jump width. The sum of Prop_Seg and Phase_Seg1 is configured by the CAN FD controller as nominal: $CANFDx\_CHy\_NBTP.NTSEG1[7:0] + 1$ data: $CANFDx\_CHy\_DBTP.DTSEG1[4:0] + 1$
Phase_Seg2	Phase_Seg2 is used to compensate for edge phase errors after the sampling point. Can be shortened by the resynchronization jump width. Phase_Seg2 is configured by the CAN FD controller as nominal: $CANFDx\_CHy\_NBTP.NTSEG2[6:0] + 1$ data: $CANFDx\_CHy\_DBTP.DTSEG2[3:0] + 1$
SJW	Resynchronization Jump Width. Used to adjust the length of Phase_Seg1 and Phase_Seg2. SJW will not be longer than either Phase_Seg1 or Phase_Seg2. SJW is configured by the CAN FD controller as nominal: $CANFDx\_CHy\_NBTP.NSJW[6:0] + 1$ data: $CANFDx\_CHy\_DBTP.DSJW[3:0] + 1$

These relations result in the following equations for the nominal and data bit times:

Nominal Bit time

$$= [\text{Sync\_Seg} + \text{Prop\_Seg} + \text{Phase\_Seg1} + \text{Phase\_Seg2}] \times t_q$$

$$= \left[1 + (\text{CANFDx\_CHy\_NBTP.NTSEG1}[7:0] + 1) + (\text{CANFDx\_CHy\_NBTP.NTSEG2}[6:0] + 1) \right] \times \left[(\text{CANFDx\_CHy\_NBTP.NBRP}[8:0] + 1) \times \text{PCLK\_CANFD}[x]\_CLOCK\_CAN[y] \text{ period} \right]$$

Data Bit time

$$= \left[1 + (\text{CANFDx\_CHy\_DBTP.DTSEG1}[4:0] + 1) + (\text{CANFDx\_CHy\_DBTP.DTSEG2}[3:0] + 1) \right] \times \left[(\text{CANFDx\_CHy\_DBTP.DBRP}[4:0] + 1) \times \text{PCLK\_CANFD}[x]\_CLOCK\_CAN[y] \text{ period} \right]$$

Note: The Information Processing Time (IPT) of the CAN FD controller is zero; this means that the data for the next bit is available at the first CAN clock edge after the sample point. Therefore, the IPT does not have to be accounted for when configuring Phase_Seg2, which is the maximum of Phase_Seg1 and the IPT.

24.3.7.2 CAN Bit Rates

The bit rate is the inverse of bit time; therefore, the nominal bit rate is

$$1 / \left[\left(1 + (\text{CANFDx\_CHy\_NBTP.NTSEG1}[7:0] + 1) + (\text{CANFDx\_CHy\_NBTP.NTSEG2}[6:0] + 1) \right) \times \left((\text{CANFDx\_CHy\_NBTP.NBRP}[8:0] + 1) \times \text{PCLK\_CANFD}[x]\_CLOCK\_CAN[y] \text{ period} \right) \right]$$

and the data bit rate is

$$1 / \left[\left(1 + (\text{CANFDx\_CHy\_DBTP.DTSEG1}[4:0] + 1) + (\text{CANFDx\_CHy\_DBTP.DTSEG2}[3:0] + 1) \right) \times \left((\text{CANFDx\_CHy\_DBTP.DBRP}[4:0] + 1) \times \text{PCLK\_CANFD}[x]\_CLOCK\_CAN[y] \text{ period} \right) \right]$$

From these formulas, we can see that the bit rates of the CAN FD controller depends on the CAN clock (PCLK_CANFD[x]_CLOCK_CAN[y]) period, and the range each parameter can be configured to. The following tables list examples of the configurable bit rates at varying CAN clock frequencies. Empty boxes indicate that the desired bit rate cannot be configured at the specified input CAN clock frequency.

Figure 24-15. Example Configuration for Nominal Bit Rates

CAN clock frequency	8MHz	10MHz	16MHz	20MHz	32MHz	40MHz
configuration	Number of tqs per bit time	CANFDx_CHy_NBTP.NBRP + 1	Number of tqs per bit time	CANFDx_CHy_NBTP.NBRP + 1	Number of tqs per bit time	CANFDx_CHy_NBTP.NBRP + 1
nominal bit rate	Number of tqs per bit time	CANFDx_CHy_NBTP.NBRP + 1	Number of tqs per bit time	CANFDx_CHy_NBTP.NBRP + 1	Number of tqs per bit time	CANFDx_CHy_NBTP.NBRP + 1
125Kbps	64tq 32tq 16tq 8tq	1 2 4 8	80tq 40tq 20tq 10tq	1 2 4 8	128tq 64tq 32tq 16tq 8tq	1 2 4 8 16 32
250Kbps	32tq 16tq 8tq	1 2 4	40tq 20tq 10tq	1 2 4 8	128tq 64tq 32tq 16tq 8tq	1 2 4 8 16
500Kbps	16tq 8tq	1 2	20tq 10tq	1 2 4	64tq 32tq 16tq 8tq	1 2 4 8
1Mbps	8tq	1	10tq	1 2	32tq 16tq 8tq	1 2 4

Figure 24-16. Example Configuration for Data Bit Rates

CAN clock frequency	8MHz		10MHz		16MHz		20MHz		32MHz		40MHz	
configuration data bit rate	Number of tqds per bit time	CANFDx_CHy - DBTP-DBRP + 1	Number of tqds per bit time	CANFDx_CHy - DBTP-DBRP + 1	Number of tqds per bit time	CANFDx_CHy - DBTP-DBRP + 1	Number of tqds per bit time	CANFDx_CHy - DBTP-DBRP + 1	Number of tqds per bit time	CANFDx_CHy - DBTP-DBRP + 1	Number of tqds per bit time	CANFDx_CHy - DBTP-DBRP + 1
500Kbps	16tqd 8tqd	1 2	20tqd 10tqd	1 2	32tqd 16tqd 8tqd	1 2 4	40tqd 20tqd 10tqd	1 2 4	32tqd 16tqd 8tqd	2 4 8	40tqd 20tqd 10tqd	2 4 8
1Mbps	8tqd	1	10tqd	1	16tqd 8tqd	1 2	20tqd 10tqd	1 2	32tqd 16tqd 8tqd	1 2 4	40tqd 20tqd 10tqd	1 2 4
2Mbps	-	-	-	-	8tqd	1	10tqd	1	16tqd 8tqd	1 2	20tqd 10tqd	1 2
4Mbps	-	-	-	-	-	-	-	-	8tqd	1	10tqd	1
5Mbps	-	-	-	-	-	-	-	-	-	-	8tqd	1
8Mbps	-	-	-	-	-	-	-	-	-	-	5tqd	1

Note: The user must configure the CAN bit timings to comply with the corresponding CAN standards to ensure proper communication on the CAN bus.

24.4 Message RAM

Message RAM (MRAM) in Traveo II family devices is shared among multiple M_TTCAN channels present in the M_TTCAN group. Refer to the device datasheet for the supported number of M_TTCAN groups, M_TTCAN channels in each group, and the total message RAM allocated to each group. Each M_TTCAN channel in the group can configure its required message RAM according to application requirements.

The message RAM stores RX/TX messages and filter configurations.


Notes:

- The message RAM should be made zero before configuration of the CAN FD controller to prevent bit errors when reading uninitialized words and ECC errors, and to avoid unexpected filter element configurations in the message RAM.
- Unused message RAM cannot be used for general purposes.

24.4.1 Message RAM Configuration

The message RAM has a width of 32 bits. The CAN FD controller can be configured to allocate up to 4480 words in the message RAM (note that the number of words that can be used will be limited by the size of the actual message RAM). It is not necessary to configure each of the sections listed in Figure 24-17, nor is there any restriction with respect to the sequence of the sections.

Figure 24-17. Message RAM Configuration



The CAN FD controller addresses the message RAM in 32-bit words, not single bytes. The configurable start addresses are 32-bit word addresses – only bits 15 to 2 are evaluated, the two least significant bits are ignored.

Notes:


- The CAN FD controller does not check for erroneous configuration of the message RAM. The configuration of the start addresses of different sections and the number of elements of each section should be done carefully to avoid falsification or loss of data.
- Message RAM is accessible by both M_TTCAN and CPU. Dynamic round-robin scheme is implemented to allocate access.

24.4.2 RX Buffer and FIFO Element

An RX buffer and FIFO element is a block of 32-bit words, which holds the data and status of a received frame that was stored in the message RAM.

Up to 64 RX buffers and two RX FIFOs can be configured in the message RAM. Each RX FIFO section can be configured to store up to 64 received messages. The structure of an RX buffer and FIFO element is shown in Figure 24-18. The element size can be configured to store CAN FD messages with up to 64 bytes data field via register CANFDx_CHy_RXESC (RX buffer/FIFO element Size Configuration).

Figure 24-18. RX Buffer and FIFO



R0 [bit31] ESI: Error State Indicator

Bit	Description
0	Transmitting node is error active.
1	Transmitting node is error passive.

R0 [bit30] XTD: Extended Identifier

Signals to the CPU whether the received frame has a standard or extended identifier.

Bit	Description
0	11-bit standard identifier.
1	29-bit extended identifier.

R0 [bit29] RTR: Remote Transmission Request

Signals to the CPU whether the received frame is a data frame or a remote frame.

Bit	Description
0	Received frame is a data frame.
1	Received frame is a remote frame.

Note: There are no remote frames in CAN FD format. In CAN FD frames (FDF = 1), the dominant RRS (Remote Request Substitution) bit replaces bit RTR (Remote Transmission Request).

R0 [bit28:0] ID[28:0]: Identifier

Standard or extended identifier depending on bit XTD. A standard identifier is stored into ID[28:18].

R1 [bit31] ANMF: Accepted Non-matching Frame

Acceptance of non-matching frames may be enabled via CANFDx_CHy_GFC.ANFS[1:0] (Accept Non-matching Frames Standard) and CANFDx_CHy_GFC.ANFE[1:0] (Accept Non-matching Frames Extended).

Bit	Description
0	Received frame matching filter index FIDX.
1	Received frame did not match any RX filter element.

R1 [bit30:24] FIDX[6:0]: Filter Index

FIDX[6:0]	Description
0-127	Index of matching RX acceptance filter element (invalid if ANMF = 1). Range is 0 to List Size Standard/Extended minus 1 (CANFDx_CHy_SIDFC.LSS - 1 resp. CANFDx_CHy_XIDFC.LSE - 1).

R1 [bit23:22] Reserved: Reserved Bits

When writing, always write '0'. The read value is undefined.

R1 [bit21] FDF: Extended Data Length

Bit	Description
0	Classic CAN frame format.
1	CAN FD frame format.

R1 [bit20] BRS: Bit Rate Switch

Bit	Description
0	Frame received without bit rate switching.
1	Frame received with bit rate switching.

R1 [bit19:16] DLC[3:0]: Data Length Code

DLC[3:0]	Description
0-8	Classic CAN + CAN FD: received frame has 0-8 data bytes.
9-15	Classic CAN: received frame has 8 data bytes. CAN FD: received frame has 12/16/20/24/32/48/64 data bytes. See Table 24-1 for details.

R1 [bit15:0] RXTS[15:0]: RX Timestamp

Timestamp Counter value captured on start of frame reception. Resolution depending on configuration of the Shared Timestamp Counter Prescaler CANFDx_TS_CTL.PRESCALE[15:0].

R2 [bit31:24]	DB3[7:0] :	Data Byte 3
R2 [bit23:16]	DB2[7:0] :	Data Byte 2
R2 [bit15:8]	DB1[7:0] :	Data Byte 1
R2 [bit7:0]	DB0[7:0] :	Data Byte 0
R3 [bit31:24]	DB7[7:0] :	Data Byte 7
R3 [bit23:16]	DB6[7:0] :	Data Byte 6
R3 [bit15:8]	DB5[7:0] :	Data Byte 5
R3 [bit7:0]	DB4[7:0] :	Data Byte 4
...	...	...
Rn [bit31:24]	DBm[7:0]:	Data Byte m
Rn [bit23:16]	DBm-1[7:0]:	Data Byte m-1
Rn [bit15:8]	DBm-2[7:0]:	Data Byte m-2
Rn [bit7:0]	DBm-3[7:0]:	Data Byte m-3

Notes:

- Depending on the configuration of the element size (defined by RX buffer/FIFO Element Size Configuration (CANFDx_CHy_RXESC)), Rn will vary from n = 3 to 17.
- m is a function of n, $m = (n - 1) \times 4 - 1$.
- The number of valid data bytes are defined by the Data Length Code.

24.4.3 TX Buffer Element

A TX buffer element is a block of 32-bit words stored in the message RAM that holds data and control information of a frame to be transmitted by the CAN FD controller.

The TX Buffers section can be configured to hold dedicated TX buffers and a TX FIFO/TX Queue. If the TX Buffers section is shared by dedicated TX buffers and a TX FIFO/TX Queue, the dedicated TX buffers start at the beginning of the TX Buffers section followed by the buffers assigned to the TX FIFO or TX Queue. The TX handler distinguishes between dedicated TX buffers and TX FIFO/TX Queue by evaluating the TX buffer configuration CANFDx_CHy_TXBC.TFQS[5:0] (Transmit FIFO/Queue Size) and CANFDx_CHy_TXBC.NDTB[5:0] (Number of Dedicated Transmit Buffers). The element size can be configured to store CAN FD messages with up to 64 bytes data field via register CANFDx_CHy_TXESC (TX Buffer Element Size Configuration).

Figure 24-19. TX Buffer Element

	31	24	23	16	15	8	7	0
T0	ESI	XTD	RTR	ID[28:0]				
T1	MM[7:0]			EFC	reserved	PDF	BRS	DLC[3:0]
T2	DB3[7:0]			DB2[7:0]		DB1[7:0]		DB0[7:0]
T3	DB7[7:0]			DB6[7:0]		DB5[7:0]		DB4[7:0]
...	...			...		...		...
Tn	DBm[7:0]			DBm-1[7:0]		DBm-2[7:0]		DBm-3[7:0]

T0 [bit31] ESI: Error State Indicator

Bit	Description
0	ESI bit in CAN FD format depends only on error passive flag.
1	ESI bit in CAN FD format transmitted recessive.

Note: The ESI bit of the transmit buffer is ORed with the error passive flag to decide the value of the ESI bit in the transmitted FD frame. As required by the CAN FD protocol specification, an error active node may optionally transmit the ESI bit recessive, but an error passive node will always transmit the ESI bit recessive.

T0 [bit30] XTD: Extended Identifier

Bit	Description
0	11-bit standard identifier.
1	29-bit extended identifier.

T0 [bit29] RTR: Remote Transmission Request

Bit	Description
0	Transmit data frame.
1	Transmit remote frame.

Note: When RTR = 1, the CAN FD controller transmits a remote frame according to ISO11898-1, even if FD Operation Enable (CANFDx_CHy_CCCR.FDOE) enables the transmission in CAN FD format.

T0 [bit28:0] ID[28:0]: Identifier

Standard or extended identifier depending on bit XTD. A standard identifier has to be written to ID[28:18].

T1 [bit31:24] MM[7:0]: Message Marker

Written by CPU during TX buffer configuration. Copied into TX Event FIFO element for identification of TX message status.

T1 [bit23] EFC: Event FIFO Control

Bit	Description
0	Don't store TX events.
1	Store TX events.

T1 [bit22] Reserved: Reserved Bit

When writing, always write '0'. The read value is undefined.

T1 [bit21] FDF: FD Format

Bit	Description
0	Frame transmitted in Classic CAN format.
1	Frame transmitted in CAN FD format.

T1 [bit20] BRS: Bit Rate Switching

Bit	Description
0	CAN FD frames transmitted without bit rate switching.
1	CAN FD frames transmitted with bit rate switching.

Note: Bits ESI, FDF, and BRS are only evaluated when CAN FD operation is enabled CANFDx_CHy_CCCR.FDOE = 1. Bit BRS is only evaluated when in addition CANFDx_CHy_CCCR.BRSE = 1. See Table 24-5 for details of bits FDF and BRS.

T1 [bit19:16] DLC[3:0]: Data Length Code

DLC[3:0]	Description
0-8	Classic CAN + CAN FD: transmit frame has 0-8 data bytes.
9-15	Classic CAN: transmit frame has 8 data bytes.

CAN FD: transmit frame has 12/16/20/24/32/48/64 data bytes.

T1 [bit15:0] Reserved: Reserved Bits

When writing, always write '0'. The read value is undefined.

T2 [bit31:24]	DB3[7:0] :	Data Byte 3
T2 [bit23:16]	DB2[7:0] :	Data Byte 2
T2 [bit15:8]	DB1[7:0] :	Data Byte 1
T2 [bit7:0]	DB0[7:0] :	Data Byte 0
T3 [bit31:24]	DB7[7:0] :	Data Byte 7
T3 [bit23:16]	DB6[7:0] :	Data Byte 6
T3 [bit15:8]	DB5[7:0] :	Data Byte 5
T3 [bit7:0]	DB4[7:0] :	Data Byte 4
...	...	...
Tn [bit31:24]	DBm[7:0]:	Data Byte m
Tn [bit23:16]	DBm-1[7:0]:	Data Byte m-1
Tn [bit15:8]	DBm-2[7:0]:	Data Byte m-2
Tn [bit7:0]	DBm-3[7:0]:	Data Byte m-3

Notes:

- Depending on the configuration of the element size (TXESC), Tn will vary from n = 3 to 17.
- m is a function of n: $m = (n - 1) \times 4 - 1$.

24.4.4 TX Event FIFO Element

Each TX Event FIFO Element stores information about transmitted messages. By reading the TX Event FIFO, the CPU gets this information in the order the messages were transmitted. Status information about the TX Event FIFO can be obtained from register CANFDx_CHy_TXEFS (TX Event FIFO Status).

	31	24	23	16	15	8	7	0		
E0	ESI	XTD	RTR	ID[28:0]						
E1	MM[7:0]			ET [1:0]	FDF	BRS	DLC[3:0]		TXTS[15:0]	

E1 [bit20] BRS: Bit Rate Switching

Bit	Description
0	Transmitting node is error active.
1	Transmitting node is error passive.

Bit	Description
0	Frame transmitted without bit rate switching.
1	Frame transmitted with bit rate switching.

E0 [bit30] XTD: Extended Identifier

Bit	Description
0	11-bit standard identifier.
1	29-bit extended identifier.

E1 [bit19:16] DLC[3:0]: Data Length Code

DLC[3:0]	Description
0-8	Classic CAN + CAN FD: frame with 0-8 data bytes transmitted.
9-15	Classic CAN: frame with 8 data bytes transmitted. CAN FD: frame with 12/16/20/24/32/48/64 data bytes transmitted. See Table 24-1 for details.

E0 [bit29] RTR: Remote Transmission Request

Bit	Description
0	Data frame transmitted.
1	Remote frame transmitted.

E1 [bit15:0] TXTS[15:0]: TX Timestamp

Timestamp Counter value captured on start-of-frame transmission. Resolution depending on configuration of the Shared Timestamp Counter Pre-scaler CANFDx_TS_CTL.PRESCALE[15:0].

E0 [bit28:0] ID[28:0]: Identifier

Standard or extended identifier depending on bit XTD. A standard identifier is stored into ID[28:18].

E1 [bit31:24] MM[7:0]: Message Marker

Copied from TX buffer into TX Event FIFO element for identification of TX message status.

E1 [bit23:22] ET[1:0]: Event Type

ET[1:0]	Description
00	Reserved.
01	TX event.
10	Transmission in spite of cancellation. Always set for transmissions in DAR mode (Disable Automatic Retransmission mode).
11	Reserved.

24.4.5 Standard Message ID Filter Element

A Standard Message ID Filter Element consists of a single 32-bit word, and can be configured as a range filter, dual filter, classic bit mask filter, or filter for a single dedicated ID, for messages with 11-bit standard IDs.


Up to 128 filter elements can be configured for 11-bit standard IDs. When accessing a Standard Message ID Filter element, its address is

Filter	List	Standard	Start	Address
(CANFDx_CHy_SIDFC.FLSSA[15:2]) + index of the filter element (0 to 127).				

E1 [bit21] FDF: FD Format

Bit	Description
0	Classic CAN frame format.
1	CAN FD frame format (new DLC-coding and CRC).

Figure 24-21. Standard Message ID Filter



S0 [bit31:30] SFT[1:0]: Standard Filter Type

SFT[1:0]	Description
00	Range filter from SFID1[10:0] to SFID2[10:0] (SFID2[10:0] ≥ received ID ≥ SFID1[10:0]).
01	Dual ID filter for SFID1[10:0] or SFID2[10:0].
10	Classic filter: SFID1[10:0] = filter, SFID2[10:0] = mask. Only those bits of SFID1[10:0] where the corresponding SFID2[10:0] bits are 1 are relevant.
11	Filter element disabled.

Note: With SFT = 11, the filter element is disabled and the acceptance filtering continues. (same behavior as with SFEC = 000)

S0 [bit29:27] SFEC[2:0]: Standard Filter Element Configuration

All enabled filter elements are used for acceptance filtering of standard frames. Acceptance filtering stops at the first matching enabled filter element or when the end of the filter list is reached.

If SFEC[2:0] = 100, 101, or 110 a match sets interrupt flag CANFDx_CHy_IR.HPM (High Priority Message) and, if enabled, an interrupt is generated. In this case register CANFDx_CHy_HPMS (High Priority Message Status) is updated with the status of the priority match.

SFEC[2:0]	Description
000	Disable filter element.
001	Store in RX FIFO 0 if filter matches.
010	Store in RX FIFO 1 if filter matches.
011	Reject ID if filter matches.
100	Set priority if filter matches.
101	Set priority and store in RX FIFO 0 if filter matches.
110	Set priority and store in RX FIFO 1 if filter matches.
111	Store into dedicated RX buffer or as debug message, configuration of SFT[1:0] ignored.

S0 [bit26:16] SFID1[10:0]: Standard Filter ID 1

This bit field has a different meaning depending on the configuration of SFEC[2:0]:

- SFEC[2:0] = 001 to 110

Set SFID1[10:0] according to the SFT[1:0] setting.

- SFEC[2:0] = 111

SFID1[10:0] defines the ID of a standard dedicated RX buffer or debug message to be stored. The received identifiers must match, no masking mechanism is used.

S0 [bit15:11] Reserved: Reserved Bits

When writing, always write '0'. The read value is undefined.

S0 [bit10:0] SFID2[10:0]: Standard Filter ID 2

This bit field has a different meaning depending on the configuration of SFEC[2:0]:

- SFEC[2:0] = 001 to 110

Set SFID2[10:0] according to the SFT[1:0] setting

- SFEC[2:0] = 111

Filter for dedicated RX buffers or for debug messages

SFID2[10:9] decides whether the received message is stored into a dedicated RX buffer or treated as message A, B, or C of the debug message sequence.

SFID2[10:9]	Description
00	Store message into a dedicated RX buffer.
01	Debug Message A.
10	Debug Message B.
11	Debug Message C.

SFID2[8:6] are reserved bits. When writing, always write '0'. The read value is undefined.

SFID2[5:0] defines the offset to the RX buffer Start Address CANFDx_CHy_RXBC.RBSA[15:2] to store a matching message.

Note: Debug message is used to debug on CAN feature.

24.4.6 Extended Message ID Filter Element


An Extended Message ID Filter Element consists of two 32-bit words, and can be configured as a range filter, dual

filter, classic bit mask filter, or filter for a single dedicated ID, for messages with 29-bit extended IDs.

Up to 64 filter elements can be configured for 29-bit extended IDs. When accessing an Extended Message ID Filter element, its address is

Filter List Extended Start Address
(CANFDx_CHy_XIDFC.FLESA[15:2]) + 2 × index of the filter element (0 to 63).

Figure 24-22. Extended Message ID Filter



F0 [bit31:29] EFEC[2:0]: Extended Filter Element Configuration

All enabled filter elements are used for acceptance filtering of extended frames. Acceptance filtering stops at the first matching enabled filter element or when the end of the filter list is reached.

If EFEC[2:0] = 100, 101, or 110 a match sets interrupt flag CANFDx_CHy_IR.HPM (High Priority Message) and, if enabled, an interrupt is generated. In this case register CANFDx_CHy_HPMS (High Priority Message Status) is updated with the status of the priority match.

EFEC[2:0]	Description
000	Disable filter element.
001	Store in RX FIFO 0 if filter matches.
010	Store in RX FIFO 1 if filter matches.
011	Reject ID if filter matches.
100	Set priority if filter matches.
101	Set priority and store in RX FIFO 0 if filter matches.
110	Set priority and store in RX FIFO 1 if filter matches.
111	Store into dedicated RX buffer or as debug message, configuration of EFT[1:0] ignored.

F0 [bit28:0] EFID1[28:0]: Extended Filter ID 1

This bit field has a different meaning depending on the configuration of EFEC[2:0].

- EFEC[2:0] = 001 to 110

Set EFID1[28:0] according to the EFT[1:0] setting.

- EFEC[2:0] = 11

EFID1[28:0] defines the ID of an extended dedicated RX buffer or debug message to be stored. The received identifiers must match, only XIDAM masking mechanism is used.

F1 [bit31:30] EFT[1:0]: Extended Filter Type

EFT[1:0]	Description
00	Range filter from EFID1[28:0] to EFID2[28:0] (EFID2[28:0] ≥ received ID ANDed with XIDAM ≥ EFID1[28:0]).
01	Dual ID filter Matches when EFID1[28:0] or EFID2[28:0] is equal to received ID ANDed with XIDAM.
10	Classic filter: EFID1[28:0] = filter, EFID2[28:0] = mask. Only those bits of EFID1[28:0] where the corresponding EFID2[28:0] bits are 1 are relevant. Matches when the received ID ANDed with XIDAM is equal to EFID1[28:0] masked by EFID2[28:0].
11	Range filter from EFID1[28:0] to EFID2[28:0] (EFID2[28:0] ≥ EFID1[28:0]), XIDAM mask not applied.

F1 [bit29] Reserved: Reserved Bit

When writing, always write '0'. The read value is undefined.

F1 [bit28:0] EFID2[28:0]: Extended Filter ID 2

This bit field has a different meaning depending on the configuration of EFEC[2:0]:

- EFEC[2:0] = 001 to 110

Set EFID2[28:0] according to the EFT[1:0] setting

- EFEC[2:0] = 111

EFID2[28:0] is used to configure this filter for dedicated RX buffers or for debug messages

EFID2[28:11] are reserved bits. When writing, always write '0'. The read value is undefined.

EFID2[10:9] decides whether the received message is stored into a dedicated RX buffer or treated as message A, B, or C of the debug message sequence.

EFID2[10:9]	Description
00	Store message into a dedicated RX buffer.
01	Debug Message A.
10	Debug Message B.
11	Debug Message C.

EFID2[8:6] are reserved bits. When writing, always write '0'. The read value is undefined.


EFID2[5:0] defines the offset to the RX Buffer Start Address CANFDx_CHy_RXBC.RBSA[15:2] to store a matching message.

Note: Debug message is used to debug on CAN feature.

24.4.7 Trigger Memory Element

Up to 64 trigger memory elements can be configured. When accessing a trigger memory element, its address is the Trigger Memory Start Address CANFDx_CHy_TTTMC.TMSA plus the index of the trigger memory element (0...63).

Figure 24-23. Trigger Memory Element



T0 Bit 31:16 TM[15:0]: Time Mark

Cycle time for which the trigger becomes active.

T0 Bit 14:8 CC[6:0]: Cycle Code

Cycle count for which the trigger is valid. Ignored for trigger types Tx_Ref_Trigger, Tx_Ref_Trigger_Gap, Watch_Trigger, Watch_Trigger_Gap, and End_of_List.

CC[6:0]	Description
0b000000x	Valid for all cycles
0b000001c	Valid every second cycle at cycle count mod2 = c
0b00001cc	Valid every fourth cycle at cycle count mod4 = cc
0b0001ccc	Valid every eighth cycle at cycle count mod8 = ccc
0b001cccc	Valid every sixteenth cycle at cycle count mod16 = ccccc
0b01ccccc	Valid every thirty-second cycle at cycle count mod32 = cccccc
0b1cccccc	Valid every sixty-fourth cycle at cycle count mod64 = ccccccc

T0 Bit 7:6 ASC[1:0]: Asynchronous Serial Communication

ASC[1:0]	Description
00	No ASC operation
01	Reserved, do not use
10	Node is ASC receiver
11	Node is ASC transmitter

Note: ASC functionality is not supported in any Traveo II device

T0 Bit 5 TMIN: Time Mark Event Internal

TMIN	Description
0	No Action
1	CANFDx_CHy_TTIR.TTMI is set when trigger memory element becomes active

T0 Bit 4 TMEX: Time Mark Event External

TMEX	Description
0	No Action
1	Pulse at output of Trigger Time Mark with the length of one PCLK_CANFD[x]_CLOCK_CAN[y] period is generated when the time mark of the trigger memory element becomes active and CANFDx_CHy_TTOCN.TTMIE = 1

T0 Bit 3:0 TYPE[3:0]: Trigger Type

TYPE [3:0]	Description
0000	Tx_Ref_Trigger - valid when not in gap
0001	Tx_Ref_Trigger_Gap - valid when in gap
0010	Tx_Trigger_Single - starts a single transmission in an exclusive time window
0011	Tx_Trigger_Continuous - starts continuous transmission in an exclusive time window
0100	Tx_Trigger_Arbitration - starts a transmission in an arbitrating time window
0101	Tx_Trigger_Merged - starts a merged arbitration window
0110	Watch_Trigger - valid when not in gap
0111	Watch_Trigger_Gap - valid when in gap
1000	Rx_Trigger - check for reception
1001	Time_Base_Trigger - only control TMIN, TMEX, and ASC
1010 ... 1111	End_of_List - illegal type, causes configuration error

Notes:

- For ASC operation (ASC = 10, 11) only trigger types Rx_Trigger and Time_Base_Trigger should be used.
- ASC operation is not supported in this device.

T1 Bit 23 FTYPE: Filter Type

FTYPE	Description
0	11-bit standard message ID
1	29-bit extended message ID

T1 Bit 22:16 MNR[6:0]: Message Number

Transmission: Trigger is valid for configured TX buffer number. Valid values are 0 to 31.

Reception: Trigger is valid for standard/extended message ID filter element number. Valid values are 0 to 63 and 0 to 127.

T1 Bits 2:0 MSC[2:0]: Message Status Count

Counts scheduling errors for periodic messages in exclusive time windows. It has no function for arbitrating messages and in event-driven CAN communication (ISO 11898-1:2015).

Notes:

- The trigger memory elements should be written when the M_TTCAN is in INIT state. Write access to the trigger memory elements outside INIT state is not allowed.
- There is an exception for TMIN and TMEX when they are defined as part of a trigger memory element of TYPE Tx_Ref_Trigger. In this case they become active at the time mark modified by the actual Reference Trigger Offset (CANFDx_CHy_TTOST.RTO).

24.4.8 ECC for Message RAM

The error correcting code (ECC) function of the message RAM enables detection and correction of the data errors in message RAM. Code uses a 7-bit parity for a 32-bit data word for the ECC functionality. It has the following features:

- Single error correction and double error detection (SEDED)
 - Single-bit error correction in memory data word
 - Detection of single- and double-bit errors in memory data word
 - Detection of error in message RAM address decoding
- Stopping CAN FD function upon detecting a double-bit (non-correctable) error
- Error injection while data transfer

The following sections describe types of ECC errors.

24.4.8.1 Correctable ECC Error

When a correctable ECC error is detected, the following will happen:

- The corrected data is returned to the read access master
 - No error signal is returned to the master
- The corrected data is written back to the message RAM, unless that write is canceled by another write to the same address.
- The ECC error is reported in the fault structures as a correctable error with the following information:
 - DATA0[15:0]: Violating address
 - DATA0[22:16]: ECC syndrome[6:0]
 - DATA0[27:24]: Master ID: 0-7 = CAN channel ID, 8 = AHB I/F

Note that for security reasons the violating data is not reported in the fault structure.

In the unlikely event of two correctable ECC errors too close together (before fault reporting and correction write back are both complete) the second ECC error is neither corrected nor reported in the fault structure. If later the same address is read again the correction and fault reporting will be done at that time.

More details are available in [Fault Subsystem chapter on page 162](#).

24.4.8.2 Non-correctable ECC Error

When data is read from message RAM and upon ECC check double-bit error is detected, the following actions are taken:

- An error is reported to master
 - In case of AHB master, bus error will occur

- In case of M_TTCAN channel, channel will shut down immediately (CANFDx_CHy_CCCR.INIT = 1)
- Interrupt (BEU) is raised
- ECC error is reported to the fault structure as non-correctable error with the following information:
 - DATA0[15:0]: Violating address
 - DATA0[22:16]: ECC syndrome[6:0]
 - DATA0[27:24]: Master ID: 0-7 = CAN channel ID, 8 = AHB I/F

For security reasons, data is not reported to the fault structure

Unlike single-bit (correctable) error, double-bit error is reported at both the master and fault structures, because read access master needs to know that the read data is not correct and fault structure needs to know all ECC errors.

ECC errors on the address bits are always non-correctable.

24.4.8.3 Address Error

An address error is detected when either M_TTCAN channel or MCU is trying to access an out of range message RAM address (address >= MRAM_SIZE). This feature is added to make software debugging easier. Address error is independent of ECC; that is, it works even if ECC is disabled. When such an address error is detected the following will happen:

- For writes, the error is not reported back to the master
 - Writes are posted and both the AHB interface and the CAN channels ignore error signaling
- For reads from a M_TTCAN channel, the error is reported back to the channel as if it is a non-correctable ECC error; this will result in the following:
 - To prevent corrupt data from being sent, channel will be shutdown (CANFDx_CHy_CCCR.INIT=1) immediately
 - An interrupt is raised (BEU)
- For reads from the AHB interface the address error results in a bus error
- For any case, read, write, and any master, the address error will be reported in the fault structure as a non-correctable error with the following information:
 - DATA0[15:0]: Violating address.
 - DATA0[27:24]: Master ID: 0-7 = CAN channel ID, 8 = AHB I/F
 - DATA0[30]: Set for a write access and cleared for a read access
 - DATA0[31]: Set to flag an address error

24.4.8.4 ECC Error Injection

For safety of the functionality, the ECC error injection feature is added to enable the software to write the ECC bits. With this feature the software can trigger correctable or non-correctable ECC errors to verify that all related hardware and software are functioning properly.

Using this feature, software is able to inject ECC error in the background while data is being fetched from the message RAM.

This feature consists of:

- An enable bit (CANFDx_ECC_ERR_INJ.ERR_EN) to enable the ECC error injection logic
- A 7-bit error parity (CANFDx_ECC_ERR_INJ.ERR_PAR) to be written instead of the generated ECC bits
- An address (CANFDx_ECC_ERR_INJ.ERR_ADDR) to specify at which message RAM address the ECC error injection is done

When a write is done to the specified error injection address, the specified error parity will be used instead of the ECC parity generated by the ECC logic. By limiting this to just one address the software can run this functional test without affecting any other message RAM accesses.

If a write back is done when a correctable error occurs in the specified error injection address, the ECC parity generated by the ECC logic will be used.

As described in the preceding sections, detection of a non-correctable error will result in either shutting down a CAN channel or shutting down a CPU (bus error). Therefore, reporting a non-correctable error back to the master will be suppressed for the targeted error injection address. This feature is necessary to allow non-correctable errors to be verified without affecting the running application.

Note that this error suppression applies to both ECC non-correctable errors and address errors. The software can easily disable this error suppression by disabling the error injection logic (CANFD_ECC_ERR_INJ.ERR_EN = 0).

24.4.8.5 ECC Parity Generation by software

To inject the ECC Error for fault generation, ECC parity must be generated by software. Follow this procedure to generate a 7-bit ECC parity.

```
CODEWORD_SW[63:0] = 64{1'b0};
CODEWORD_SW[31:0] = ACTUALWORD[31:0];
CODEWORD_SW[63:32] = ADDR[31:0];
CODEWORD_SW[63:57] = {7{1'b0}};
```

```
ECC_P0_SW = 64b00000011_01111111_00110110_11011011_00100010_01010100_00101010_10101011;
ECC_P1_SW = 64b00000101_10111101_11101011_01011010_01000100_10011001_01001101_00110101;
ECC_P2_SW = 64b00001001_11011101_11011100_11101110_00001000_11100010_01110001_11000110;
ECC_P3_SW = 64b00010001_11101110_10111011_10101001_10001111_00000011_10000001_11111000;
ECC_P4_SW = 64b00100001_11110110_11010111_01110101_11110000_00000011_11111110_00000000;
ECC_P5_SW = 64b01000001_11111011_01101101_10110100_11111111_11111100_00000000_00000000;
ECC_P6_SW = 64b10000001_00000011_11111111_11111000_00010001_00101100_10010110_01011111;
```

As shown here, Reduction XOR of the ANDed result of CODEWORD_SW[63:0] and respective ECC constants will give a single parity bit.

```
parity[0] = ^ (CODEWORD_SW[63:0] & ECC_P0_SW)
parity[1] = ^ (CODEWORD_SW[63:0] & ECC_P1_SW)
...
parity[6] = ^ (CODEWORD_SW[63:0] & ECC_P6_SW)
```

parity[6:0] gives seven bits parity for 32 bits ACTUALWORD[31:0].

24.4.9 Message RAM OFF

Message RAM can be turned off to save power by setting CANFDx_CTL.MRAM_OFF bit. Default value of this bit is '0' and message RAM is retained in this configuration during DeepSleep power mode.

All the M_TTCAN channels must be powered down before setting CANFDx_CTL.MRAM_OFF bit. See [Power Down \(Sleep Mode\) on page 334](#) to power down the M_TTCAN channels. When message RAM is OFF, any access to message RAM may raise Address Error (MRAM_SIZE = 0).

After switching the message RAM on again, software needs to allow a certain power-up time before message RAM can be used, that is, before STOP_REQ can be de-asserted. Check the RAM_PWR_DELAY_CTL register to see the required time for message RAM power up process.

24.4.10 RAM Watchdog (RWD)

The RAM watchdog monitors the READY output of the Message RAM. A Message RAM access starts the Message RAM Watchdog Counter with the value configured by CANFDx_CHy_RWD.WDC. The counter is reloaded with CANFDx_CHy_RWD.WDC when the Message RAM signals successful completion by activating its READY output. In case there is no response from the Message RAM until the counter has counted down to zero, the counter stops and interrupt flag CANFDx_CHy_IR.WDI is set. The RAM Watchdog Counter is clocked by the Host clock (CLK_SYS). Refer to the Registers TRM for more information about the CANFDx_CHy_RWD register.

24.5 TTCAN Operation

24.5.1 Reference Message

A reference message is a data frame characterized by a specific CAN identifier. It is received and accepted by all nodes except the time master (sender of the reference message).

For Level 1, the data length must be at least one. For Level 0 and Level 2, the data length must be at least four; otherwise, the message is not accepted as a reference message. The reference message may be extended by other data up to the sum of eight CAN data bytes. All bits of the identifier except the three LSBs characterize the message as a reference message. The last three bits specify the priorities of up to eight potential time masters. Reserved bits are transmitted as logical 0 and are ignored by the receivers. The reference message is configured using the CANFDx_CHy_TTRMC register.

The time master transmits the reference message. If the reference message is disturbed by an error, it is retransmitted immediately. In a retransmission, the transmitted Master_Ref_Mark is updated. The reference message is sent periodically, but it is allowed to stop the periodic transmission (Next_is_Gap bit). It can initiate event-synchronized transmission at the start of the next basic cycle by the current time master or by one of the other potential time masters.

The node transmitting the reference message is the current time master. The time master is allowed to transmit other messages. If the current time master fails, its function is replicated by the potential time master with the highest priority. Nodes that are neither time master nor potential time master are time-receiving nodes.

24.5.1.1 Level 1

Level 1 operation is configured via CANFDx_CHy_TTOCF.OM = 01 and CANFDx_CHy_TTOCF.GEN. External clock synchronization is not available in Level 1.

The information related to the reference message is stored in the first data byte as shown in Table 24-8. Cycle_Count is optional.

Table 24-8. First Byte of Level 1 Reference Message

Bits	7	6	5	4	3	2	1	0
First Byte	Next_is_Gap	Reserved	Cycle_Count [5:0]					

24.5.1.2 Level 2

Level 2 operation is configured via CANFDx_CHy_TTOCF.OM = 10 and CANFDx_CHy_TTOCF.GEN.

The information related to the reference message is stored in the first four data bytes as shown in Table 24-9. Cycle_Count and the lower four bits of NTU_Res are optional. The M_TTCAN does not evaluate NTU_Res[3:0] from received reference messages, it always transmits these bits as zero.

Table 24-9. First Four Bytes of Level 2 Reference Message

Bits	7	6	5	4	3	2	1	0
First Byte	Next_is_Gap	Reserved	Cycle_Count[5:0]					
Second Byte	NTU_Res[6:4]			NTU_Res[3:0]				Disc_Bit
Third Byte	Master_Ref_Mark[7:0]							
Fourth Byte	Master_Ref_Mark[15:8]							

24.5.1.3 Level 0

Level 0 operation is configured via CANFDx_CHy_TTOCF.OM = 11. External event-synchronized time-triggered operation is not available in Level 0.

The information related to the reference message is stored in the first four data bytes as shown in Table 24-10. In Level 0, Next_is_Gap is always zero. Cycle_Count and the lower four bits of NTU_Res are optional. The M_TTCAN does not evaluate NTU_Res[3:0] from received reference messages; it always transmits these bits as zero.

Table 24-10. First four Bytes of Level 0 Reference Message

Bits	7	6	5	4	3	2	1	0
First Byte	Nex- t_is_Gap	Reserved	Cycle_Count[5:0]					
Second Byte	NTU_Res[6:4]				NTU_Res[3:0]			Disc_Bit
Third Byte	Master_Ref_Mark[7:0]							
Fourth Byte	Master_Ref_Mark[15:8]							

24.5.2 TTCAN Configuration

24.5.2.1 TTCAN Timing

The Network Time Unit (NTU) is the unit in which all times are measured. The NTU is a constant of the whole network and is defined by the network system designer. In TTCAN Level 1 the NTU is the nominal CAN bit time. In TTCAN Level 0 and Level 2 the NTU is a fraction of the physical second.

The NTU is the time base for the local time. The integer part of the local time (16-bit value) is incremented once for each NTU. Cycle time and global time are both derived from local time. The fractional part (3-bit value) of local time, cycle time, and global time is not readable.

In TTCAN Level 0 and Level 2, the length of the NTU is defined by the Time Unit Ratio (TUR). The TUR is a non-integer number given by the formula $TUR = CANFDx\_CHy\_TURNA.NAV / CANFDx\_CHy\_TURCF.DC$. The length of the NTU is given by the formula $NTU = CAN\ Clock\ Period \times TUR$.

The TUR Numerator Configuration NC is an 18-bit number, $CANFDx\_CHy\_TURCF.NCL[15:0]$ can be programmed in the range 0x0000-0xFFFF. $CANFDx\_CHy\_TURCF.NCH[17:16]$ is hard-wired to 0b01. When the number 0xnnnn is written to $CANFDx\_CHy\_TURCF.NCL[15:0]$, $CANFDx\_CHy\_TURNA.NAV$ starts with the value $0x10000 + 0x0nnnn = 0x1nnnn$. The TUR Denominator Configuration $CANFDx\_CHy\_TURCF.DC$ is a 14-bit number. $CANFDx\_CHy\_TURCF.DC$ may be programmed in the range 0x0001 - 0x3FFF; 0x0000 is an illegal value.

In Level 1, NC must be $\geq 4 \times CANFDx\_CHy\_TURCF.DC$. In Level 0 and Level 2 NC must be $\geq 8 \times CANFDx\_CHy\_TURCF.DC$ to allow the 3-bit resolution for the internal fractional part of the NTU.

A hardware reset presets $CANFDx\_CHy\_TURCF.DC$ to 0x1000 and $CANFDx\_CHy\_TURCF.NCL$ to 0x10000, resulting in an NTU consisting of 16 CAN clock periods. Local time and application watchdog are not started before either the $CANFDx\_CHy\_CCCR.INIT$ is reset or $CANFDx\_CHy\_TURCF.ELT$ is set. $CANFDx\_CHy\_TURCF.ELT$ may not be set before the NTU is configured. Setting $CANFDx\_CHy\_TURCF.ELT$ to '1' also locks the write access to register $CANFDx\_CHy\_TURCF$.

At startup $CANFDx\_CHy\_TURNA.NAV$ is updated from NC ($= CANFDx\_CHy\_TURCF.NCL + 0x10000$) when $CANFDx\_CHy\_TURCF.ELT$ is set. In TTCAN Level 1 there is no drift compensation. $CANFDx\_CHy\_TURNA.NAV$ does not change during operation, it always equals NC.

In TTCAN Level 0 and Level 2, there are two possibilities for $CANFDx\_CHy\_TURNA.NAV$ to change. When operating as time slave or backup time master, and when $CANFDx\_CHy\_TTOCF.ECC$ is set, $CANFDx\_CHy\_TURNA.NAV$ is updated automatically to the value calculated from the monitored global time speed, as long as the M_TTCAN is in synchronization states In_Schedule or In_Gap. When it loses synchronization, it returns to NC. When operating as the actual time master, and when $CANFDx\_CHy\_TTOCF.EECS$ is set, the host may update $CANFDx\_CHy\_TURCF.NCL$. When the host sets $CANFDx\_CHy\_TTOCN.ECS$, $CANFDx\_CHy\_TURNA.NAV$ will be updated from the new value of NC at the next reference message. The status flag $CANFDx\_CHy\_TTOST.WECS$ is set when $CANFDx\_CHy\_TTOCN.ECS$ is set and is cleared when $CANFDx\_CHy\_TURNA.NAV$ is updated. $CANFDx\_CHy\_TURCF.NCL$ is write-locked while $CANFDx\_CHy\_TTOST.WECS$ is set.

In TTCAN Level 0 and Level 2, the clock calibration process adapts $CANFDx\_CHy\_TURNA.NAV$ in the range of the synchronization deviation limit (SDL) of $NC \pm 2 (CANFDx\_CHy\_TTOCF.LDSDL + 5)$.

$CANFDx\_CHy\_TURCF.NCL$ should be programmed to the largest applicable numerical value to achieve the best accuracy in the calculation of $CANFDx\_CHy\_TURNA.NAV$.

The synchronization deviation (SD) is the difference between NC and $CANFDx\_CHy\_TURNA.NAV$ ($SD = |NC - CANFDx\_CHy\_TURNA.NAV|$). It is limited by the SDL, which is configured by its dual logarithm $CANFDx\_CHy\_TTOCF.LDSDL$ ($SDL = 2 (CANFDx\_CHy\_TTOCF.LDSDL + 5)$) and should not exceed the clock tolerance given by the CAN bit timing configuration. SD is calculated at each new basic cycle. When the calculated $CANFDx\_CHy\_TURNA.NAV$ deviates by more than SDL from NC, or if the Disc_Bit in the reference message is set, the drift compensation is suspended, $CANFDx\_CHy\_TTIR.GTE$ is set, and $CANFDx\_CHy\_TTOSC.QCS$ is reset; if Disc_Bit = '1', $CANFDx\_CHy\_TTIR.GTD$ is set.

TUR configuration examples are shown in Table 24-11.

Table 24-11. TUR Configuration Examples

TUR	8	10	24	50	510	125000	32.5	100/12	529/17
NC	0x1FFF8	0x1FFFE	0x1FFF8	0x1FFEA	0x1FFFE	0x1FFE0	0x1FFE0	0x19000	0x10880
CANFDx- CHy_TURCF.DC	0x3FFF	0x3333	0x1555	0x0A3D	0x0101	0x0001	0x0FC0	0x3000	0x0880

CANFDx_CHy_TTOCN.ECS schedules NC for activation by the next reference message. CANFDx_CHy_TTOCN.SGT schedules CANFDx_CHy_TTGTP.TP for activation by the next reference message. Setting of CANFDx_CHy_TTOCN.ECS and CANFDx_CHy_TTOCN.SGT requires CANFDx_CHy_TTOCF.EECS to be set (external clock synchronization enabled) while the M_TTCAN is actual time master.

The M_TTCAN module provides an application watchdog to verify the function of the application program. The host has to serve this watchdog regularly; otherwise, all CAN bus activity is stopped. The Application Watchdog Limit CANFDx_CHy_TTOCF.AWL specifies the number of NTUs the watchdog has to be served. The maximum number of NTUs is 256. The Application Watchdog is served by reading register CANFDx_CHy_TTOST. CANFDx_CHy_TTOST.AWE indicates whether the watchdog is served in time. In case the application failed to serve the application watchdog, interrupt flag CANFDx_CHy_TTIR.AW is set. For software development, the application watchdog may be disabled by programming CANFDx_CHy_TTOCF.AWL to 0x00 (see [24.3.1.10 Application Watchdog](#)).

24.5.2.2 Message Scheduling

CANFDx_CHy_TTOCF.TM controls whether the M_TTCAN operates as a potential time master or as a time slave. If it is a potential time master, the three LSBs of the reference message identifier CANFDx_CHy_TTRMC.RID define the master priority, 0 being the highest and 7 the lowest priority. Two nodes in the network may not use the same master priority. CANFDx_CHy_TTRMC.RID is used for recognition of reference messages. CANFDx_CHy_TTRMC.RMPS is not relevant for time slaves.

The Initial Reference Trigger Offset CANFDx_CHy_TTOCF.IRTO is a 7-bit-value that defines (in NTUs) how long a backup time master waits before it starts the transmission of a reference message, when a reference message is expected but the bus remains idle. The recommended value for CANFDx_CHy_TTOCF.IRTO is the master priority multiplied with a factor depending on the expected clock drift between the potential time masters in the network. The sequential order of the backup time masters, when one of them starts the reference message if the current time master fails, should correspond to their master priority, even with maximum clock drift.

CANFDx_CHy_TTOCF.OM decides whether the node operates in TTCAN Level 0, Level 1, or Level 2. In one network, all potential time masters should operate on the same level. Time slaves may operate on Level 1 in a Level 2 network, but not vice versa. The configuration of the TTCAN operation mode via CANFDx_CHy_TTOCF.OM is the last step in the setup. When CANFDx_CHy_TTOCF.OM = 00 (event-driven CAN communication), the M_TTCAN operates according to ISO 11898-1:2015, without time triggers. When

CANFDx_CHy_TTOCF.OM = 01 (Level 1), the M_TTCAN operates according to ISO 11898-4, but without the possibility to synchronize the basic cycles to external events, the Next_is_Gap bit in the reference message is ignored. When CANFDx_CHy_TTOCF.OM = 10 (Level 2), the M_TTCAN operates according to ISO 11898-4, including the event-synchronized start of a basic cycle. When CANFDx_CHy_TTOCF.OM = 11 (Level 0), the M_TTCAN operates as event-driven CAN but maintains a calibrated global time base similar to Level 2.

CANFDx_CHy_TTOCF.EECS enables the external clock synchronization, allowing the application program of the current time master to update the TUR configuration during time-triggered operation, to adapt the clock speed and (in Level 0,2 only) the global clock phase to an external reference.

CANFDx_CHy_TTMLM.ENTT in the TT Matrix Limits register specifies the number of expected Tx_Triggers in the system matrix. This is the sum of Tx_Triggers for exclusive single arbitrating and merged arbitrating windows, excluding the Tx_Ref_Triggers. Note that this is usually not the number of Tx_Trigger memory elements; the number of basic cycles in the system matrix and the trigger's repeat factors must be taken into account. An inaccurate configuration of CANFDx_CHy_TTMLM.ENTT will result in either a TX Count Underflow (CANFDx_CHy_TTIR.TXU = 1 and CANFDx_CHy_TTOST.EL = 01, severity 1) or in a TX Count Overflow (CANFDx_CHy_TTIR.TXO = 1 and CANFDx_CHy_TTOST.EL = 10, severity 2).

Note: In case the first reference message seen by a node does not have Cycle_Count zero, this node may finish its first matrix cycle with its TX count resulting in a TX Count Underflow condition. As long as a node is in state, synchronizing its Tx_Triggers will not lead to transmissions.

CANFDx_CHy_TTMLM.CCM specifies the number of the last basic cycle in the system matrix. The counting of basic cycles starts at 0. In a system matrix consisting of eight basic cycles CANFDx_CHy_TTMLM.CCM would be 7. CANFDx_CHy_TTMLM.CCM is ignored by time slaves, a receiver of a reference message considers the received cycle count as the valid cycle count for the actual basic cycle.

CANFDx_CHy_TTMLM.TXEW specifies the length of the TX enable window in NTUs. The TX enable window is the period at the beginning of a time window where a transmission may be started. If the sample point of the first bit of a transmit message is not inside the TX enable window, the transmission cannot be started in that time window at all. An example is because of an overlap from the previous time window's message. CANFDx_CHy_TTMLM.TXEW should be chosen based on the network's synchronization quality and the relation between the length of the time windows and the length of messages.

24.5.2.3 Trigger Memory

The trigger memory is part of the message RAM. It stores up to 64 trigger elements. A trigger memory element consists of Time Mark TM, Cycle Code CC, Trigger Type TYPE, Filter Type FTYPE, Message Number MNR, Message Status Count MSC, Time Mark Event Internal TMIN, Time Mark Event External TMEX, and Asynchronous Serial Communication ASC (see [24.4.7 Trigger Memory Element](#)).

The time mark defines at which cycle time a trigger becomes active. The trigger elements in the trigger memory must be sorted by their time marks. The trigger element with the lowest time mark is written to the first trigger memory word. Message number and cycle code are ignored for triggers of type Tx_Ref_Trigger, Tx_Ref_Trigger_Gap, Watch_Trigger, Watch_Trigger_Gap, and End_of_List.

When the cycle time reaches the time mark of the actual trigger, the FSE switches to the next trigger and starts to read it from the trigger memory. For a transmit trigger, the TX handler starts to read the message from the message RAM as soon as the FSE switches to its trigger. The RAM access speed defines the minimum time step between a transmit trigger and its preceding trigger, the TX handler should be able to prepare the transmission before the transmit trigger's time mark is reached. The RAM access speed also limits the number of non-matching (with regard to their cycle code) triggers between two matching triggers, the next matching trigger must be read before its time mark is reached. If the reference message is n NTU long, a trigger with a time mark less than n will never become active and will be treated as a configuration error.

The starting point of cycle time is the sample point of the reference message's start-of-frame bit. The next reference message is requested when cycle time reaches the Tx_Ref_Trigger's time mark. The M_TTCAN reacts to the transmission request at the next sample point. A new Sync_Mark is captured at the start-of-frame bit, but the cycle time is incremented until the reference message is successfully transmitted (or received) and the Sync_Mark is taken as the new Ref_Mark. At that point, cycle time is restarted. As a consequence, cycle time can never (with the exception of initialization) be seen at a value less than n , with n being the length of the reference message measured in NTU.

Length of a basic cycle: Tx_Ref_Trigger time mark + 1 NTU + 1 CAN bit time

The trigger list will be different for all nodes in the TTCAN network. Each node knows only the Tx_Triggers for its own transmit messages, the Rx_Triggers for the receive messages that are processed by this node, and the triggers concerning the reference messages.

Trigger Types

Tx_Ref_Trigger (TYPE = 0000) and Tx_Ref_Trigger_Gap (TYPE = 0001) cause the transmission of a reference

message by a time master. A configuration error (CANFDx_CHy_TTOST.EL = 11, severity 3) is detected when a time slave encounters a Tx_Ref_Trigger(_Gap) in its trigger memory. Tx_Ref_Trigger_Gap is only used in external event-synchronized time-triggered operation mode. In that mode, Tx_Ref_Trigger is ignored when the M_TTCAN synchronization state is In_Gap (CANFDx_CHy_TTOST.SYS = 10).

Tx_Trigger_Single (TYPE = 0010), Tx_Continuous (TYPE = 0011), Tx_Trigger_Arbitration (TYPE = 0100), and Tx_Trigger_Merged (TYPE = 0101) cause the start of a transmission. They define the start of a time window.

Tx_Trigger_Single starts a single transmission in an exclusive time window when the message buffer's Transmission Request Pending bit is set. After successful transmission, the Transmission Request Pending bit is reset.

Tx_Trigger_Continuous starts a transmission in an exclusive time window when the message buffer's transmission Request Pending bit is set. After successful transmission, the Transmission Request Pending bit remains set, and the message buffer is transmitted again in the next matching time window.

Tx_Trigger_Arbitration starts an arbitrating time window, Tx_Trigger_Merged a merged arbitrating time window. The last Tx_Trigger of a merged arbitrating time window must be of type Tx_Trigger_Arbitration. A Configuration Error (CANFDx_CHy_TTOST.EL = 11, severity 3) is detected when a trigger of type Tx_Trigger_Merged is followed by any other Tx_Trigger than one of type Tx_Trigger_Merged or Tx_Trigger_Arbitration. Several Tx_Triggers may be defined for the same TX message buffer. Depending on their cycle code, they may be ignored in some basic cycles. The cycle code should be considered when the expected number of Tx_Triggers (CANFDx_CHy_TTMLM.ENTT) is calculated.

Watch_Trigger (TYPE = 0110) and Watch_Trigger_Gap (TYPE = 0111) check for missing reference messages. They are used by both time masters and time slaves. Watch_Trigger_Gap is only used in external event-synchronized time-triggered operation mode. In that mode, a Watch_Trigger is ignored when the M_TTCAN synchronization state is In_Gap (CANFDx_CHy_TTOST.SYS = 10).

Rx_Trigger (TYPE = 1000) is used to check for the reception of periodic messages in exclusive time windows. Rx_Triggers are not active until state In_Schedule or In_Gap is reached. The time mark of an Rx_Trigger should be placed after the end of that message transmission, independent of time window boundaries. Depending on their cycle code, Rx_Triggers may be ignored in some basic cycles. At the Rx_Trigger time mark, it is checked whether the last received message before this time mark and after start of cycle or previous Rx_Trigger matches the acceptance filter element referenced by MNR. Accepted

messages are stored in one of two receive FIFOs, according to the acceptance filtering, independent of the Rx_Trigger. Acceptance filter elements that are referenced by Rx_Triggers should be placed at the beginning of the filter list to ensure that the filtering is finished before the Rx_Trigger time mark is reached.

Time_Base_Trigger (TYPE = 1001) is used to generate internal/external events depending on the configuration of ASC, TMIN, and TMEX.

End_of_List (TYPE = 1010...1111) is an illegal trigger type, a configuration error (CANFDx_CHy_TTOST.EL = 11, severity 3) is detected when an End_of_List trigger is encountered in the trigger memory before the Watch_Trigger or Watch_Trigger_Gap.

Restrictions for the Node's Trigger List

Two triggers may not be active at the same cycle time and cycle count, but triggers that are active in different basic cycles (different cycle code) may share the same time mark.

Rx_Triggers and Time_Base_Triggers may not be placed inside the TX enable windows of Tx_Trigger_Single/Continuous/Arbitration, but they may be placed after Tx_Trigger_Merged.

Triggers that are placed after the Watch_Trigger (or the Watch_Trigger_Gap when CANFDx_CHy_TTOST.SYS = 10) will never become active. The watch triggers themselves will not become active when the reference messages are transmitted on time.

All unused trigger memory words (after the Watch_Trigger or after the Watch_Trigger_Gap when CANFDx_CHy_TTOST.SYS = 10) must be set to trigger type End_of_List.

A typical trigger list for a potential time master will begin with a number of Tx_Triggers and Rx_Triggers followed by the Tx_Ref_Trigger and Watch_Trigger. For networks with external event-synchronized time-triggered communication, this is followed by the Tx_Ref_Trigger_Gap and the Watch_Trigger_Gap. The trigger list for a time slave will be the same but without the Tx_Ref_Trigger and the Tx_Ref_Trigger_Gap.

At the beginning of each basic cycle, that is at each reception or transmission of a reference message, the trigger list is processed starting with the first trigger memory element. The FSE looks for the first trigger with a cycle code that matches the current cycle count. The FSE waits until cycle time reaches the trigger's time mark and activates the trigger. Later, the FSE looks for the next trigger in the list with a cycle code that matches the current cycle count.

Special consideration is needed for the time around Tx_Ref_Trigger and Tx_Ref_Trigger_Gap. In a time master competing for master ship, the effective time mark of a Tx_Ref_Trigger may be decremented to be the first node to

start a reference message. In backup time masters the effective time mark of a Tx_Ref_Trigger or Tx_Ref_Trigger_Gap is the sum of its configured time mark and the Reference Trigger Offset CANFDx_CHy_TTOCF.IRTO. If error level 2 is reached (CANFDx_CHy_TTOST.EL = 10), the effective time mark is the sum of its time mark and 0x127. No other trigger elements should be placed in this range; otherwise, the time marks may appear out of order and are flagged as a configuration error. Trigger elements that are coming after Tx_Ref_Trigger may never become active as long as the reference messages come in time.

There are interdependencies between the following parameters:

- Host clock frequency
- Speed and waiting time for Trigger RAM accesses
- Length of the acceptance filter list
- Number of trigger elements
- Complexity of cycle code filtering in the trigger elements
- Offset between time marks of the trigger elements

Examples of Trigger Handling

The following example shows how the trigger list is derived from a node's system matrix. Assume that node A is a first time master; a section of the system matrix shown in [Table 24-12](#).

Table 24-12. System Matrix Node A

Cycle Count	Time Mark1	Time Mark2	Time Mark3	Time Mark4	Time Mark5	Time Mark6	Time Mark7
0	Tx7					TxRef	Error
1	Rx3		Tx2, Tx4			TxRef	Error
2						TxRef	Error
3	Tx7		Rx5			TxRef	Error
4	Tx7			Rx6		TxRef	Error

The cycle count starts with 0 – 0, 1, 3, 7, 15, 31, 63 (the number of basic cycles in the system matrix is 1, 2, 4, 8, 16, 32, 64). The maximum cycle count is configured by CANFDx_CHy_TTMLM.CCM. The Cycle Code (CC) is composed of repeat factor (value of most significant '1') and the number of the first basic cycle in the system matrix (bit field after most significant '1').

Example: When CC is 0b0010011 (repeat factor: 16, first basic cycle: 3) and maximum cycle count of CANFDx_CHy_TTMLM.CCM = 0x3F, matches occur at cycle counts 3, 19, 35, 51.

A trigger element consists of Time Mark (TM), Cycle Code (CC), Trigger Type (TYPE), and Message Number (MNR). For transmission, MNR references the TX buffer number (0..31). For reception, MNR references the number of the filter element (0..127) that matched during acceptance filtering. Depending on the configuration of the Filter Type

FTYPE, the 11-bit or 29-bit message ID filter list is referenced.

In addition, a trigger element can be configured for Asynchronous Serial Communication (ASC), generation of Time Mark Event Internal (TMIN), and Time Mark Event External (TMEX). The Message Status Count (MSC) holds the counter value (0..7) for scheduling errors for periodic messages in exclusive time windows when the time mark of the trigger element becomes active.

Table 24-13. Trigger List Node A

Trigger	Time Mark TM[15:0]	Cycle Code CC[6:0]	Trigger Type TYPE[3:0]	Mess. No. MNR[6:0]
0	Mark1	0b0000100	Tx_Trigger_Single	7
1	Mark1	0b1000000	Rx_Trigger	3
2	Mark1	0b1000011	Tx_Trigger_Single	7
3	Mark3	0b1000001	Tx_Trigger_Merged	2
4	Mark3	0b1000011	Rx_Trigger	5
5	Mark4	0b1000001	Tx_Trigger_Arbitration	4
6	Mark4	0b1000100	Rx_Trigger	6
7	Mark6	n.a.	Tx_Ref_Trigger	0 (Ref)
8	Mark7	n.a.	Watch_Trigger	n.a.
9	n.a.	n.a.	End_of_List	n.a.

Tx_Trigger_Single, Tx_Trigger_Continuous, Tx_Trigger_Merged, Tx_Trigger_Arbitration, Rx_Trigger, and Time_Base_Trigger are only valid for the specified cycle code. For all other trigger types the cycle code is ignored.

The FSE starts the basic cycle by scanning the trigger list starting from zero until a trigger with time mark that is greater than the cycle time is reached, CC matches the actual cycle count, or a trigger of type Tx_Ref_Trigger, Tx_Ref_Trigger_Gap, Watch_Trigger, or Watch_Trigger_Gap is encountered.

When the cycle time reaches TM, the action defined by TYPE and MNR is started. There is an error in the configuration when it reaches End_of_List.

At Mark6, the reference message (always TxRef) is transmitted. After transmission, the FSE returns to the beginning of the trigger list. When it reaches Watch Trigger at Mark7, the node is unable to transmit the reference message; error treatment is then started.

Detection of Configuration Errors

A configuration error is signaled via CANFDx_CHy_TTOST.EL = 11 (severity 3) when:

- The FSE comes to a trigger in the list with a cycle code that matches the current cycle count but with a time mark that is less than the cycle time.
- The previous active trigger was a Tx_Trigger_Merged and the FSE comes to a trigger in the list with a cycle

code that matches the current cycle count but that is neither a Tx_Trigger_Merged nor a Tx_Trigger_Arbitration nor a Time_Base_Trigger nor an Rx_Trigger.

- The FSE of a node with CANFDx_CHy_TTOCF.TM = 0 (time slave) encounters a Tx_Ref_Trigger or a Tx_Ref_Trigger_Gap.
- Any time mark placed inside the TX enable window (defined by CANFDx_CHy_TTMLM.TXEW) of a Tx_Trigger with a matching cycle code.
- A time mark is placed near the time mark of a Tx_Ref_Trigger and the Reference Trigger Offset CANFDx_CHy_TTOST.RTO causes a reversal of their sequential order measured in cycle time.

24.5.2.4 TTCAN Schedule Initialization

The synchronization to the M_TTCAN message schedule starts when CANFDx_CHy_CCCR.INIT is reset. The M_TTCAN can operate time-triggered (CANFDx_CHy_TTOCF.GEN = 0) or external event-synchronized time-triggered (CANFDx_CHy_TTOCF.GEN = 1). All nodes start with cycle time zero at the beginning of their trigger list with CANFDx_CHy_TTOST.SYS = 00 (out of synchronization); no transmission is enabled with the exception of the reference message. Nodes in external event-synchronized time-triggered operation mode will ignore Tx_Ref_Trigger and Watch_Trigger and use Tx_Ref_Trigger_Gap and Watch_Trigger_Gap instead until the first reference message decides whether a gap is active.

Time Slaves

After configuration, a time slave will ignore its Watch_Trigger and Watch_Trigger_Gap when it does not receive any message before reaching the Watch_Triggers. When it reaches Init_Watch_Trigger, interrupt flag CANFDx_CHy_TTIR.IWT is set, the FSE is frozen, and the cycle time will become invalid. However, the node will still be able to take part in CAN bus communication (to give acknowledge or to send error flags). The first received reference message will restart the FSE and the cycle time.

Note: Init_Watch_Trigger is not part of the trigger list. It is implemented as an internal counter that counts up to 0xFFFF = maximum cycle time.

When a time slave receives any message but the reference message before reaching the Watch_Triggers, it will assume a fatal error (CANFDx_CHy_TTOST.EL = 11, severity 3), set interrupt flag CANFDx_CHy_TTIR.WT, switch off its CAN bus output, and enter the bus monitoring mode (CANFDx_CHy_CCCR.MON set to '1'). In the bus monitoring mode, it is still able to receive messages, but cannot send any dominant bits and therefore, cannot acknowledge.

Note: To leave the fatal error state, the host must set CANFDx_CHy_CCCR.INIT = '1'. After reset of

CANFDx_CHy_CCCR.INIT, the node restarts TTCAN communication.

When no error is encountered during synchronization, the first reference message sets CANFDx_CHy_TTOST.SYS = 01 (synchronizing), the second sets the TTCAN synchronization state (depending on its Next_is_Gap bit) to CANFDx_CHy_TTOST.SYS = 11 (In_Schedule) or CANFDx_CHy_TTOST.SYS = 10 (In_Gap), enabling all Tx_Triggers and Rx_Triggers.

Potential Time Master

After configuration, a potential time master will start the transmission of a reference message when it reaches its Tx_Ref_Trigger (or its Tx_Ref_Trigger_Gap when in external event-synchronized time-triggered operation). It will ignore its Watch_Trigger and Watch_Trigger_Gap when it does not receive any message or transmit the reference message successfully before reaching the Watch_Triggers (the reason assumed is that all other nodes still in reset or configuration and does not acknowledge). When it reaches Init_Watch_Trigger, the attempted transmission is aborted, interrupt flag CANFDx_CHy_TTIR.IWT is set, the FSE is frozen, and the cycle time will become invalid, but the node will still be able to take part in CAN bus communication (to acknowledge or send error flags). Resetting CANFDx_CHy_TTIR.IWT will re-enable the transmission of reference messages until the next time Init_Watch_Trigger condition is met, or another CAN message is received. The FSE will be restarted by the reception of a reference message.

When a potential time master reaches the Watch_Triggers after it has received any message but the reference message, it will assume a fatal error (CANFDx_CHy_TTOST.EL = 11, severity 3), set interrupt flag CANFDx_CHy_TTIR.WT, switch off its CAN bus output, and enter the bus monitoring mode (CANFDx_CHy_CCCR.MON set to '1'). In bus monitoring mode, it is still able to receive messages, but it cannot send any dominant bits and therefore, cannot acknowledge.

When no error is detected during initialization, the first reference message sets CANFDx_CHy_TTOST.SYS = 01 (synchronizing), the second sets the TTCAN synchronization state (depending on its Next_is_Gap bit) to CANFDx_CHy_TTOST.SYS = 11 (In_Schedule) or CANFDx_CHy_TTOST.SYS = 10 (In_Gap), enabling all Tx_Triggers and Rx_Triggers.

A potential time master is current time master (CANFDx_CHy_TTOST.MS = 11) when it is the transmitter of the last reference message; otherwise, it is the backup time master (CANFDx_CHy_TTOST.MS = 10).

When all potential time masters have finished configuration, the node with the highest time master priority in the network will become the current time master.

24.5.3 TTCAN Gap Control

All functions related to gap control apply only when the M_TTCAN is operated in external event-synchronized time-triggered mode (CANFDx_CHy_TTOCF.GEN = 1). In this operation mode the TTCAN message schedule may be interrupted by inserting gaps between the basic cycles of the system matrix. All nodes connected to the CAN network should be configured for external event-synchronized time-triggered operation.

During a gap, all transmissions are stopped and the CAN bus remains idle. A gap is finished when the next reference message starts a new basic cycle. The gap starts at the end of a basic cycle that was started by a reference message with bit Next_is_Gap = '1'; for example, gaps are initiated by the current time master.

The current time master has two options to initiate a gap. A gap can be initiated under software control when the application program writes CANFDx_CHy_TTOCN.NIG = 1. The Next_is_Gap bit will be transmitted as '1' with the next reference message. A gap can also be initiated under hardware control when the application program writes CANFDx_CHy_TTOCN.GCS = 1. When a reference message is started and CANFDx_CHy_TTOCN.GCS is set, Next_is_Gap = '1' will be set.

As soon as that reference message is completed, the CANFDx_CHy_TTOST.WFE bit will announce the gap to the time master and slaves. The current basic cycle will continue until its last time window. The time after the last time window is the gap time.

For the actual time master and the potential time masters, CANFDx_CHy_TTOST.GSI will be set when the last basic cycle has finished and the gap time starts. In nodes that are time slaves, the CANFDx_CHy_TTOST.GSI bit will remain at '0'.

When a potential time master is in synchronization state In_Gap (CANFDx_CHy_TTOST.SYS = 10), it has four options to intentionally finish a gap:

- Under software control by writing CANFDx_CHy_TTOCN.FGP = 1.
- Under hardware control (CANFDx_CHy_TTOCN.GCS = 1), CANFDx_CHy_TTOCN.FGP will automatically be set when an edge from HIGH to LOW at the internal event trigger input pin is detected and restarts the schedule.
- The third option is a time-triggered restart. When CANFDx_CHy_TTOCN.TMG = 1, the next register time mark interrupt (CANFDx_CHy_TTIR.RTMI = 1) will set CANFDx_CHy_TTOCN.FGP and start the reference message.
- Any potential time master will finish a gap when it reaches its Tx_Ref_Trigger_Gap, assuming that the event to synchronize to did not occur on time.

None of these options can cause a basic cycle to be interrupted with a reference message.

Setting CANFDx_CHy_TTOCN.FGP after the gap time has started will start the transmission of a reference message immediately and will thereby synchronize the message schedule. When CANFDx_CHy_TTOCN.FGP is set before the gap time has started (while the basic cycle is still in progress), the next reference message is started at the end of the basic cycle, at the Tx_Ref_Trigger – there will be no gap time in the message schedule.

In time-triggered operation, bit Next_is_Gap = '1' in the reference message will be ignored, and the bits CANFDx_CHy_TTOCN.NIG, CANFDx_CHy_TTOCN.FGP, and CANFDx_CHy_TTOCN.TMG will be considered.

24.5.4 Stop Watch

The stop watch function enables capturing of M_TTCAN internal time values (local time, cycle time, or global time) triggered by an external event.

To enable the stop watch function, the application program must first define local time, cycle time, or global time as stop watch source via CANFDx_CHy_TTOCN.SWS. When CANFDx_CHy_TTOCN.SWS is not equal to '00' and TT Interrupt Register flag CANFDx_CHy_TTIR.SWE is '0', the actual value of the time selected by CANFDx_CHy_TTOCN.SWS will be copied into CANFDx_CHy_TTCPT.SWV on the next rising/falling edge (as configured via CANFDx_CHy_TTOCN.SWP) on stop watch trigger. This will set interrupt flag CANFDx_CHy_TTIR.SWE. After the application program has read CANFDx_CHy_TTCPT.SWV, it may enable the next stop watch event by resetting CANFDx_CHy_TTIR.SWE to '0'.

24.5.5 Local Time, Cycle Time, Global Time, and External Clock Synchronization

There are two possible levels in time-triggered CAN: Level 1 and Level 2. Level 1 provides only time-triggered operation using cycle time. Level 2 additionally provides increased synchronization quality, global time, and external clock synchronization. In both levels, all timing features are based on a local time base – the local time.


The local time is a 16-bit cyclic counter, it is incremented once each NTU. Internally the NTU is represented by a 3-bit counter, which can be regarded as a fractional part (three binary digits) of the local time. Generally, the 3-bit NTU counter is incremented eight times each NTU. If the length of the NTU is shorter than eight CAN clock periods (as may be configured in Level 1, or as a result of clock calibration in Level 2), the length of the NTU fraction is adapted, and the NTU counter is incremented only four times each NTU.

Figure 24-24 describes the synchronization of the cycle time and global time, performed in the same manner by all TTCAN nodes, including the time master. Any message

received or transmitted invokes a capture of the local time taken at the message's frame synchronization event. This frame synchronization event occurs at the sample point of each Start-of-Frame (SoF) bit and causes the local time to be stored as Sync_Mark. Sync_Marks and Ref_Marks are captured including the 3-bit fractional part.

Whenever a valid reference message is transmitted or received, the internal Ref_Mark is updated from the Sync_Mark. The difference between Ref_Mark and Sync_Mark is the Cycle Sync Mark (Cycle Sync Mark = Sync_Mark – Ref_Mark) stored in register CANFDx_CHy_TTCSM. The most significant 16 bits of the difference between Ref_Mark and the actual value of the local time is the cycle time (Cycle Time = Local Time – Ref_Mark).

Figure 24-24. Cycle Time and Global Time Synchronization



The cycle time that can be read from CANFDx_CHy_TTCTC.CT is the difference of the node's local time and Ref_Mark, both synchronized into the host clock domain and truncated to 16 bits.

The global time exists for TTCAN Level 0 and Level 2 only, in Level 1 it is invalid. The node's view of the global time is the local image of the global time in (local) NTUs. After configuration, a potential time master will use its own local time as global time. This is done by transmitting its own Ref_Marks as Master_Ref_Marks in the reference message (bytes 3 and 4). The global time that can be read from CANFDx_CHy_TTLGT.GT is the sum of the node's local time and its local offset, both synchronized into the host clock domain and truncated to 16 bit. The fractional part is used for clock synchronization only.

A node that receives a reference message calculates its local offset to the global time by comparing its local Ref_Mark with the received Master_Ref_Mark (see Figure 24-24). The node's view of the global time is local time + local offset. In a potential time master that has never received another time master's reference message, Local_Offset will be zero. When a node becomes the current time master after having received other reference messages first, Local_Offset will be frozen at its last value. In the time receiving nodes, Local_Offset may be subject to small adjustments, due to clock drift, when another node becomes time master, or when there is a global time discontinuity, signaled by Disc_Bit in the reference message. With the exception of global time discontinuity, the global time provided to the application program by register CANFDx_CHy_TTLGT is smoothed by a low-pass filtering to have a continuous monotonic value.

Figure 24-25. TTCAN Level 0 and Level 2 Drift Compensation




Figure 24-25 illustrates how in TTCAN Level 0 and Level 2 the receiving node compensates the drift between its own local clock and the time master's clock by comparing the length of a basic cycle in local time and in global time. If there is a difference between the two values, and the Disc_Bit in the reference message is not set, a new value for CANFDx_CHy_TURNA.NAV is calculated. If the synchronization deviation $(SD) = |NC - CANFDx\_CHy\_TURNA.NAV| \leq SDL$, the new value for CANFDx_CHy_TURNA.NAV takes effect. Otherwise, the automatic drift compensation is suspended.

In TTCAN Level 0 and Level 2, CANFDx_CHy_TTOST.QCS indicates whether the automatic drift compensation is active or suspended. In TTCAN Level 1, CANFDx_CHy_TTOST.QCS is always '1'.

The current time master may synchronize its local clock speed and the global time phase to an external clock source. This is enabled by bit CANFDx_CHy_TTOCF.EECS.

The stop watch function (see [Stop Watch on page 368](#)) may be used to measure the difference in clock speed between the local clock and the external clock. The local clock speed is adjusted by first writing the newly calculated Numerator Configuration Low to CANFDx_CHy_TURCF.NCL (CANFDx_CHy_TURCF.DC cannot be updated during operation). The new value takes effect by writing CANFDx_CHy_TTOCN.ECS to '1'.

The global time phase is adjusted by first writing the phase offset into the TT Global Time Preset register (CANFDx_CHy_TTGTP). The new value takes effect by writing CANFDx_CHy_TTOCN.SGT to '1'. The first reference message transmitted after the global time phase adjustment will have the Disc_Bit set to '1'.

CANFDx_CHy_TTOST.QGTP shows whether the node's global time is in phase with the time master's global time. CANFDx_CHy_TTOST.QGTP is permanently '0' in TTCAN Level 1 and when the SDL is exceeded in TTCAN Level 0,2 (CANFDx_CHy_TTOST.QCS = 0). It is temporarily '0' while the global time is low-pass filtered to supply the application with a continuous monotonic value. There is no low-pass filtering when the last reference message contains a Disc_Bit = '1' or when CANFDx_CHy_TTOST.QCS = 0.

24.5.6 Synchronization Triggers

One of the benefits of TTCAN is that it can make communication latency deterministic. To maintain this property across multiple CAN networks (or a Flexray network) these networks must be synchronized. M_TTCAN includes several trigger inputs and outputs to enable this synchronization.

Each M_TTCAN channel has trigger input and trigger output connected to trigger multiplexer. Using trigger functionality each channel has the possibility to not just synchronize with any other M_TTCAN channel, but also to other working

network (such as the Flexray network). For more information refer to the [Trigger Multiplexer chapter on page 493](#).

Stop watch and Event trigger inputs for the M_TTCAN channel are connected through the CAN_TT_TR_IN signal coming from the trigger multiplexer. Output trigger from the channel such as Time Mark Trigger and Register Time Mark triggers are connected through CAN_TT_TR_OUT to the trigger multiplexer.

Using this infrastructure, synchronously running networks are achievable.

24.5.7 TTCAN Error Level

The ISO 11898-4 specifies four levels of error severity:

- S0 - No Error
- S1 - Warning

Only notification of application, reaction application-specific.
- S2 Error

Notification of application. All transmissions in exclusive or arbitrating time windows are disabled (that is, no data or remote frames may be started). Potential time masters still transmit reference messages with the Reference Trigger Offset CANFDx_CHy_TTOST.RTO set to the maximum value of 127.
- S3 - Severe Error

Notification of application. All CAN bus operations are stopped; that is, transmission of dominant bits is not allowed and CANFDx_CHy_CCCR.MON is set. The S3 error condition remains active until the application updates the configuration (sets CANFDx_CHy_CCCR.CCE).

If several errors are detected at the same time, the highest severity prevails. When an error is detected, the application is notified by CANFDx_CHy_TTIR.ELC. The error level is monitored by CANFDx_CHy_TTOST.EL.

The M_TTCAN signals the following error conditions as required by ISO 11898-4:

Config_Error (S3)

Sets error level CANFDx_CHy_TTOST.EL to '11' when a merged arbitrating time window is not properly closed or when there is a Tx_Trigger with a time mark beyond the Tx_Ref_Trigger.

Watch_Trigger_Reached (S3)

Sets error level CANFDx_CHy_TTOST.EL to '11' when a watch trigger is reached because the reference message is missing.

Application_Watchdog (S3)

Sets error level CANFDx_CHy_TTOST.EL to '11' when the application fails to serve the application watchdog. The application watchdog is configured via CANFDx_CHy_TTOCF.AWL. It is served by reading the CANFDx_CHy_TTOST register. When the watchdog is not served in time, bit CANFDx_CHy_TTOST.AWE and interrupt flag CANFDx_CHy_TTIR.AW are set, all TTCAN communication is stopped, and the M_TTCAN is set into bus monitoring mode (CANFDx_CHy_CCCR.MON set to '1').

CAN_Bus_Off (S3)

Entering CAN_Bus_Off state sets error level CANFDx_CHy_TTOST.EL to '11'. CAN_Bus_Off state is signaled by CANFDx_CHy_PSR.BO = 1 and CANFDx_CHy_CCCR.INIT = 1.

Scheduling_Error_2 (S2)

Sets error level CANFDx_CHy_TTOST.EL to '10' if the MSC of one Tx_Trigger has reached 7. In addition, interrupt flag CANFDx_CHy_TTIR.SE2 is set. CANFDx_CHy_TTOST.EL is reset to 00 at the beginning of a matrix cycle when no Tx_Trigger has an MSC of 7 in the preceding matrix cycle.

Tx_Overflow (S2)

Sets error level CANFDx_CHy_TTOST.EL to '10' when the TX count is equal or higher than the expected number of Tx_Triggers CANFDx_CHy_TTMLM.ENTT and a Tx_Trigger event occurs. In addition, interrupt flag CANFDx_CHy_TTIR.TXO is set. CANFDx_CHy_TTOST.EL is reset to 00 when the TX count is no more than CANFDx_CHy_TTMLM.ENTT at the start of a new matrix cycle.

Scheduling_Error_1 (S1)

Sets error level CANFDx_CHy_TTOST.EL to '01' if within one matrix cycle the difference between the maximum MSC and the minimum MSC for all trigger memory elements (of exclusive time windows) is larger than 2, or if one of the MSCs of an exclusive Rx_Trigger has reached 7. In addition, interrupt flag CANFDx_CHy_TTIR.SE1 is set. If within one matrix cycle none of these conditions is valid, CANFDx_CHy_TTOST.EL is reset to 00.

Tx_Underflow (S1)

Sets error level CANFDx_CHy_TTOST.EL to '01' when the TX count is less than the expected number of Tx_Triggers CANFDx_CHy_TTMLM.ENTT at the start of a new matrix cycle. In addition, interrupt flag CANFDx_CHy_TTIR.TXU is set. CANFDx_CHy_TTOST.EL is reset to 00 when the TX count is at least CANFDx_CHy_TTMLM.ENTT at the start of a new matrix cycle.

24.5.8 TTCAN Message Handling

24.5.8.1 Reference Message

For potential time masters, the identifier of the reference message is configured via CANFDx_CHy_TTRMC.RID. No dedicated TX buffer is required for transmission of the reference message. When a reference message is transmitted, the first data byte (TTCAN Level 1) and the first four data bytes (TTCAN Level 0 and Level 2) will be provided by the FSE.

If the Payload Select reference message CANFDx_CHy_TTRMC.RMPS is set, the rest of the reference message's payload (Level 1: bytes 2-8, Level 0 and Level 2: bytes 5-6) is taken from TX Buffer 0. In this case, the data length DLC code from message buffer 0 is used.

Table 24-14. Number of Data Bytes Transmitted with a Reference Messages

CANFDx_CHy_TTRMC.RMPS	CANFDx_CHy_TXBRP.TRP0	Level 0	Level 1	Level 2
0	0	4	1	4
0	1	4	1	4
1	0	4	1	4
1	1	4 + MB0	1 + MB0	4 + MB0

To send additional payload with the reference message in Level 1, a DLC > 1 should be configured. For Level 0 and Level 2 a DLC > 4 is required. In addition, the transmission request pending bit CANFDx_CHy_TXBRP.TRP0 of message buffer 0 must be set (see [Table 24-14](#)). If CANFDx_CHy_TXBRP.TRP0 is not set when a reference message is started, the reference message is transmitted with the data bytes supplied by the FSE only.

For acceptance filtering of reference messages the Reference Identifier CANFDx_CHy_TTRMC.RID is used.

24.5.8.2 Message Reception

Message reception is done via the two RX FIFOs in the same way as for event-driven CAN communication.

The MSC is part of the corresponding trigger memory element and must be initialized to zero during configuration. It is updated while the M_TTCAN is in synchronization states In_Gap or In_Schedule. The update happens at the message's Rx_Trigger. At this point, it is checked at which acceptance filter element the latest message received in this basic cycle is matched. The matching filter number is stored as the acceptance filter result. If this is the same as the filter number defined in this trigger memory element, the MSC is decremented by one. If the acceptance filter result is not the same filter number as defined for this filter element, or if the acceptance filter result is cleared, the MSC is incremented

by one. At each Rx_Trigger and at each start of cycle, the last acceptance filter result is cleared.

The time mark of an Rx_Trigger should be set to a value that ensures reception and acceptance filtering for the targeted message is completed. This should consider the RAM access time and the order of the filter list. It is recommended, that filters used for Rx_Triggers are placed at the beginning of the filter list. It is not recommended to use an Rx_Trigger for the reference message.

24.5.8.3 Message Transmission

For time-triggered message transmission, the M_TTCAN supplies 32 dedicated TX buffers (see [TTCAN Configuration on page 362](#)). A TX FIFO or TX queue is not available when the M_TTCAN is configured for time-triggered operation (CANFDx_CHy_TTOCF.OM = 01 or 10).

Each Tx_Trigger in the trigger memory points to a particular TX buffer containing a specific message. There may be more than one Tx_Trigger for a given TX buffer if that TX buffer contains a message that is to be transmitted more than once in a basic cycle or matrix cycle.

The application program must update the data regularly and on time, synchronized to the cycle time. The host CPU should ensure that no partially updated messages are transmitted. To assure this the host should proceed in the following way:

Tx_Trigger_Single/Tx_Trigger_Merged/
Tx_Trigger_Arbitration:

- Check whether the previous transmission has completed by reading TXBTO
- Update the TX buffer's configuration and/or payload
- Issue an Add Request to set the TX Buffer Request Pending bit

Tx_Trigger_Continuous:

- Issue a Cancellation Request to reset the TX Buffer Request Pending bit
- Check whether the cancellation has finished by reading CANFDx_CHy_TXBCF
- Update TX buffer configuration and/or payload
- Issue an Add Request to set the TX Buffer Request Pending bit

The message MSC stored with the corresponding Tx_Trigger provides information on the success of the transmission.

The MSC is incremented by one when the transmission cannot be started because the CAN bus was not idle within the corresponding transmit enable window or when the message was started but could not be completed successfully. The MSC is decremented by one when the message is transmitted successfully or when the message could have been started within its transmit enable window

but was not started because transmission was disabled (M_TTCAN in Error Level S2 or host has disabled this particular message).

The TX buffers may be managed dynamically – several messages with different identifiers may share the same TX buffer element. In this case the host must ensure that no transmission request is pending for the TX buffer element to be reconfigured by checking CANFDx_CHy_TXBRP.

If a TX buffer with pending transmission request should be updated, the host must first issue a cancellation request and check whether the cancellation has completed by reading CANFDx_CHy_TXBCF before it starts updating.

The TX handler will transfer a message from the message RAM to its intermediate output buffer at the trigger element, which becomes active immediately before the Tx_Trigger element that defines the beginning of the transmit window. During and after transfer time, the transmit message may not be updated and its CANFDx_CHy_TXBRP bit may not be changed. To control this transfer time, an additional trigger element may be placed before the Tx_Trigger. An example is a Time_Base_Trigger, which does not cause any other action. The difference in time marks between the Tx_Trigger and the preceding trigger should be large enough to guarantee that the TX handler can read four words from the message RAM even at high RAM access load from other modules.

Transmission in Exclusive Time Windows

A transmission is started time-triggered when the cycle time reaches the time mark of a Tx_Trigger_Single or Tx_Trigger_Continuous. There is no arbitration on the bus with messages from other nodes. The MSC is updated according to the result of the transmission attempt. After successful transmission started by a Tx_Trigger_Single, the respective TX Buffer Request Pending bit is reset. After successful transmission started by a Tx_Trigger_Continuous the respective TX Buffer Request Pending bit remains set. When the transmission is not successful due to disturbances, it will be repeated the next time one of its Tx_Triggers becomes active.

Transmission in Arbitrating Time Windows

A transmission is started time-triggered when the cycle time reaches the time mark of a Tx_Trigger_Arbitration. Several nodes may start to transmit at the same time. In this case the message has to arbitrate with the messages from other nodes. The MSC is not updated. When the transmission is not successful (lost arbitration or disturbance), it will be repeated the next time one of its Tx_Triggers becomes active.

Transmission in Merged Arbitrating Time Windows

The purpose of a merged arbitrating time window is to enable multiple nodes to send a limited number of frames, which are transmitted in immediate sequence, the order given by CAN arbitration. It is not intended for burst

transmission by a single node. Because the node does not have exclusive access within this time window, all requested transmissions may not be successful.

Messages that have lost arbitration or were disturbed by an error, may be retransmitted inside the same merged arbitrating time window. The retransmission will not be started if the corresponding Transmission Request Pending flag was reset by a successful TX cancellation.

In single transmit windows, the TX handler transmits the message indicated by the message number of the trigger element. In merged arbitrating time windows, it can handle up to three message numbers from the trigger list. Their transmissions will be attempted in the sequence defined by the trigger list. If the time mark of a fourth message is read before the first is transmitted (or canceled by the host), the fourth request will be ignored.

The transmission inside a merged arbitrating time window is not time-triggered. The transmission of a message may start before its time mark, or after the time mark if the bus was not idle.

The messages transmitted by a specific node inside a merged arbitrating time window will be started in the order of their Tx_Triggers. Therefore, a message with low CAN priority may prevent the successful transmission of a following message with higher priority, if there is competing bus traffic. This should be considered for the configuration of the trigger list. Time_Base_Triggers may be placed between consecutive Tx_Triggers to define the time until the data of the corresponding TX buffer needs to be updated.

24.5.9 TTCAN Interrupt and Error Handling

The TT Interrupt Register CANFDx_CHy_TTIR consists of four segments. Each interrupt can be enabled separately by the corresponding bit in the TT Interrupt Enable register CANFDx_CHy_TTIE. The flags remain set until the host clears them. A flag is cleared by writing a '1' to the corresponding bit position.

The first segment consists of flags CER, AW, WT, and IWT. Each flag indicates a fatal error condition where the CAN communication is stopped. With the exception of IWT, these error conditions require a reconfiguration of the M_TTCAN module before the communication can be restarted.

The second segment consists of flags ELC, SE1, SE2, TXO, TXU, and GTE. Each flag indicates an error condition where the CAN communication is disturbed. If they are caused by a transient failure, such as by disturbances on the CAN bus, they will be handled by the TTCAN protocol's failure handling and do not require intervention by the application program.

The third segment consists of flags GTD, GTW, SWE, TTMI, and RTMI. The first two flags are controlled by global time events (Level 0 and Level 2 only) that require a reaction by

the application program. With a Stop Watch Event, internal time values are captured. The Trigger Time Mark Interrupt notifies the application that a specific Time_Base_Trigger is reached. The Register Time Mark Interrupt signals that the time referenced by CANFDx_CHy_TTOCN.TMC (cycle, local, or global) equals time mark CANFDx_CHy_TTTMK.TM. It can also be used to finish a gap.

The fourth segment consists of flags SOG, CSM, SMC, and SBC. These flags provide a means to synchronize the application program to the communication schedule.

24.5.10 Level 0

TTCAN Level 0 is not part of ISO11898-4. This operation mode makes the hardware, that in TTCAN Level 2 maintains the calibrated global time base, also available for event-driven CAN according to ISO 11898-1:2015.

Level 0 operation is configured via CANFDx_CHy_TTOCF.OM = 11. In this mode, M_TTCAN operates in event-driven CAN communication; there is no fixed schedule, the configuration of CANFDx_CHy_TTOCF.GEN is ignored. External event-synchronized operation is not available in Level 0. A synchronized time base is maintained by transmission of reference messages.

In Level 0 the trigger memory is not active and need not be configured. The time mark interrupt flag (CANFDx_CHy_TTIR.TTMI) is set when the cycle time has reached CANFDx_CHy_TTOCF.IRTO × 0x200. It reminds the host to set a transmission request for message buffer 0. The Watch_Trigger interrupt flag (CANFDx_CHy_TTIR.WT) is set when the cycle time has reached 0xFF00. These values were chosen to have enough margin for a stable clock calibration. There are no further TT-error-checks.

Register time mark interrupts (CANFDx_CHy_TTIR.RTMI) are also possible.

The reference message is configured as for Level 2 operation. Received reference messages are recognized by the identifier configured in register CANFDx_CHy_TTRMC. For the transmission of reference messages only message buffer 0 may be used. The node transmits reference messages any time the host sets a transmission request for message buffer 0; there is no reference trigger offset.

Level 0 operation is configured via:

- CANFDx_CHy_TTRMC
- CANFDx_CHy_TTOCF except EVTP, AWL, GEN
- CANFDx_CHy_TTMLM except ENTT, TXEW
- CANFDx_CHy_TURCF

Level 0 operation is controlled via:

- CANFDx_CHy_TTOCN except NIG, TMG, FGP, GCS, TTMIE
- CANFDx_CHy_TTGTP
- CANFDx_CHy_TTTMK
- CANFDx_CHy_TTIR excluding bits CER, AW, IWT SE2, SE1, TXO, TXU, SOG (no function)
- CANFDx_CHy_TTIR – the following bits have changed function:
 - TTMI not defined by trigger memory - activated at cycle time CANFDx_CHy_TTOCF.IRTO × 0x200
 - WT not defined by trigger memory - activated at cycle time 0xFF00


Level 0 operation is signaled via:

- CANFDx_CHy_TTOST excluding bits AWE, WFE, GSI, GFI, RTO (no function)

24.5.10.1 Synchronizing

Figure 24-26 describes the states and state transitions in TTCAN Level 0 operation. Level 0 has no In_Gap state.

Figure 24-26. Level 0 Schedule Synchronization State Machine



T0: transition condition always taking prevalence

T1: Init state left, cycle time is zero

T2: at least two successive reference messages observed
(last reference message did not contain a set Disc_Bit bit)

24.5.10.2 Handling Error Levels

During Level 0 operation only the following error conditions may occur:

- Watch_Trigger_Reached (S3), reached cycle time 0xFF00
- CAN_Bus_Off (S3)

Because S1 and S2 errors are not possible, the error level can only switch between S0 (No Error) and S3 (Severe Error). In TTCAN Level 0 an S3 error is handled differently.


When S3 error is reached, both CANFDx_CHy_TTOST.SYS and CANFDx_CHy_TTOST.MS are reset, and interrupt flags CANFDx_CHy_TTIR.GTE and CANFDx_CHy_TTIR.GTD are set.

When S3 (CANFDx_CHy_TTOST.EL = 11) is entered, bus monitoring mode is, contrary to TTCAN Level 1 and Level 2, not entered. S3 error level is left automatically after transmission (time master) or reception (time slave) of the next reference message.

24.5.10.3 Master Slave Relation

Figure 24-27 describes the master slave relation in TTCAN Level 0. In case of an S3 error, the M_TTCAN returns to state Master_Off.

Figure 24-27. Level 0 Master to Slave Relation



T0: transition condition always taking prevalence

T1: reference message observed when not potential master

T2: reference message observed with master priority != own master priority, error state != S3

T3: reference message observed with master priority = own master priority, error state != S3

T4: reference message observed with own master priority

T5: reference message observed with master priority higher than own master priority

T6: error state S3

T7: error state S3

T8: error state S3

24.5.11 Synchronization to External Time Schedule

This feature can be used to synchronize the phase of the M_TTCAN's schedule to an external schedule (for example, that of a second TTCAN network). It is applicable only when the M_TTCAN is current time master (CANFDx_CHy_TTOST.MS = 11).

External synchronization is controlled by the CANFDx_CHy_TTOCN.ESCN bit. If CANFDx_CHy_TTOCN.ESCN is set, at rising edge of the internal event trigger pin, the M_TTCAN compares its actual

cycle time with the target phase value configured by CANFDx_CHy_TTGTP.CTP.

Before setting CANFDx_CHy_TTOCN.ESCN, the host should adapt the phases of the two time schedules, for example, by using the TTCAN gap control (see 24.5.3 TTCAN Gap Control). When the host sets CANFDx_CHy_TTOCN.ESCN, CANFDx_CHy_TTOST.SPL is set.

If the difference between the cycle time and target phase value CANFDx_CHy_TTGTP.CTP at the trigger is greater than 9 NTU, the phase lock bit CANFDx_CHy_TTOST.SPL is reset, and interrupt flag CANFDx_CHy_TTIR.CSM is set. CANFDx_CHy_TTOST.SPL is also reset (and

CANFDx_CHy_TTIR.CSM is set), when another node becomes time master.

If both CANFDx_CHy_TTOST.SPL and CANFDx_CHy_TTOCN.ESCN are set, and if the difference between the cycle time and the target phase value CANFDx_CHy_TTGTP.CTP is less or equal 9 NTU, the phase lock bit CANFDx_CHy_TTOST.SPL remains set, and the measured difference is used as reference trigger offset value to adjust the phase at the next transmitted reference message.

Note: The rising edge detection at the internal pin is enabled at the start of each basic cycle. The first rising edge triggers the compare of the actual cycle time with CANFDx_CHy_TTGTP.CTP. All further edges until the beginning of the next basic cycle are ignored.


24.6 Setup Procedures

This section provides example procedures for configurations of M_TTCAN group and flow for respective M_TTCAN channels.

24.6.1 General Program Flow

This is a general flow to configure the M_TTCAN module.


Figure 24-28. General Program Flow



24.6.2 Clock Stop Request

To save power, the application can stop providing clock to unused M_TTCAN channels by following these steps.


Figure 24-29. Clock Stop Request Procedure



To resume providing clock, the CANFDx_CTL.STOP_REQ bit should be reset.


24.6.3 Message RAM OFF Operation

Figure 24-30. Message RAM OFF Operation



24.6.4 Message RAM ON Operation

Figure 24-31. Message RAM On Procedure




After switching message RAM ON again, software needs to allow a certain power-up time before message RAM can be used; that is, before STOP_REQ can be de-asserted. The power-up time is equivalent to the system SRAM power-up time specified in the CPUSS.RAM_PWR_DELAY_CTL register.

24.6.5 Consolidated Interrupt Handling

When using consolidated interrupt for the M_TTCAN group, follow the procedure given in [Figure 24-32](#).


Figure 24-32. Consolidated Interrupt Processing



24.6.6 Procedures Specific to M_TTCAN Channel


This section describes sample procedures per channel. If several M_TTCAN channels are used, the application should configure each channel as shown in [Figure 24-33](#). The figure shows the general program flow (per channel).

Figure 24-33. Configuration Sequence Specific to Channel



24.6.6.1 CAN Bus Configuration


Figure 24-34. Configuration Required for CAN Bus



24.6.6.2 Message RAM Configuration

The following flow chart shows an overview of the message RAM configuration.

Figure 24-35. Message RAM Configuration Overview



Each configuration mentioned in the overview is detailed in the following figures.

Figure 24-36. ID Filter List Configuration






Figure 24-37. RX FIFO and RX Buffer Configuration





```
graph TD; Start([Start]) --> Step1[Write to Tx Buffer Configuration(CANFDx_CHy_TXBC)  
• Tx FIFO/Queue Mode(TFQM)  
• Transmit FIFO/Queue Size(TFQS[5:0])  
• Number of Dedicated Transmit Buffers(NDTB[5:0])  
• Tx Buffers Start Address(TBSA[15:2])]; Step1 --> Step2[Write to Tx Buffer Element Size Configuration(CANFDx_CHy_TXESC)  
• Tx Buffer Data Field Size(TBDS[2:0])]; Step2 --> Step3[Write to Tx Event FIFO Configuration(CANFDx_CHy_TXEFC)  
• Event FIFO Watermark(EFWM[5:0])  
• Event FIFO Size(EFS[5:0])  
• Event FIFO Start Address(EFSA[15:2])]; Step3 --> End([End]);
```

The flowchart illustrates the sequence of operations for configuring the CANFD Tx Buffer. It begins with a 'Start' terminal, followed by three main configuration steps, and ends with an 'End' terminal. The steps are connected by downward arrows, indicating a sequential process. The first step configures the Tx Buffer Configuration (CANFDx_CHy_TXBC), the second configures the Tx Buffer Element Size Configuration (CANFDx_CHy_TXESC), and the third configures the Tx Event FIFO Configuration (CANFDx_CHy_TXEFC). Each step includes a list of specific configuration parameters and their bit fields. A vertical label 'in arbitrary order' is positioned to the right of the flowchart, indicating that the steps can be performed in any sequence.

Start

Write to Tx Buffer Configuration(CANFDx_CHy_TXBC)

- Tx FIFO/Queue Mode(TFQM)
- Transmit FIFO/Queue Size(TFQS[5:0])
- Number of Dedicated Transmit Buffers(NDTB[5:0])
- Tx Buffers Start Address(TBSA[15:2])

Write to Tx Buffer Element Size Configuration(CANFDx_CHy_TXESC)

- Tx Buffer Data Field Size(TBDS[2:0])

Write to Tx Event FIFO Configuration(CANFDx_CHy_TXEFC)

- Event FIFO Watermark(EFWM[5:0])
- Event FIFO Size(EFS[5:0])
- Event FIFO Start Address(EFSA[15:2])

End


in arbitrary order

```

graph TD
    Start([Start]) --> Request{ID Filter configuration is requested?}
    Request -- No --> End([End])
    Request -- Yes --> Join1(( ))
    Join1 --> Finished{ID Filter configuration finished?}
    Finished -- Yes --> End
    Finished -- No --> Type{Standard ID or Extended ID}
    Type -- Standard ID --> StdElement[Std_element = index of next free Standard Message ID Filter Element  
(ID Filter elements should be controlled by software)]
    StdElement --> ReadStd[Read Filter List Standard Start Address  
SA=CANFDx_Chy_SIDFC.FLSSA[15:2]]
    ReadStd --> CalcAPStd[Calculate access pointer(AP)  
AP=SA+Std_element]
    CalcAPStd --> WriteStd[Write Standard ID Filter information to Message RAM  
Message RAM(AP)=Standard ID Filter information(S0)]
    WriteStd --> Join2(( ))
    Type -- Extended ID --> ExtElement[Ext_element = index of next free Extended Message ID Filter Element  
(ID Filter elements should be controlled by software)]
    ExtElement --> ReadExt[Read Filter List Extended Start Address  
SA=CANFDx_Chy_XIDFC.FLESA[15:2]]
    ReadExt --> CalcAPExt[Calculate access pointer(AP)  
AP=SA+(Ext_element×2)]
    CalcAPExt --> WriteExt[Write Extended ID Filter information to Message RAM  
Message RAM(AP) =Extended ID Filter information(F0)  
Message RAM(AP+1)=Extended ID Filter information(F1)]
    WriteExt --> Join2
    Join2 --> Join1
  
```


24.6.6.3 Interrupt Configuration

Figure 24-40. Interrupt Configuration



24.6.6.4 Transmit Frame Configuration

Figure 24-41. Transmit Frame Configuration



24.6.6.5 Interrupt Handling

When consolidated interrupts are configured, INTR0/1_CAUSE register will be read to find out the source M_TTCAN channel for the triggered interrupt. Figure 24-42 shows a general interrupt handling flow chart.

Figure 24-42. Interrupt Handling




Figure 24-43. Bus OFF Error Handling




Figure 24-44. Bit Error Handling




Figure 24-45. Message RAM Access Failure Handling




Figure 24-46. TX Event FIFO Handling




Figure 24-47. Dedicated RX Buffer Handling




Figure 24-48. High Priority Message Handling





Figure 24-49. RX FIFO Top Pointer Handling



24.7 Registers

Register	Name	Description
CANFDx_CHy_CREL	Core Release Register	Displays the revision of the CAN FD controller.
CANFDx_CHy_ENDN	Endian Register	Checks the endianness of the CAN FD controller when accessed by the CPU.
CANFDx_CHy_DBTP	Data Bit Timing and Prescaler Register	Configures the data bit time and enables Transmitter Delay Compensation.
CANFDx_CHy_TEST	Test Register	Monitors the CANx_y_RX/CANx_y_TX pins. It is also used to enable the Loop Back modes.
CANFDx_CHy_RWD	RAM Watchdog	Monitors the message RAM to see if it is ready to be accessed.
CANFDx_CHy_CCCR	CC Control Register	Configures various operating modes of the CAN FD controller.
CANFDx_CHy_NBTP	Nominal Bit Timing and Prescaler Register	Configures the nominal bit time of the CAN FD controller.
CANFDx_CHy_TSCC	Timestamp Counter Configuration	Holds the settings for Timestamp Generation.
CANFDx_CHy_TOCC	Timeout Counter Configuration	Holds the settings for the Timeout Counter.
CANFDx_CHy_TOCV	Timeout Counter Value	Holds the value of the Timeout Counter.
CANFDx_CHy_ECR	Error Counter Register	Holds the values of the Error Counters
CANFDx_CHy_PSR	Protocol Status Register	Displays the CAN protocol status of the CAN FD controller.
CANFDx_CHy_TDCR	Transmitter Delay Compensation Register	Configures the offset value and the filter window length for Transmitter Delay Compensation
CANFDx_CHy_IR	Interrupt Register	Holds the flags that are set when one of the listed conditions is detected (edge-sensitive).
CANFDx_CHy_IE	Interrupt Enable	The settings in this register determine which status changes in the Interrupt Register (IR) will be signaled on an interrupt line
CANFDx_CHy_ILS	Interrupt Line Select	Assigns an interrupt generated by a specific interrupt flag from the Interrupt Register (IR) to one of the two CAN FD controller interrupt lines (canfd_int0/1).
CANFDx_CHy_ILE	Interrupt Line Enable	This register can separately enable/disable each of the two interrupt lines to the CPU.
CANFDx_CHy_GFC	Global Filter Configuration	Global settings for Message ID filtering.
CANFDx_CHy_SIDFC	Standard ID Filter Configuration	Settings for 11-bit standard Message ID filtering.
CANFDx_CHy_XIDFC	Extended ID Filter Configuration	Settings for 29-bit extended Message ID filtering.
CANFDx_CHy_XIDAM	Extended ID AND Mask	Defines the valid bits of a 29-bit ID for acceptance filtering.
CANFDx_CHy_HPMS	High Priority Message Status	This register is updated every time a Message ID filter element configured to generate a priority event matches.
CANFDx_CHy_NDAT1	New Data 1	Holds flags that are set when the respective dedicated RX buffer receives a frame.
CANFDx_CHy_NDAT2	New Data 2	Holds flags that are set when the respective dedicated RX buffer receives a frame.
CANFDx_CHy_RXF0C	RX FIFO 0 Configuration	Settings for the RX FIFO 0.
CANFDx_CHy_RXF0S	RX FIFO 0 Status	Status of the RX FIFO 0.
CANFDx_CHy_RXF0A	RX FIFO 0 Acknowledge	Acknowledges that the CPU has read a message or a sequence of messages from the RX FIFO 0 to indicate to the CAN FD controller that the corresponding message RAM area may be released.
CANFDx_CHy_RXBC	RX Buffer Configuration	Defines the start address of the RX Buffer section in the message RAM.
CANFDx_CHy_RXF1C	RX FIFO 1 Configuration	Settings for the RX FIFO 1.
CANFDx_CHy_RXF1S	RX FIFO 1 Status	Status of the RX FIFO 1.

Register	Name	Description
CANFDx_CHy_RXF1A	RX FIFO 1 Acknowledge	Acknowledges that the CPU has read a message or a sequence of messages from the RX FIFO 1 to indicate to the CAN FD controller that the corresponding message RAM area may be released.
CANFDx_CHy_RXESC	RX Buffer/FIFO Element Size Configuration	Configures the number of data bytes belonging to an RX buffer and FIFO element.
CANFDx_CHy_TXBC	TX Buffer Configuration	Settings for TX buffers stored in the message RAM
CANFDx_CHy_TXFQS	TX FIFO/Queue Status	Related to the pending TX requests listed in the TX Buffer Request Pending register (CANFDx_CHy_TXBRP).
CANFDx_CHy_TXESC	TX Buffer Element Size Configuration	Configures the number of data bytes belonging to a TX buffer element
CANFDx_CHy_TXBRP	TX Buffer Request Pending	Holds the status of the transmission requests of each corresponding TX Buffer
CANFDx_CHy_TXBAR	TX Buffer Add Request	Requests the transmission of each corresponding TX buffer.
CANFDx_CHy_TXBCR	TX Buffer Cancellation Request	Cancels transmission requests of each corresponding TX buffer.
CANFDx_CHy_TXBTO	TX Buffer Transmission Occurred	Displays whether the corresponding TX buffer is transmitted.
CANFDx_CHy_TXBCF	TX Buffer Cancellation Finished	Signals whether the cancellation request of the corresponding TX buffer is successful.
CANFDx_CHy_TXBTIE	TX Buffer Transmission Interrupt Enable	The settings in this register determine which TX buffer will assert an interrupt upon transmission.
CANFDx_CHy_TXBCIE	TX Buffer Cancellation Finished Interrupt Enable	The settings in this register determine which TX buffer will assert an interrupt upon completion of a transmission cancellation request.
CANFDx_CHy_TXEFC	TX Event FIFO Configuration	Settings for the TX Event FIFO.
CANFDx_CHy_TXEFS	TX Event FIFO Status	Status of the TX Event FIFO.
CANFDx_CHy_TXEFA	TX Event FIFO Acknowledge	Acknowledges that the CPU has read an event from the TX Event FIFO to indicate to the CAN FD controller that the corresponding message RAM area may be released.
CANFDx_CHy_TTTMC	TT Trigger Memory Configuration	Configures memory element and memory start address.
CANFDx_CHy_TTRMC	TT Reference Message Configuration	Configures the reference message such as reference identifier, reference payload type, and type of identifier.
CANFDx_CHy_TTOCF	TT Operation Configuration	Configures fundamentals for time-triggered operations such as TTCAN operation level, time master, and clock calibration.
CANFDx_CHy_TTMLM	TT Matrix Limits	Configures cycle counts and synchronization for the clock start, and enables TX Window.
CANFDx_CHy_TURCF	TUR Configuration	Configures numerator and denominator for time unit configuration.
CANFDx_CHy_TTOCN	TT Operation Control	Controls main TTCAN operation.
CANFDx_CHy_TTGTP	TT Global Time Preset	Sets preset value and defines target of cycle time when a rising edge of TTCAN event is expected.
CANFDx_CHy_TTTMK	TT Time Mark	Configures number of cycles in which time mark will be valid.
CANFDx_CHy_TTIR	TT Interrupt Register	Flags in the TTIR is set when particular conditions are met.
CANFDx_CHy_TTIE	TT Interrupt Enable	Provides possibility to enable interrupt for several status changes.
CANFDx_CHy_TTILS	TT Interrupt Line Select	User can select dedicated Interrupt0 or Interrupt1 line for specific interrupt source.
CANFDx_CHy_TTOST	TT Operation Status	Status register for TT operation.
CANFDx_CHy_TURNA	TUR Numerator Actual	Shows actual numerator value for time unit configuration.
CANFDx_CHy_TTLGT	TT Local and Global Time	Shows non-fractional part of the global and local time.
CANFDx_CHy_TTCTC	TT Cycle Time and Count	Read-only register that shows current cycle count and non-fraction part of the cycle time.
CANFDx_CHy_TTCPT	TT Capture Time	Read-only register that shows current cycle count and stop watch value.

Register	Name	Description
CANFDx_CHy_TTCSM	TT Cycle Sync Mark	Read-only register that shows cycle sync mark in terms of cycle time.
CANFDx_CHy_RXFTOP_CTL	Receive FIFO Top control	Enables Receive FIFO Top Control logic for both FIFOs.
CANFDx_CHy_RXFTOP0_STAT	Receive FIFO 0 Top Status	This is a pointer to the next word in the message buffer defined by FIFO start address.
CANFDx_CHy_RXFTOP0_DATA	Receive FIFO 0 Top Data	Data placed at the address by CANFDx_CHy_RXFTOP0_STAT.
CANFDx_CHy_RXFTOP1_STAT	Receive FIFO 1 Top Status	This is a pointer to the next word in the message buffer defined by FIFO start address.
CANFDx_CHy_RXFTOP1_DATA	Receive FIFO 1 Top Data	Data placed at the address by CANFDx_CHy_RXFTOP1_STAT.
CANFDx_CTL	Global CAN Control Register	Provides clock control to the respective TTCAN channels.
CANFDx_STATUS	Global CAN Status Register	Read-only register that shows the acknowledge from the respective TTCAN channel for the clock stop request.
CANFDx_INTR0_CAUSE	Consolidated Int0 Cause Register	Shows pending interrupt0 for each TTCAN channel.
CANFDx_INTR1_CAUSE	Consolidated Int1 Cause Register	Shows pending interrupt1 for each TTCAN channel.
CANFDx_TS_CTL	Time Stamp Control Register	Configuration for the Timestamp prescaler and counter enable is done in this register.
CANFDx_TS_CNT	Time Stamp Count Register	Shows timestamp counter value.
CANFDx_ECC_CTL	ECC Control Register	Configures ECC for message RAM.
CANFDx_ECC_ERR_INJ	ECC Error Injection Register	ECC error can be injected to a particular word address in the message RAM using this register.

Note: 'x' in CANFDx signifies the CAN macro instance and 'y' in CANFDx_CHy signifies the channel under the CAN instance.

25. Timer, Counter, and PWM



The Timer, Counter, and Pulse Width Modulator (TCPWM) block in Traveo II implements a 16- or 32-bit timer, counter, pulse width modulator (PWM), pseudo random PWM, shift register, and quadrature decoder functionality. TCPWM includes up to four counter groups where each group can include up to 256 counters. The counter can be used to measure the period and pulse width of an input signal (timer), find the number of times an event occurs (counter), generate PWM signals, or decode quadrature signals. The TCPWM block works in Active and Sleep modes.

This chapter explains the features, implementation, and operational modes of the TCPWM block.


25.1 Features

The TCPWM block has the following features:

- Supports up to four counter groups (device specific)
- Each counter group consists up to 256 counters (counter group specific)
- Each counter can run in one of seven function modes:
 - Timer-counter with compare
 - Timer-counter with capture
 - Quadrature decoding
 - Pulse width modulation
 - PWM with dead time
 - Pseudo-random PWM
 - Shift register mode
- 16-bit or 32-bit counters (counter group specific)
- Up, down, and up/down counting modes
- Clock prescaling (division by 1, 2, 4, ... 64, 128)
- Up to two capture and compare functions (counter group specific)
- Double buffering of all compare/capture and period registers
- Two output trigger signals for each counter to indicate underflow, overflow, and capture/compare events; they can also directly be connected with the line output signal
- Supports interrupt on:
 - Terminal Count - Depends on the mode; typically occurs on overflow or underflow
 - Capture/Compare - The count is captured in the capture registers or the counter value equals the value in the compare register
- Line out selection feature for stepper motor application including two complementary output lines with dead time insertion
- Selectable start, reload, stop, count, and two capture event signals for each TCPWM with rising edge, falling edge, both edges, and level trigger options
- Each counter with up to 254 (device specific) synchronized input trigger signals and two constant input signals: '0' and '1'.
- Two types of input triggers for each counter:
 - General-purpose triggers used by all counters
 - One-to-one triggers for specific counter
- Synchronous operation of multiple counters
- Debug mode support

25.2 Block Diagram

Figure 25-1. TCPWM Block Diagram



In the Traveo II device, there are up to four TCPWM counter groups each supporting up to 256 counters; they can have a counter width of 16-bit or 32-bit. In addition, counter groups can also include a second capture and compare function. Refer to the device datasheet to find dedicated counter group configurations.

Note: This document does not discuss the specific counter group configuration in detail. If a second capture/compare feature is mentioned, refer to the device datasheet to know if these functions are available in the particular device.

All register names and related bit fields are related to one counter example. Find the register prefixes for dedicated counters in the *Traveo II Body Controller High Registers TRM*.

Each counter can have 254 input trigger signals and two constant input signals, '0' and '1'; all of them are synchronized with CLK_PERI clock.

The TCPWM block has these interfaces:

- **Bus interface:** Connects the block to the CPU subsystem via AHB-Lite interface.
- **I/O signal interface:** Consists of input triggers (such as reload, start, stop, count, and capture0/1) and output signals (such as LINE_OUT, LINE_COMPL_OUT, TR_OUT0, and TR_OUT1).
- **Interrupts:** Provides interrupt request signals from each counter, based on terminal count (TC), Compare/Capture CC0_match, or Compare/Capture CC1_match event.

- **System interface:** Consists of control signals such as clock and reset from the system resources subsystem (SRSS).

The TCPWM block can be configured by writing to the TCPWM registers. See [“TCPWM Registers” on page 454](#) for more information on all registers required for this block.

25.2.1 Enabling and Disabling Counters in TCPWM Block

A counter can be enabled by writing '1' to the corresponding ENABLE bit of the CTRL register; it can be disabled by writing '0' to the same bit.

Note: The counter must be configured before enabling it. Disabling the counter retains the values in the registers.

25.2.2 Clocking

The TCPWM receives a single clock, CLK_PERI. Furthermore, it receives a system clock enable signal clock_sys_en to generate internal CLK_SYS and a counter clock enable signal clock_counter_en for PCLK_TCPWM[x]_CLOCKS[y] of each counter.

Each TCPWM counter can have its own clock source. The only source for the clock is from the configurable peripheral clock dividers generated by the clocking system; see the [Clocking System chapter on page 197](#) for details. To select a clock divider for a particular counter inside a TCPWM, use the CLOCK_CTL register from the PERI register space. In this section the clock to the counter will be called PCLK_TCPWM[x]_CLOCKS[y]. Event generation is

performed on the PCLK_TCPWM[x]_CLOCKS[y]. Another clock, CLK_SYS, is used for the pulse width of the output triggers. CLK_SYS is synchronous to CLK_PERI, but can be divided using CLOCK_CTL from the PERI_GROUP_STRUCT register.

25.2.2.1 Clock Prescaling

PCLK_TCPWM[x]_CLOCKS[y] can be further divided inside each counter, with values of 1, 2, 4, 8...64, 128. This division is called prescaling. The prescaling is set in the DT_LINE_OUT_L [7:0] field of the DT register. The lower three bits of this field determine prescaling of the selected counter clock.


Note: Clock prescaling is not available in quadrature mode and pulse width modulation mode with dead time.

25.2.2.2 Count Event

The counter functionality is performed on an “active count” prescaled clock, which is gated by a “count event” signal. For example, a counter increments or decrements by '1' every counter clock cycle in which a count event is detected.

Note: Count events are not supported in quadrature and pulse-width modulation pseudo-random modes; the PCLK_TCPWM[x]_CLOCKS[y] is used in these cases instead of the active count prescaled clock.

Figure 25-2. Counter Clock Generation



All status or output change can only happen at active count prescaled counter clock. In the other words, if a count event is inactive, counter, status, interrupt, and all outputs will not change value. For example, if a count event in pass-through mode becomes low when counter goes to '0' in down count mode, the tc event and underflow event will be generated at the next prescaled counter clock after count event goes high. The only exception is immediate kill mode. Kill input will suppress the PWM output immediately regardless of active count prescaled counter clock.

inputs are available. Input trigger 0 is always constant '0' and input trigger 1 is always constant '1'.

Each counter can select any of the 256 trigger signals to be the source for any of the following events:

- Capture 0 and Capture 1
- Count
- Reload
- Stop/Kill
- Start

Note: The TR_CMD register can be used to trigger the Reload, Stop, Start, and Capture0/1 respectively from software.

25.2.3 Trigger Inputs

Each TCPWM block has 254 Trigger_In signals and constant '0' and '1' signals, which come from other on-chip resources such as other TCPWMs, SCBs, and DMA. The Trigger_In signals are shared with all counters inside one TCPWM block.

Two types of trigger signals are synchronized and can be used by the counters to generate events.

- General-purpose triggers. These can be used by all counters. These triggers are generated by different blocks in the system and are distributed by the trigger infrastructure (peripheral trigger multiplexers).
- One-to-one triggers. A separate set exists for each counter, only connected to that counter. These triggers are used for direct trigger connections from trigger sources (such as ADC channels) to associated TCPWM counters.

Use the trigger mux registers TR_IN_SEL0 and TR_IN_SEL1 to configure which signals get routed to the Trigger_In for each TCPWM block. See Table 25-1 for all possible multiplexer settings selecting an input trigger event for a TCPWM block. For each event two constant trigger

Table 25-1. Multiplexer Selection for Input Trigger Events

Input Trigger Selection Register	Bit Field	Bits	Description
TR_IN_SEL0	CAPTURE0_SEL	7:0	Selects one of the up to 256 input triggers as a capture0 trigger. In the PWM, PWM_DT, and PWM_PR modes this trigger is used to switch the values if the compare and period registers with their buffer counterparts.
	COUNT_SEL	15:8	Selects one of the 256 input triggers as a count trigger. In QUAD mode, this is the first phase (phi A)
	RELOAD_SEL	23:16	Selects one of the 256 input triggers as a reload trigger. In QUAD mode, this is the index or revolution pulse
	STOP_SEL	31:24	Selects one of the 256 input triggers as a stop trigger. In PWM, PWM_DT, and PWM_PR modes, this is the kill trigger
TR_IN_SEL1	START_SEL	7:0	Selects one of the 256 input triggers as a start trigger. In QUAD mode, this is the second phase (phi B)
	CAPTURE1_SEL	15:8	Selects one of the up to 256 input triggers as a capture1 trigger

The following sections describe each TCPWM mode and the function of each input event in detail.

Typical operation uses the reload event once to initialize and start the counter and the stop event to stop the counter. When the counter is stopped, the start event can be used to start the counter with its counter value unmodified from when it was stopped.

If stop, reload, and start events coincide, the following precedence relationship holds:

- A stop event has higher priority than a reload event.
- A reload event has higher priority than a start event.


As a result, when a reload or start event coincides with a stop event, the reload or start event has no effect.

Before going to the counter each Trigger_IN can pass through a positive edge detector, negative edge detector, both edge detector, or pass straight through to the counter. This is controlled using TR_IN_EDGE_SEL register.

Multiple detected events are treated as follows:

- In the rising edge and falling edge modes, multiple events are effectively reduced to a single event. As a result, events may be lost.
- In the rising/falling edge mode, an even number of events are not detected and an odd number of events are reduced to a single event. This is because the rising/falling edge mode is typically used for capture events to determine the width of a pulse. The current functionality will ensure that the alternating pattern of rising and falling is maintained.

Figure 25-3. TCPWM Input Events



- In PWM mode, if PWM_IMM_KILL = 0, the rising edge kill asserts while count event is inactive, line output will not be suppressed until the next prescaled counter clock after the count event becomes active.
- In capture mode, if rising edge capture0 inputs while count event is inactive, CC0 and CC0_BUFF will get updated at the next prescaled counter clock after count event becomes active.

Typically, the count event is a constant '1' and prescaling is off. In this case, the active count prescaled counter clock is the same as the counter clock. In other cases, edge detection may detect multiple events (on the counter clock) before the next active count prescaled counter clock on which the detected event is used. Multiple detected events are treated as follows:

- In the rising edge and falling edge modes, multiple events are effectively reduced to a single event. As a result, events may be lost.
- In the rising/falling edge mode, an even number of events is not detected and an odd number of events is reduced to a single event. This is because the rising/falling edge mode is typically used for capture events to determine the width of a pulse. The current functionality will ensure that the alternating pattern of rising, falling, rising, falling, and so on is maintained.

A pass-through event will not be remembered by CLK_PERI; it will affect the functionality if it lasts and can be detected by a counter operation clock. If the pulse width of a pass-through event is less than a counter operation clock cycle, it may get lost. Pass-through detection may result in an event that is active for multiple counter clocks. This may result in undesirable behavior of the counter and its associated trigger outputs. Pass-through event detection should only be used for stop and count event types in most function modes. Pass-through mode can also be used in switch events in the PWM/PWM_DT/PWM_PR mode, if it selects the constant high as the source. In quadrature mode, both start and count event is used with pass through in X1/X2/X4 mode.


25.2.4 Synchronization of Multiple Counters

The previous sections described hardware-based event generation. In addition, software-based event generation is supported: the reload, start, stop, capture0, and capture1 events can be generated by writing to the TR_CMD registers. These are counter specific registers and allow software-based event generation only for a single counter.

Synchronized software-based event generation (such as starting multiple counters synchronously) is possible by selecting the same trigger signal in all desired counters (via TR_IN_SEL0 and TR_IN_SEL1 registers) and generating a trigger by the TR_CMD register in the PERI block.

The following figure illustrates an example of how to synchronously start multiple counters with TR CMD of the PERI block.

Figure 25-5. PERI TR CMD Synchronously Starts Counters




The following example describes the required steps to start counters synchronously:

- Configure CLOCK_CTL[55]/[54]/[53] to select same clock divider clock_div_8[1].

- Configure DIV_8_CTL[1] and DIV_CMD to generate clock divider enable signal for TCPWM. Hence, three counters will have the same clock divider enable signal; in other words, their clocks are synchronous.
 - Enable the three counters one by one - counter is enabled but will not run until start or reload event is detected.
 - Configure the three counters' start event, all selecting tr_all_cnt_in[0]; this means the three counters will run synchronously when tr_all_cnt_in[0] asserts.
 - Use TR_CMD to generate trigger pulse on TR_GROUP[9].TR_OUT[0]. All three counters will run synchronously.
- Note:** All registers listed here belongs to the PERI block (see the [Trigger Multiplexer chapter on page 493](#)).
- A software-based event is set after writing TR_CMD respective bit to '1', and cleared by hardware on the next "active count" prescaled counter clock. For some events in a specific mode, it is cleared on the next "active count" prescaled counter clocks that have a tc event. The event detection setting (TR_IN_EDGE_SEL) does not have an effect on a software-based event.

Figure 25-6. Software Trigger Command Detection by Active Count Prescaled Counter Clock



25.2.5 Trigger Outputs

Each counter has two trigger output signals (TR_OUT0 and TR_OUT1) to indicate following events. They can be routed through the trigger mux to other peripherals on the device. The bit field OUT0 in TR_OUT_SEL register is used to select one of the internal events to generate output trigger 0 (TR_OUT0), respectively the bit field OUT1 is selecting one of the internal events to generate TR_OUT1. It allows also to disable the output triggers.

- **Overflow (OV):** An overflow event indicates that in up counting mode, COUNTER equals PERIOD register, and is changed to a different value.
- **Underflow (UN):** An underflow event indicates that in a down counting mode, COUNTER equals 0, and is changed to a different value.
- **TC (Terminal Count):** A TC event is the logical OR of the underflow and overflow events
- **CC0/1_MATCH:** This event is generated when the counter is running and one of the following conditions occur:
 - Counter equals the compare value. This event is either generated when the match is about to occur (COUNTER does not equal CC0/1 and is changed to CC0/1) or when the match is about to not occur (COUNTER equals CC0/1 and is changed to a different value).
 - A capture event has occurred and the CC0 (CC1) and CC0_BUFF (CC1_BUFF) registers are updated.
- **LINE_OUT:** A PWM output signal
- **DISABLED:** Output trigger is disabled

The selection of the events for the output trigger generation is done by the TR_OUT_SEL register. It also allows disabling the output triggers.

Note: These signals only remain high for two cycles of CLK_SYS. For reliable operation, the condition that causes this trigger should be a maximum of one quarter of the CLK_SYS. For example, if the CLK_SYS is running at 24 MHz, the condition causing the trigger should occur at a frequency equal to or less than 6 MHz.

When LINE_OUT is selected for output triggers, output trigger will bypass two cycle pulses generation logic and directly output LINE_OUT.

The generated triggers have two main uses:

- **Initiating a DW/DMA data transfer.** For example, in PWM mode with an up counting timer, the overflow can be used to transfer new period and compare values from memory to the counters' PERIOD_BUFF and CC0_BUFF registers.
- **Reconstruction of a PWM signal in a programmable digital component.** As documented in [25.3.4 Pulse Width Modulation \(PWM\) Mode](#), the PWM line output signal is derived from the cc0_match (cc1_match), underflow, and overflow internal events. By making these internal events available as output triggers, other components can reconstruct and potentially modify the PWM signal (note the mentioned frequency restrictions).

25.2.6 Internal Events

25.2.6.1 Underflow Event

An underflow event indicates that in down counting, COUNTER equals zero, and is changed to a different value. Reload will also generate underflow event in some specific mode. [Table 25-3](#) summarizes the underflow generation of each function mode.

Table 25-3. Underflow Generation

MODE	UP	DOWN	UPDN1	UPDN2
TIMER	Counter is decrementing and changes from a state in which COUNTER equals 0. Reload event in DOWN, UPDN1, and UPDN2 modes.			
CAPTURE	Counter is decrementing and changes from a state in which COUNTER equals 0. Reload event in DOWN, UPDN1, and UPDN2 modes.			
QUAD	QUAD_RANGE0: Not used QUAD_RANGE0_CMP: Not used QUAD_RANGE1_CMP: Counter value COUNTER equals 0 and is decrementing. QUAD_RANGE1_CAPT: Counter value COUNTER equals 0 and is decrementing.			
PWM	Counter is decrementing and changes from a state in which COUNTER equals 0. Reload event in DOWN, UPDN1, and UPDN2 modes.			
PWM_DT	Counter is decrementing and changes from a state in which COUNTER equals 0. Reload event in DOWN, UPDN1, and UPDN2 modes.			
PWM_PR	Not used			
SR	Not used			

25.2.6.2 Overflow Event

An overflow event indicates that in up counting, COUNTER equals PERIOD, and is changed to a different value. Reload will also generate overflow event in some specific mode. [Table 25-4](#) summarizes the overflow generation of each function mode.

Table 25-4. Overflow Generation

MODE	UP	DOWN	UPDN1	UPDN2
TIMER	Counter is incrementing and changes from a state in which COUNTER equals PERIOD. Reload event in UP count mode.			
CAPTURE	Counter is incrementing and changes from a state in which COUNTER equals PERIOD. Reload event in UP count mode.			
QUAD	QUAD_RANGE0: Not used QUAD_RANGE0_CMP: Not used QUAD_RANGE1_CMP: Counter value COUNTER equals PERIOD and is incrementing. QUAD_RANGE1_CAPT: Counter value COUNTER equals PERIOD and is incrementing.			
PWM	Counter is incrementing and changes from a state in which COUNTER equals PERIOD. Reload event in UP count mode.			
PWM_DT	Counter is incrementing and changes from a state in which COUNTER equals PERIOD. Reload event in UP count mode.			
PWM_PR	Not used			
SR	Not used			

25.2.6.3 TC Event

A tc (terminal count) event is the logical OR of the underflow and overflow events. An exception is that reload event will generate an underflow or overflow, but not a tc event. In quadrature mode, index will generate a tc event. [Table 25-5](#) summarizes the tc generation of each function mode.

Table 25-5. TC Generation

MODE	UP	DOWN	UPDN1	UPDN2
TIMER	Overflow	Underflow	Underflow	Logic OR of overflow and underflow
CAPTURE	Overflow	Underflow	Underflow	Logic OR of overflow and underflow
QUAD	QUAD_RANGE0: ■ Index event. QUAD_RANGE0_CMP: ■ Counter value COUNTER equals 0 or 0xFFFF/0xFFFFFFFF in “wraparound capture” mode. ■ Index or capture0 event in “index capture” mode. QUAD_RANGE1_CMP: ■ Counter value COUNTER equals 0 and decrementing (underflow), or PERIOD and incrementing (overflow). ■ Index event. QUAD_RANGE1_CAPT: ■ Same as QUAD_RANGE1_CMP.			
PWM	Overflow	Underflow	Underflow	Logic OR of overflow and underflow
PWM_DT	Overflow	Underflow	Underflow	Logic OR of overflow and underflow
PWM_PR	Counter changes from a state in which COUNTER equals PERIOD.			
SR	Not used			

25.2.6.4 cc0_match (cc1_match) Event

A cc0_match event indicates that the COUNTER equals CC0. This event is either generated when COUNTER is about to change to CC0, or when COUNTER equals CC0 and is about to change to a different value. A special case is for 0 or 100 percent duty cycle generation in PWM mode; for more details, see the [25.3.4 Pulse Width Modulation \(PWM\) Mode](#). In other specific operation modes, the event is used to indicate that the CC0/CC0_BUFF registers are updated. cc1_match is generated per state of COUNTER and CC1, other behavior is same as cc0_match. [Table 25-6](#) and [Table 25-7](#) summarize the cc0/1_match generation of each function mode.

Table 25-6. cc0_match Generation

MODE	UP	DOWN	UPDN1	UPDN2
Timer	Counter changes from a state in which COUNTER equals CC0.			
CAPTURE	Capture0 event			
QUAD	QUAD_RANGE0: ■ Counter value COUNTER equals 0 or 0xFFFF. ■ Index event. QUAD_RANGE0_CMP: ■ Counter changes to a state in which COUNTER equals CC0. QUAD_RANGE1_CMP: ■ Same as QUAD_RANG0_CMP. QUAD_RANGE1_CAPT: ■ Capture0 event.			
PWM	Counter changes to a state in which COUNTER equals CC0.		COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC0. If a second compare function is present in a counter group, CC0_MATCH_DOWN_EN/CC0_MATCH_UP_EN will enable/disable cc0_match generation.	
PWM_DT	Counter changes to a state in which COUNTER equals CC0.		COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC0. If a second compare function is present in a counter group, CC0_MATCH_DOWN_EN/CC0_MATCH_UP_EN will enable/disable cc0_match generation.	
PWM_PR	Counter changes from a state in which COUNTER equals CC0.			
SR	Counter changes to a state in which COUNTER equals CC0.			

Table 25-7. cc1_match Generation

MODE	UP	DOWN	UPDN1	UPDN2
Timer	Counter changes from a state in which COUNTER equals CC1.			
CAPTURE	Capture1 event			
QUAD	QUAD_RANGE0: ■ Not used. QUAD_RANGE0_CMP: ■ Counter changes to a state in which COUNTER equals CC1. QUAD_RANGE1_CMP: ■ Same as QUAD_RANG0_CMP. QUAD_RANGE1_CAPT: ■ Capture1 event.			
PWM	Counter changes to a state in which COUNTER equals CC1.		COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC1. If a second compare function is present in a counter group, CC1_MATCH_DOWN_EN/CC1_MATCH_UP_EN will enable/disable cc1_match generation.	
PWM_DT	Counter changes to a state in which COUNTER equals CC1.		COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC1. If a second compare function is present in a counter group, CC1_MATCH_DOWN_EN/CC1_MATCH_UP_EN will enable/disable cc1_match generation.	
PWM_PR	Counter changes from a state in which COUNTER equals CC1.			
SR	Counter changes to a state in which COUNTER equals CC1.			

25.2.7 Interrupts

The TCPWM block provides a dedicated interrupt output for each counter. Interrupts are counter mode specific and can be generated for a Terminal Count (TC) or Compare/Capture0/1 (CC0/1) event. A TC is the logical OR of the OV and UN events.

Four registers are used to handle interrupts in this block, as shown in [Table 25-8](#).

Table 25-8. Interrupt Register

Interrupt Registers	Bits	Name	Description
INTR (Interrupt request register)	0	TC	This bit is set to '1', when a terminal count is detected. Write '1' to clear this bit.
	1	CC0_MATCH	This bit is set to '1' when the counter value matches capture/compare0 (CC0) register value. Write '1' to clear this bit.
	2	CC1_MATCH	This bit is set to '1' when the counter value matches capture/compare1 (CC1) register value. Write '1' to clear this bit.
INTR_SET (Interrupt set request register)	0	TC	Write '1' to set the corresponding bit in the interrupt request register. When read, this register reflects the interrupt request register status.
	1	CC0_MATCH	Write '1' to set the corresponding bit in the interrupt request register. When read, this register reflects the interrupt request register status.
	2	CC1_MATCH	Write '1' to set the corresponding bit in the interrupt request register. When read, this register reflects the interrupt request register status.
INTR_MASK (Interrupt mask register)	0	TC	Mask bit for the corresponding TC bit in the interrupt request register.
	1	CC0_MATCH	Mask bit for the corresponding CC_MATCH0 bit in the interrupt request register.
	2	CC1_MATCH	Mask bit for the corresponding CC_MATCH1 bit in the interrupt request register.
INTR_MASKED (Interrupt masked request register)	0	TC	Logical AND of the corresponding TC request and mask bits.
	1	CC0_MATCH	Logical AND of the corresponding CC_MATCH0 request and mask bits.
	2	CC1_MATCH	Logical AND of the corresponding CC_MATCH1 request and mask bits.

25.2.8 Debug Mode

The TCPWM counters support debugging. It can be configured per counter if the counter operation continues or pauses in debug state (for example, after running to a break point). This feature is especially intended when using a TCPWM counter as an OS timer. It is realized by gating the PCLK_TCPWM[x]_CLOCKS[y] when entering debug state by setting the DBG_FREEZE_EN bit to '1' in the CTRL register and asserting a debug pause trigger.

In a multicore environment 'debug state' means that at least one of the CPUs is in the debug state. In cases where only one CPU is debugged but another or multiple other CPUs are continuously running, the user can configure the counter via the debugger to continue or pause depending on which CPU is using the counter.

Note: The trigger input cannot be asserted when the counter is in debug state.

25.2.9 PWM Outputs

The PWM, PWM_DT, PWM_PR, and SR operation modes produce two output signals:

- A PWM LINE_OUT output signal
- A complementary PWM LINE_COMPL_OUT output signal (inverted version of LINE_OUT)

Note that in PWM and PWM_DT modes the CC0_match, CC1_match, underflow, and overflow internal event conditions are used to drive LINE_OUT and LINE_COMPL_OUT, by configuring the TR_PWM_CTRL register (Table 25-9). In PWM_PR and SR modes, line output is not controlled by TR_PWM_CTRL.

Table 25-9. Configuring Output Line for OV, UN, and CC0/1 Conditions

Field	Bit	Value	Event	Description
CC0_MATCH_MODE Default Value = 3	1:0	0	Set LINE_OUT to '1	Configures output line on a compare match (CC0) event
		1	Clear LINE_OUT to '0	
		2	Invert LINE_OUT	
		3	No change	
OVERFLOW_MODE Default Value = 3	3:2	0	Set LINE_OUT to '1	Configures output line on an overflow (OV) event
		1	Clear LINE_OUT to '0	
		2	Invert LINE_OUT	
		3	No change	
UNDERFLOW_MODE Default Value = 3	5:4	0	Set LINE_OUT to '1	Configures output line on an underflow (UN) event
		1	Clear LINE_OUT to '0	
		2	Invert LINE_OUT	
		3	No change	
CC1_MATCH_MODE Default Value = 3	7:6	0	Set LINE_OUT to '1	Configures output line on a compare match (CC1) event
		1	Clear LINE_OUT to '0	
		2	Invert LINE_OUT	
		3	No change	

The generation of PWM output signals is a multi-step process. Both LINE_OUT and LINE_COMPL_OUT are generated from the PWM signal line. The PWM signal line is generated as per the state of cc0_match, cc1_match, underflow, and overflow internal events, as specified by the counter's TR_PWM_CTRL register. For each internal event, the TR_PWM_CTRL register specifies how the event affects the output LINE_OUT

- The output is set to '0'
- The output is set to '1'
- The output is inverted
- The output is not affected


In case the internal cc0_match event generates at the same time when internal underflow or overflow event generates, cc0_match will take effect after LINE_OUT changes state per settings of underflow/overflow. cc1_match will take effect after cc0_match. Table 25-10 lists some examples to show the mechanism.

Table 25-10. LINE_OUT Construction Example

Coincide Case	Overflow	Underflow	CC0_match	CC1_match	LINE_OUT
CC0_match and overflow	CLEAR	Don't care	INVERT	Don't care	1 (SET)
CC0_match and underflow	Don't care	SET	INVERT	Don't care	0 (CLEAR)
CC0_match and CC1_match	Don't care	Don't care	SET	CLEAR	CLEAR
CC0_match and CC1_match and overflow	INVERT	Don't care	INVERT	INVERT	INVERT

The following figure illustrates the process of LINE_OUT and LINE_COMPL_OUT generation.

Figure 25-7. PWM Output Generation Process in PWM/PWM_DT/PWM_PR/SR Mode



Generally, LINE_OUT output reflects the state of PWM signal line and LINE_COMPL_OUT output reflects the inverted version of line. The line behavior depends on different function modes. Furthermore, some other factors will impact LINE_OUT and LINE_COMPL_OUT according to different function modes; they include 'line selection', 'dead time insertion', 'kill function', and 'line polarity'.

■ PWM signal line generation

- ❑ In PWM/PWM_DT mode, line is constructed by internal events underflow, overflow, cc0_match, and cc1_match per settings of TR_PWM_CTRL.
- ❑ In PWM_PR mode, line reflects the state of comparison b/w COUNTER and CC0.
- ❑ In SR mode, line is the shift output of shift register (COUNTER).

■ Line selection (available in counter groups supporting Advanced Motor Control)

- ❑ LINE_OUT and LINE_COMPL_OUT can individually select different output according to the LINE_SEL.OUT_SEL and LINE_SEL.COMPL_SEL register settings. This functionality only works in PWM and PWM_PR modes.
- ❑ LINE_OUT and LINE_COMPL_OUT can individually selects Low, High, Line, Inverted line, and Hi-Z. When it selects Hi-Z, line_out_en and line_compl_out_en will be low.

■ Dead time insertion

- ❑ Dead time insertion functionality is mutually exclusive with line selection functionality, it only works in PWM_DT mode.
- ❑ Dead time works on both line signal and inverted version of line signal.

■ PWM disable mode

Specifies the behavior of the line_out and line_out_compl_out PWM outputs while the TCPWM counter is disabled (ENABLED bit set to '0' in the CTRL register) or stopped. The PWM output behavior is determined by the PWM_DISABLE_MODE bit field in the CTRL register. There are four options:

- ❑ Z (PWM_DISABLE_MODE = 0)

When the counter is disabled the line_out and line_compl_out PWM outputs are not driven by the TCPWM. Instead, the port default level configuration applies, for example, "Z" (high impedance). When the counter is stopped on a stop event, the PWM outputs are deactivated and the polarity is defined by the QUAD_ENCODING_MODE bit field in the CTRL register.

- ❑ Retain (PWM_DISABLE_MODE = 1)

When the counter is disabled or stopped on a stop event, the PWM outputs are retained (keep their previous levels). While the counter is disabled or stopped the PWM outputs can be changed via LINE_SEL (this is only valid for counter groups with parameter GRP_SMC_PRESENT = 1).

- ❑ Low (PWM_DISABLE_MODE = 2)

When the counter is disabled or stopped on a stop event, the line_out PWM output is driven as a fixed '0' and the line_compl_out PWM output is driven as a fixed '1'.

- ❑ High (PWM_DISABLE_MODE = 3)

When the counter is disabled or stopped on a stop event, the line_out PWM output is driven as a fixed '1' and the line_compl_out PWM output is driven as a fixed '0'.

- Kill function
 - Kill function works in PWM, PWM_DT, and PWM_PR modes. It does not work in SR mode.
 - Kill works on both line and inverted version of line, and there are several kill function modes supported.
- Polarity for LINE_OUT and LINE_COMPL_OUT
 - Polarity inversion is used to determine the LINE_OUT and LINE_COMPL_OUT output signal values.
 - CTRL.QUADRATURE_ENCODING_MODE[0] is for LINE_OUT polarity and CTRL.QUADRATURE_ENCODING_MODE[1] is for LINE_COMPL_OUT polarity.
 - When the counter is not enabled in reset state or not running (temporarily stopped or killed), the PWM output signals values are determined by their respective polarity settings.

Details of PWM line signal generation is described in separate function mode sections later in the document.

Besides LINE_OUT and LINE_COMPL_OUT, each counter provides line_out_en and line_compl_out_en, which reflect the counter enable state. These two output enable signals can be used to disable GPIO output after counter is disabled.

TCPWM block has four ports LINE_OUT[counter group number*256-1:0], LINE_COMPL_OUT[counter group number*256-1:0], line_out_en[counter group number*256-1:0], and line_compl_out_en[counter group number*256-1:0].

25.2.10 Power Modes

The TCPWM block works in Active and Sleep modes. The TCPWM block is powered from VCCACT. The retention MMIO registers are powered in DeepSleep with VCCRET, but unpowered in Hibernate mode. The configuration registers and other logic are powered in DeepSleep mode to keep the states of configuration registers. See [Table 25-11](#) for details.

Table 25-11. Power Modes in TCPWM Block

Power Mode	Block Status
Active	This block is fully operational in this mode with clock running and power switched on.
Sleep	The CPU is in sleep but the block is still functional in this mode. All counter clocks are on.
DeepSleep	Both power and clocks to the block are turned off, but configuration registers retain their states.
Hibernate	In this mode, the power to this block is switched off. Configuration registers will lose their state. No CLK_PERI is provided.

25.3 Operation Modes

The counter block can function in seven operational modes, as shown in [Table 25-12](#). The MODE [26:24] field of the counter control register (CTRL) configures the counter in the specific operational mode.

Table 25-12. Operational Mode Configuration

Mode	MODE Field [26:24]	Description
Timer	000	The counter increments or decrements by '1' at every counter clock cycle in which a count event is detected. The Compare/Capture register is used to compare the count.
Capture	010	The counter increments or decrements by '1' at every counter clock cycle in which a count event is detected. A capture event copies the counter value into the capture register.
Quadrature	011	Quadrature decoding. The counter is decremented or incremented based on two phase inputs according to an X1, X2, and X4 decoding scheme or to the rotary count mode.
PWM	100	Pulse width modulation.
PWM_DT	101	Pulse width modulation with dead time insertion.
PWM_PR	110	Pseudo-random PWM using a 16- or 32-bit linear feedback shift register (LFSR) with programmable length to generate pseudo-random noise.
SR	111	Shift register mode

The counter can be configured to count up, down, and up/down by setting the UP_DOWN_MODE[17:16] field in the CTRL register, as shown in [Table 25-13](#).

Table 25-13. Counting Mode Configuration (except Quadrature mode)

Counting Modes	UP_DOWN_MODE [17:16]	Description
UP Counting Mode	00	Increments the counter until the period value is reached. A Terminal Count (TC) condition is generated when the counter changes from the period value.
DOWN Counting Mode	01	Decrements the counter from the period value until 0 is reached. A TC condition is generated when the counter changes from a value of '0'.
UP/DOWN Counting Mode 1	10	Increments the counter until the period value is reached, and then decrements the counter until '0' is reached. A TC condition is generated only when the counter changes from a value of '0'.
UP/DOWN Counting Mode 2	11	Similar to up/down counting mode 1 but a TC condition is generated when the counter changes from '0' and when the counter value changes from the period value.

In Quadrature mode this field acts as QUAD_RANGE_MODE field selecting between different counter ranges, reload value, and compare/capture behavior.

Table 25-14. Counting Mode Configuration for Quadrature Mode

Counting Modes	UP_DOWN_MODE [17:16]	Description
UP Counting Mode	00	Increments the counter until the period value is reached. A TC condition is generated when the counter changes from the period value.
DOWN Counting Mode	01	Decrements the counter from the period value until 0 is reached. A TC condition is generated when the counter changes from a value of '0'.
UP/DOWN Counting Mode 1	10	Increments the counter until the period value is reached, and then decrements the counter until '0' is reached. A TC condition is generated only when the counter changes from a value of '0'.
UP/DOWN Counting Mode 2	11	Similar to up/down counting mode 1 but a TC condition is generated when the counter changes from '0' and when the counter value changes from the period value.

25.3.1 Timer Mode

The timer mode is commonly used to measure the time of occurrence of an event or to measure the time difference between two events. The timer functionality increments/decrements a counter between 0 and the value stored in the PERIOD register. When the counter is running, the count value stored in the COUNTER register is compared with the compare/capture register (CC0 and CC1). When COUNTER equals CC0, the cc0_match event is generated, even-handedly when COUNTER equals CC1, the cc1_match event is generated.

Timer functionality is typically used for one of the following:

- Timing a specific delay - the count event is a constant '1'.

- Counting the occurrence of a specific event - the event should be connected as an input trigger and selected for the count event.

Table 25-15. Timer Mode Trigger Input Description

Trigger Inputs	Usage
Reload	Initializes and starts the counter. Behavior is dependent on UP_DOWN_MODE: ■ COUNT_UP: The counter is set to '0' and count direction is set to 'up'. ■ COUNT_DOWN: The counter is set to PERIOD and count direction is set to 'down'. ■ COUNT_UPDN1/2: The counter is set to '1' and count direction is set to 'up'. Can be used when the counter is running or not running.
Start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is dependent on UP_DOWN_MODE. When the counter is not running: ■ COUNT_UP: The count direction is set to 'up'. ■ COUNT_DOWN: The count direction is set to 'down'. ■ COUNT_UPDN1/2: The count direction is not modified. Note that when the counter is running, the start event has no effect. Can be used when the counter is running or not running.
Stop	Stops the counter.
Count	Count event increments/decrements the counter.
Capture0	Not used.
Capture1	Not used.

Incrementing and decrementing the counter is controlled by the count event and the counter clock, PCLK_TCPWM[x]_CLOCKS[y]. Typical operation will use a constant '1' count event and PCLK_TCPWM[x]_CLOCKS[y] without prescaling. Advanced operations are also possible; for example, the counter event configuration can decide to count the rising edges of a synchronized input trigger.

Table 25-16. Timer Mode Supported Features

Supported Features	Description
Clock prescaling	Prescales the PCLK_TCPWM[x]_CLOCKS[y].
One shot	Counter is stopped by hardware, on a tc event. In COUNT_UPDN2, counter is stopped on tc event when underflow.
Auto reload CC	CC0 and CC0_BUFF are exchanged on a cc0_match event (when specified by CTRL.AUTO_RELOAD_CC, no input event is required). CC1 and CC1_BUFF are exchanged on a cc1_match event (when specified by CTRL.AUTO_RELOAD_CC).
Up/down modes	Specified by UP_DOWN_MODE: ■ COUNT_UP: The counter counts from 0 to PERIOD. ■ COUNT_DOWN: The counter counts from PERIOD to 0. ■ COUNT_UPDN1/2: The counter counts from 1 to PERIOD and back to 0.

Table 25-17 lists the trigger outputs and the conditions when they are triggered.

Table 25-17. Timer Mode Trigger Outputs

Trigger Outputs	Description
cc0_match	Counter changes from a state in which COUNTER equals CC0.


Table 25-17. Timer Mode Trigger Outputs

Trigger Outputs	Description
cc1_match	Counter changes from a state in which COUNTER equals CC1.
Underflow (UN)	Counter is decrementing and changes from a state in which COUNTER equals 0.
Overflow (OV)	Counter is incrementing and changes from a state in which COUNTER equals PERIOD.
TC	Specified by UP_DOWN_MODE: ■ COUNT_UP: tc event is the same as the overflow event. ■ COUNT_DOWN: tc event is the same as the underflow event. ■ COUNT_UPDN1: tc event is the same as the underflow event. ■ COUNT_UPDN2: tc event is the same as the logical OR of the overflow and underflow events. Reload will generate underflow/overflow, but not generate tc.

Table 25-18. Timer Mode PWM Outputs

PWM Outputs	Description
LINE_OUT	Not used.
LINE_COMPL_OUT	Not used.

Figure 25-8. Timer Functionality



Notes:

- The triggers tr_out0 and tr_out1 are generated based on the internal events cc0_match, cc1_match, underflow, overflow, and tc respectively (selection is done by the TR_OUT_SEL register).
- The timer functionality only uses PERIOD (and not PERIOD_BUFF).
- It is not recommended to write to COUNTER when the counter is running.

Figure 25-9 illustrates a timer in up-counting mode. The counter is initialized (to 0) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 4, an overflow and tc event are generated.
- When the counter changes from a state in which COUNTER is 2, a cc0_match event is generated.
- PERIOD is 4, resulting in an effective repeating counter pattern of $4+1 = 5$ counter clock periods. The CC0 register is 2, and sets the condition for a cc0_match event.

A constant count event of '1' and PCLK_TCPWM[x]_CLOCKS[y] without prescaling is used in the following scenarios. If the count event is '0' and a reload event is triggered, the reload will only be registered on the first clock edge when the count event is '1'.

Figure 25-9. Timer in Up-counting Mode




Figure 25-10 illustrates a timer in “one-shot” operation mode. Note that the counter is stopped on a tc event.

Figure 25-10. Timer in One-shot Mode




Figure 25-11 illustrates clock prescaling. Note that the counter is only incremented every other counter cycle.

Figure 25-11. Timer Clock Prescaling




Figure 25-12 illustrates a counter that is initialized and started (reload event), stopped (stop event), and continued/started (start event). Note that the counter does not change value when it is not running (STATUS.RUNNING).

Figure 25-12. Counter Start/Stopped/Continued




Figure 25-13 illustrates a timer that uses CC0/1 and CC0/1_BUFF registers. Note that CC0/1 and CC0/1_BUFF register contents are exchanged on a cc0/1_match event.

Figure 25-13. Use of CC0 and CC0_BUFF Register Bits




Figure 25-14 illustrates a timer in down-counting mode. The counter is initialized (to PERIOD) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 0, an underflow and tc event are generated.
- When the counter changes from a state in which COUNTER is 2, a cc0_match event is generated.
- PERIOD is 4, resulting in an effective repeating counter pattern of $4+1 = 5$ counter clock periods.

MODE = TIMER
UP_DOWN_MODE = COUNT_DOWN

COUNTER starts with PERIOD

period is PERIOD+1

reload

COUNTER

PERIOD = 4

CC0 = 2

Underflow (UN)

Overflow (OV)

Terminal Count (TC)


cc0_match

no TC event

CC0 event on leaving the COUNTER value

Notes:

- Figure 25-15. Timer in Up/Down Counting Mode 1



Notes:

- When CC0 is 0, the cc0_match event is generated at the start of the period (when the counter changes from a state in which COUNTER is 0).
- When CC0 is PERIOD, the cc0_match event is generated at the middle of the period (when the counter changes from a state in which COUNTER is PERIOD).

Figure 25-16. Up/Down Counting Mode with Different CC Values





Figure 25-17 illustrates a timer in up/down counting mode 2. This mode is same as up/down counting mode 1, except for the TC event, which is generated when either underflow or overflow event occurs.

Figure 25-17. Up/Down Counting Mode 2



25.3.1.1 Configuring Counter for Timer Mode

The steps to configure the counter for Timer mode of operation and the affected register bits are as follows.

1. Disable the counter by writing '0' to the ENABLE bit of the CTRL register.
2. Select Timer mode by writing '000' to the MODE[26:24] field of the CTRL register.
3. Set the required 16- or 32-bit period in the PERIOD register.
4. Set the 16- or 32-bit compare value in the CC0 register and the buffer compare value in the CC0_BUFF register.
5. Set AUTO_RELOAD_CC0 field of the CTRL register, if required to switch values at every CC condition.
6. Set clock prescaling by writing to the DT_LINE_OUT_L[7:0] field of the DT register.
7. Set the direction of counting by writing to the UP_DOWN_MODE[17:16] field of the CTRL register.
8. The timer can be configured to run either in continuous mode or one-shot mode by writing 0 or 1, respectively to the ONE_SHOT[18] field of the CTRL register.
9. Set the TR_IN_SEL0 or TR_IN_SEL1 register to select the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
10. Set the TR_IN_EDGE_SEL register to select the edge of the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
11. If required, set the interrupt upon TC or CC0_MATCH or CC1_MATCH condition.

12. Enable the counter by writing '1' to ENABLED bit of the CTRL register. A start trigger must be provided through firmware (START bit in the TR_CMD register) to start the counter if the hardware start signal is not enabled.

25.3.2 Capture Mode

The capture functionality increments and decrements a counter between 0 and PERIOD. When the capture event is activated the counter value COUNTER is copied to CC0/1 (and CC0/1 is copied to CC0/1_BUFF).

The capture functionality can be used to measure the width of a pulse (connected as one of the input triggers and used as capture event).

The capture event can be triggered through the capture trigger input or through a firmware write to CAPTURE0/1 bit in the TR_CMD command register.

Table 25-19. Capture Mode Trigger Input Description

Generated Events	Usage
Reload	Sets the counter value and starts the counter. Behavior is dependent on UP_DOWN_MODE: ■ COUNT_UP: The counter is set to '0' and count direction is set to 'up'. ■ COUNT_DOWN: The counter is set to PERIOD and count direction is set to 'down'. ■ COUNT_UPDN1/2: The counter is set to '1' and count direction is set to 'up'. Can be used when the counter is running or not running.
Start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is dependent on UP_DOWN_MODE: ■ COUNT_UP: The count direction is set to 'up'. ■ COUNT_DOWN: The count direction is set to 'down'. ■ COUNT_UPDN1/2: The count direction is not modified. Note that when the counter is running, the start event has no effect. Can be used when the counter is running or not running.
Stop	Stops the counter.
Count	Count event increments/decrements the counter.
Capture0	Copies the counter value to CC0 and copies CC0 to CC0_BUFF.
Capture1	Copies the counter value to CC1 and copies CC1 to CC1_BUFF.

Table 25-20. Supported Features of CAPTURE

Supported Features	Description
Clock prescaling	Prescales the PCLK_TCPWM[x]_CLOCKS[y].
One shot	Counter is stopped by hardware, after a single period of the counter: ■ COUNT_UP: on an overflow event. ■ COUNT_DOWN, COUNT_UPDN1/2: on an underflow event.
Up/down modes	Specified by UP_DOWN_MODE: ■ COUNT_UP: The counter counts from 0 to PERIOD. ■ COUNT_DOWN: The counter counts from PERIOD to 0. ■ COUNT_UPDN1/2: The counter counts from 1 to PERIOD and back to 0.

Table 25-21. Internal Events of CAPTURE

Internal Events	Description
CC0_match	CC0 is copied to CC0_BUFF and counter value is copied to CC0 (cc0_match equals capture event).
CC1_match	CC1 is copied to CC1_BUFF and counter value is copied to CC1 (cc1_match equals capture event).
Underflow (UN)	Counter is decrementing and changes from a state in which COUNTER equals 0.
Overflow (OV)	Counter is incrementing and changes from a state in which COUNTER equals PERIOD.
TC	Specified by UP_DOWN_MODE: ■ COUNT_UP: tc event is the same as the overflow event. ■ COUNT_DOWN: tc event is the same as the underflow event. ■ COUNT_UPDN1: tc event is the same as the underflow event. ■ COUNT_UPDN2: tc event is the same as the logical OR of the overflow and underflow events. Reload will generate underflow/overflow, but not generate tc

Table 25-22. Capture Mode PWM Outputs

PWM Outputs	Description
LINE_OUT	Not used.
LINE_COMPL_OUT	Not used.

Figure 25-18. Capture Functionality




Figure 25-19 illustrates capture behavior in the up-counting mode.

Notes:

- The capture event detection uses rising edge detection. As a result, the capture event is remembered until the next active count prescaled counter clock.
- When a capture event occurs, COUNTER is copied into CC0/1. CC0/1 is copied to CC0/1_BUFF register.
- A cc_match event is generated when the counter value is captured.

8. Set the TR_IN_EDGE_SEL register to select the edge that causes the event (Reload, Start, Stop, Capture0/1, and Count).
9. If required, set the interrupt upon TC or CC0_MATCH or CC1_MATCH condition.
10. Enable the counter by writing '1' to the ENABLED bit in CTRL register. A start trigger must be provided through firmware (START bit in TR_CMD register) to start the counter if the hardware start signal is not enabled.

25.3.3 Quadrature Decoder Mode

Quadrature functionality increments and decrements a counter between 0 and 0xFFFF or 0xFFFFFFFF (32-bit mode) or PERIOD (depending on QUAD_RANGE_MODE). Counter updates are under control of quadrature signal inputs: index, phiA, and phiB. The index input is used to indicate an absolute position. The phiA and phiB inputs are used to determine a change in position (the rate of change in position can be used to derive speed).

Table 25-23 shows an overview of supported range modes, which varies between different maximum counter values uses capture and compare functionalities.

Table 25-23. Quadrature Mode Functionality Overview

Supported Range Modes (QUAD_RANGE_MODE)	Description
QUAD_RANGE0	Counter range is between 0x0000 and 0xFFFF/0xFFFFFFFF (32-bit mode).
QUAD_RANGE0_CMP	Counter range is between 0x0000 and 0xFFFF/0xFFFFFFFF (32-bit mode). In this mode a compare function is supported during quadrature decoding using the CC0/CC0_BUFF (CC1/CC1_BUFF) registers and the cc0_match (cc1_match) event.
QUAD_RANGE1_CMP	The compare functionality is the same as for QUAD_RANGE0_CMP mode. The counter range can be set between 0x0000 and PERIOD.
QUAD_RANGE1_CAPT	Counter range is between 0x0000 and PERIOD. Quadrature functionality in QUAD_RANGE1_CAPT mode provides the same functionality as the QUAD_RANGE1_CMP mode with the only difference that 1 or 2 capture functions are available instead of 1 or 2 compare functions.

The quadrature inputs are mapped onto triggers (as described in Table 25-24).

Table 25-24. Quadrature Mode Trigger Input Description

Trigger Input	Usage
reload/index	This event acts as a quadrature index input. It initializes the counter to the counter midpoint 0x8000 (16-bit) or 0x80000000 (32-bit mode) and starts the quadrature functionality. Rising edge event detection or falling edge detection mode should be used. ■ QUAD_RANGE0: initialize counter to 0x8000/0x80000000 (midpoint) ■ QUAD_RANGE0_CMP: initialize counter to 0x8000/0x80000000 (midpoint) ■ QUAD_RANGE1_CMP: initialize counter to 0. ■ QUAD_RANGE1_CAPT: initialize counter to 0.
start/phiB	This event acts as a quadrature phiB input. Pass-through (no edge detection) event detection mode should be used for X1, X2, or X4.

Table 25-24. Quadrature Mode Trigger Input Description

Trigger Input	Usage
stop	Stops the quadrature functionality. When quadrature stops, reload must be used to start the quadrature.
count/phiA	This event acts as a quadrature phiA input. Pass-through (no edge detection) event detection mode should be used for X1, X2, or X4.
Capture0	■ QUAD_RANGE0: Not used ■ QUAD_RANGE0_CMP: Second index ■ QUAD_RANGE1_CMP: Not used. ■ QUAD_RANGE1_CAPT: Capture event to copy COUNTER to CC0 and CC0 to CC0_BUFF.
Capture1	Available only in counter groups with second capture function. ■ QUAD_RANGE0: Not used ■ QUAD_RANGE0_CMP: Not used ■ QUAD_RANGE1_CMP: Not used. ■ QUAD_RANGE1_CAPT: Second capture event to copy COUNTER to CC1 and CC1 to CC1_BUFF.

Table 25-25. Quadrature Mode Supported Features

Supported Features	Description
Supported encoding modes (QUAD_ENCODING_MODE)	Four encoding schemes for the phiA and phiB inputs are supported (as specified by QUAD_ENCODING_MODE [21:20] bit field in the CTRL register): ■ X1 encoding. ■ X2 encoding. ■ X4 encoding. ■ Up/down rotary count mode

Note: Clock prescaling is not supported and the count event is used as a quadrature input phiA. Thus, the quadrature functionality operates on the counter clock (PCLK_TCPWM[x]_CLOCK[y]), rather than on an active count prescaled counter clock.

Table 25-26 summarize the trigger outputs dependent on different QUAD range modes.

Table 25-26. Quadrature Mode Trigger Output Description

Trigger Outputs	QUAD Range Mode	Description
cc0_match	QUAD_RANGE0	Counter value COUNTER equals 0 or 0xFFFF/0xFFFFFFFF (32-bit mode) reload/index event
	QUAD_RANGE0_CMP	Counter changes to a state in which COUNTER equals CC0
	QUAD_RANGE1_CMP	Same as QUAD_RANGE0_CMP
	QUAD_RANGE1_CAPT	Capture0 event
cc1_match	QUAD_RANGE0	Not used
	QUAD_RANGE0_CMP	Counter changes to a state in which COUNTER equals CC1
	QUAD_RANGE1_CMP	Same as QUAD_RANGE0_CMP
	QUAD_RANGE1_CAPT	Capture1 event
underflow	QUAD_RANGE0	Not used
	QUAD_RANGE0_CMP	Not used
	QUAD_RANGE1_CMP	Counter value COUNTER equals 0 and is decrementing
	QUAD_RANGE1_CAPT	Counter value COUNTER equals 0 and is decrementing
overflow	QUAD_RANGE0	Not used
	QUAD_RANGE0_CMP	Not used
	QUAD_RANGE1_CMP	Counter value COUNTER equals PERIOD and is incrementing
	QUAD_RANGE1_CAPT	Counter value COUNTER equals PERIOD and is incrementing

Table 25-26. Quadrature Mode Trigger Output Description


Trigger Outputs	QUAD Range Mode	Description
tc	QUAD_RANGE0	Index event
	QUAD_RANGE0_CMP	Counter value COUNTER equals 0 or 0xFFFF/0xFFFFFFFF Index or capture on index event (specified by AUTO_RELOAD_PERIOD in the CTRL register)
	QUAD_RANGE1_CMP	Counter value COUNTER equals 0 and is decrementing (underflow) or PERIOD and is incrementing (overflow) Index event
	QUAD_RANGE1_CAPT	Same as QUAD_RANGE1_CMP

Table 25-27. Quadrature Mode PWM Outputs

PWM Outputs	Description
LINE_OUT	Not used.
LINE_COMPL_OUT	Not used.

Counter increments (incr1 event) and decrements (decr1 event) are determined by the quadrature encoding scheme as illustrated by Figure 25-21.


Figure 25-21. Quadrature Mode Waveforms (X1, X2, and X4 mode)



Note: The x1 encoding scheme is identical to the up/down counting functionality as follows: Rising edges of input phiA increment or decrement the counter depending on the state of input phiB (direction input).

With UP_DOWN encoding (up/down rotary count mode) the counter is incremented by phiA and decremented by phiB as illustrated by Figure 25-22.


Figure 25-22. Up/down Rotary Mode



The state of phiA/phiB determines the increment/decrement according to settings of the encoding mode; the phiA/phiB is detected by counter clock. The increment/decrement occurs at the next counter clock rising edge after edges of phiA/phiB. For example, the condition of increment in X1 mode is, at the first counter clock rising edge, the phiA is low, at the second

counter clock rising edge, phiA is high (now the counter will recognize that rising edge occurs on phiA). In the meanwhile, (second counter clock rising edge), phiB is low; then the counter will do an increment. Hence to get correct quadrature encoding as Figure 25-22, the rising/falling edge of phiA/phiB must be detected in a different counter clock cycle. Figure 25-23 shows the phiA/phiB detection at a different counter clock cycle.


Figure 25-23. phiA/phiB Detection at Counter Clock



25.3.3.1 Quadrature QUAD_RANGE0 Mode

In this mode the counter range is between 0x0000 and 0xFFFF/0xFFFFFFFF (32-bit mode)

Figure 25-24. Quadrature (QUAD_RANGE0 mode) Function Diagram



Quadrature functionality in QUAD_RANGE0 mode (16-bit example) is described as a software-generated reload event starts quadrature operation. As a result, COUNTER is set to 0x8000, which is the counter midpoint (the COUNTER is set to 0x7FFF if the reload event coincides with a decrement event; the COUNTER is set to 0x8001 if the reload event coincides with an increment event). Note that a software-generated reload event is typically generated only once, when the counter is not running. All other reload/index events are hardware-generated reload events as a result of the quadrature index signal.

During quadrature operation:

- The counter value COUNTER is incremented or decremented based on the specified quadrature encoding scheme.
- On a reload/index event, CC0 is copied to CC0_BUFF, COUNTER is copied to CC0, and COUNTER is set to 0x8000. In addition, the tc and cc0_match events are generated.
- When the counter value COUNTER is 0x0000, CC0 is copied to CC0_BUFF, COUNTER (0x0000) is copied to CC0, and COUNTER is set to 0x8000. In addition, the cc0_match event is generated.
- When the counter value COUNTER is 0xFFFF, CC0 is copied to CC0_BUFF, COUNTER (0xFFFF) is copied to CC0, and COUNTER is set to 0x8000. In addition, the cc0_match event is generated.

The software interrupt handler uses the tc and cc0_match interrupt cause fields to distinguish between a reload/index event and a situation in which a minimum/maximum counter value was reached (about to wrap around). The CC0 and CC0_BUFF registers are used to determine when the interrupt causing event occurred.

Note that a counter increment/decrement can coincide with a reload/index/tc event or with a situation `cc0_match` event. Under these circumstances, the counter value is set to either `0x8000+1` (increment) or `0x8000-1` (decrement).

Figure 25-25 illustrates quadrature functionality as a function of the reload/index, `incr1`, and `decr1` events. Note that the first reload/index event copies the counter value `COUNTER` to `CC0`.

Figure 25-25. Overflow Coincides with Increment




Figure 25-26 to Figure 25-28 illustrate quadrature functionality for different event scenarios (including scenarios with coinciding events). In all scenarios, the first reload/index event is generated by software when the counter is not yet running.

Figure 25-26. Underflow Coincides with Decrement




Figure 25-27. Underflow Coincides with Increment and Overflow Coincides with Decrement





Figure 25-28. Index, Decrement (increment), and Underflow (overflow) Coincides



The QUAD_RANGE0 functionality has the advantage that the interrupts when reaching minimum/maximum values are far apart in time, so such interrupts are unlikely to get lost. This mode is preferred when interrupts are used for example to implement a higher range counter in software. Because the hardware and software counters are not updated in an atomic operation, this is not recommended for applications with real-time requirements (such as motor control).

A disadvantage of this mode is that a physical angle position of the quadrature encoder can have multiple counter representations, so software needs to do module and subtract operations to calculate the absolute angle position.

```
x = COUNTER; // read COUNTER register
if (x >= 0x8000)
    pos = (x - 0x8000) mod NR_COUNTS; // NR_COUNTS = encoders number of counts for one revolution
else
    pos = NR_COUNTS - ((0x8000 - x) mod NR_COUNTS);
```

25.3.3.2 Configuring Counter for Quadrature Mode (QUAD_RANGE0 mode)


The steps to configure the counter for quadrature mode of operation and the affected register bits are as follows.

1. Disable the counter by writing '0' to the ENABLE bit of the CTRL register.
2. Select Quadrature mode by writing '011' to the MODE[26:24] field of the CTRL register.
3. Set the required encoding mode by writing to the QUADRATURE_MODE[21:20] field of the CTRL register.
4. Set the TR_IN_SEL0 or TR_IN_SEL1 register to select the trigger that causes the event (Index and Stop).
5. Set the TR_IN_EDGE_SEL register to select the edge that causes the event (Index and Stop).
6. Set the Quadrature mode QUAD_RANGE0 by writing with the value '0' to the UP_DOWN_MODE [17:16] field in the CTRL register.
7. If required, set the interrupt upon TC or CC0_MATCH condition.
8. Enable the counter by writing '1' to the ENABLE bit of the CTRL register. A start trigger must be provided through firmware (START bit in TR_CMD register) to start the counter if the hardware start signal is not enabled.

25.3.3.3 Quadrature QUAD_RANGE0_CMP Mode

In this mode the counter range is also between 0x0000 and 0xFFFF/0xFFFFFFFF (32-bit mode). It allows a compare function during quadrature decoding using the CC0/CC0_BUFF registers and the cc0_match event.

Figure 25-29. Quadrature (QUAD_RANGE0_CMP mode) Function Diagram



Quadrature functionality in QUAD_RANGE0_CMP mode provides the same functionality as the QUAD_RANGE0 mode, except for the following differences:


- PERIOD and PERIOD_BUFF are used instead of CC0 and CC0_BUFF to capture the counter value when a reload/index event occurs or the minimum/maximum value is reached (about to wrap around).
 - When the 'capture on index' function is selected (via overloaded AUTO_RELOAD_PERIOD bit) and a reload/index event occurs, PERIOD is copied to PERIOD_BUFF, COUNTER is copied to PERIOD, and COUNTER is set to 0x8000. In addition, the tc event is generated.
 - When the 'capture on wrap-around' function is selected (via overloaded AUTO_RELOAD_PERIOD bit) and the counter value COUNTER is 0x0000 or 0xFFFF, PERIOD is copied to PERIOD_BUFF, COUNTER (0x0000 or 0xFFFF) is copied to PERIOD, and COUNTER is set to 0x8000. In addition, the tc event is generated.
- Capture0 can be used as the second index event. This event acts as a second quadrature index input. It has the same function as the reload/index0 event. Both events are OR combined.
- CC0 (CC1) and CC0_BUFF (CC1_BUFF) are used for compare functionality.
 - A cc0_match (cc1_match) event is generated when the counter changes to a state in which COUNTER equals CC0 (CC1).
 - CC0 (CC1) and CC0_BUFF (CC1_BUFF) are exchanged on a cc0_match (cc1_match) event (when specified by AUTO_RELOAD_CC bit).

Note that 'capture on index' and 'capture on wraparound' functions are separated to prevent PERIOD and PERIOD_BUFF from being overwritten before software has read them in case a wraparound is followed by multiple index events in a short time (quadrature encoder is moved back and forth around its index point). If both functions are needed, the two counters can be used synchronously (or a counter group, which includes two compare functions) in QUAD_RANGE0_CMP mode, one with 'capture on index', the other with 'capture on wraparound' behavior selected. Note also that multiple compare values can be realized by multiple synchronous counters in QUAD_RANGE0_CMP mode with different CC0 values. Except the differences mentioned above, the QUAD_RANGE0_CMP mode behaves as the QUAD_RANGE0 mode including behavior with coinciding events. [Figure 25-30](#) illustrates an example scenario with decrementing counter and additional compare functionality.

25.3.3.4 Quadrature QUAD_RANGE1_CMP Mode

In this mode the counter range is between 0x0000 and PERIOD.

Figure 25-31. Quadrature (QUAD_RANGE1_CMP) Function Diagram



Quadrature functionality in QUAD_RANGE1_CMP mode is described as a software-generated reload event starts quadrature operation. As a result, COUNTER is set to 0x0000 (the COUNTER is set to PERIOD if the reload event coincides with a decrement event; the COUNTER is set to 0x0001 if the reload event coincides with an increment event). Note that a software-generated reload event is generated only once, when the counter is not running. All other reload/index events are hardware-generated reload events as a result of the quadrature index signal.

During quadrature operation:


- The counter value COUNTER is incremented or decremented based on the specified quadrature encoding scheme.
- On a reload/index event, COUNTER is set to 0x0000. In addition, the tc event is generated.
- When COUNTER is 0x0000 and decrementing, COUNTER is set to PERIOD. In addition, the tc event and underflow event are generated.
- When COUNTER equals PERIOD and is incrementing, COUNTER is set to 0x0000. In addition, the tc event and overflow event are generated.

CC0 and CC0_BUFF are used for compare functionality.

- A cc0/1_match event is generated when the counter changes to a state in which COUNTER equals CC0/1.
- CC0/1 and CC0/1_BUFF are exchanged on a cc0/1_match event (when specified by AUTO_RELOAD_CC bit in the CTRL register).

Note that a counter increment/decrement can coincide with a reload/index/tc event. In this case, the counter value is set to either 0x0000+1 (increment) or PERIOD (decrement). The following figure illustrates quadrature functionality as a function of the reload/index, incr1 and decr1 events.

Figure 25-32. Quadrature Index, incr1 and tc (overflow) Generation



The following figures illustrate quadrature functionality for different event scenarios (including scenarios with coinciding events). In all scenarios, the first reload/index event is generated by software when the counter is not yet running.

Figure 25-33. Quadrature Index, decr1 and tc (underflow) Generation




Figure 25-34. No tc (underflow) after COUNTER = 0x0000 and no tc (overflow) after COUNTER = PERIOD




Figure 25-35. Index, Decrement, Underflow, and cc0_match Coincide





Figure 25-36. Index, Increment, Overflow, and cc0_match Coincide



The QUAD_RANGE1_CMP functionality allows the COUNTER register to reflect the current angle position of the rotary encoder; that is, no MOD or SUB calculations need to be done in software on the COUNTER value to get the current angle position. This allows a DMA copy of the current angle position from the COUNTER register; for example, to send a buffer of CAN/UART/field bus interface to synchronize with other devices. However, a disadvantage of this mode is that fast sequences of tc interrupts can occur (when encoder moves back and forth around start position). It is recommended to not use the tc interrupt in this mode.

25.3.3.5 Quadrature QUAD_RANGE1_CAPT Mode

In this mode the counter range is also between 0x0000 and PERIOD. Quadrature functionality in QUAD_RANGE1_CAPT mode provides the same functionality as the QUAD_RANGE1_CMP mode with the only difference that one or two capture functions are available instead of one or two compare functions.

Figure 25-37. Quadrature (QUAD_RANGE1_CAPT) Function Diagram




Figure 25-38 illustrates an example scenario with decrementing counter and capture functionality.

Figure 25-39. PWM Function Diagram




Table 25-28. PWM Mode Trigger Input Description

Trigger Inputs	Usage
reload	Sets the counter value and starts the counter. Behavior is dependent on UP_DOWN_MODE: ■ COUNT_UP: The counter is set to '0' and count direction is set to 'up'. ■ COUNT_DOWN: The counter is set to PERIOD and count direction is set to 'down'. ■ COUNT_UPDN1/2: The counter is set to '1' and count direction is set to 'up'. Can only be used when the counter is not running.
start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is dependent on UP_DOWN_MODE: ■ COUNT_UP: The count direction is set to 'up'. ■ COUNT_DOWN: The count direction is set to 'down'. ■ COUNT_UPDN1/2: The count direction is set to 'up'. Note that when the counter is running, the start event has no effect. Can be used when the counter is running or not running.
stop/kill	Stops the counter. Different stop/kill modes exist.
count	Count event increments/decrements the counter.
Capture0	This event acts as a switch event. When this event is active, the CC0/CC0_BUFF, CC1/CC1_BUFF, PERIOD/PERIOD_BUFF, and LINE_SEL/LINE_SEL_BUFF registers are exchanged on a tc event (when specified by AUTO_RELOAD_CC bit, AUTO_RELOAD_PERIOD bit, and AUTO_RELOAD_LINE_SEL bit in the CTRL register). A switch event requires rising, falling, or rising/falling edge event detection mode. Pass-through mode is not supported, unless the selected event is a constant '0' or '1'. When a switch event is detected and the counter is running, the event is kept pending until the next tc event. The switch event will be cleared and has no effect if it is detected when counter is not running.
Capture1 (stop1/kill1)	This event acts as a second stop/kill event. It has the same function as the stop0/kill0 event. Both events are OR combined. Note: Having two stop/kill events for a PWM allows selecting a common trigger for one stop/kill event from a PERI trigger multiplexer (allowing synchronous stop/kill operation of multiple PWMs) while selecting a dedicated ADC out-of-range trigger for the other stop/kill event (allowing real-time hardware stop of a PWM when current PWM driven signal is out of range).

Table 25-29. PWM Mode Supported Features

Supported Features	Description
Clock prescaling	Prescales the PCLK_TCPWM[x]_CLOCK[y].
One shot	Counter is stopped by hardware, after a single period of the counter: ■ COUNT_UP: on an overflow event. ■ COUNT_DOWN and COUNT_UPDN1/2: on an underflow event.
Auto reload CC	CC0/1 and CC0/1_BUFF are exchanged on a switch event and tc event (when specified by AUTO_RELOAD_CC bit in CTRL register).
Auto reload PERIOD	PERIOD and PERIOD_BUFF are exchanged on a switch event and tc event (when specified by CTRL.AUTO_RELOAD_PERIOD). Note: When COUNT_UPDN2 mode exchanges PERIOD and PERIOD_BUFF at a tc event that coincides with an overflow event, software should ensure that the PERIOD and PERIOD_BUFF values are the same.
Auto reload LINE_SEL	LINE_SEL and LINE_SEL_BUFF are exchanged on a switch event and tc event (when specified by the AUTO_RELOAD_LINE_SEL bit in the CTRL register).
Up/down modes	Specified by UP_DOWN_MODE: ■ COUNT_UP: The counter counts from 0 to PERIOD. ■ COUNT_DOWN: The counter counts from PERIOD to 0. ■ COUNT_UPDN1/2: The counter counts from 1 to PERIOD and back to 0.
Kill modes	Specified by PWM_IMM_KILL, PWM_SYNC_KILL, and PWM_STOP_ON_KILL.

Note that the PWM mode does not support dead time insertion. This functionality is supported by the separate PWM_DT mode.

Table 25-30. PWM Mode Trigger Output Description

Trigger Output	Description
cc0_match	Specified by UP_DOWN_MODE: ■ COUNT_UP and COUNT_DOWN: The counter changes to a state in which COUNTER equals CC0. ■ COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC0.
cc1_match	Specified by UP_DOWN_MODE: ■ COUNT_UP and COUNT_DOWN: The counter changes to a state in which COUNTER equals CC1. ■ COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC1.
Underflow (UN)	Counter is decrementing and changes from a state in which COUNTER equals 0. Reload event generate underflow in COUNT_DOWN, COUNT_UPDN1, or COUNT_UPDN2 mode.
Overflow (OV)	Counter is incrementing and changes from a state in which COUNTER equals PERIOD. Reload event generate overflow in COUNT_UP mode.
tc	Specified by UP_DOWN_MODE: ■ COUNT_UP: tc event is the same as the overflow event. ■ COUNT_DOWN: tc event is the same as the underflow event. ■ COUNT_UPDN1: tc event is the same as the underflow event. ■ COUNT_UPDN2: tc event is the same as the logical OR of the overflow and underflow events. Reload will generate underflow/overflow, but not the tc output trigger.

Table 25-31. PWM Mode PWM Outputs

PWM Outputs	Description
LINE_OUT	PWM line output.
LINE_COMPL_OUT	Complementary PWM line output.

Note that the cc0_match event generation in COUNT_UP and COUNT_DOWN modes are different from the generation in other functional modes or counting modes. This is to ensure that 0 percent and 100 percent duty cycles can be generated.

PWM behavior depends on the PERIOD and CC0 registers. Software can update the PERIOD_BUFF and CC0_BUFF registers, without affecting the PWM behavior. The switch/capture event can be used to switch the values of the compare and buffered compare registers. It also switches the values of the period and buffered period registers. This is the main rationale for double buffering these registers. Table 25-32 summarizes the kill mode supported in PWM mode.

Table 25-32. Kill Modes of PWM

Kill Mode	Settings	Kill-behavior
No-IMM-Async	PWM_IMM_KILL = 0; PWM_SYNC_KILL = 0; PWM_STOP_ON_KILL = 0; STOP_EDGE = NO_EDGE_DET	PWM output is suppressed: ■ At next active count clock after kill input is active. PWM output suppress is removed: ■ At next active count clock after kill input is inactive
IMM-Async	PWM_IMM_KILL = 1; PWM_SYNC_KILL = 0; PWM_STOP_ON_KILL = 0; STOP_EDGE = NO_EDGE_DET	PWM output is suppressed: ■ Immediately after kill input is active. PWM output suppress is removed: ■ At next active count clock after kill input is inactive.
No-IMM-Sync	PWM_IMM_KILL = 0; PWM_SYNC_KILL = 1; PWM_STOP_ON_KILL = 0; STOP_EDGE = RISING	PWM output is suppressed: ■ At next active count clock after kill input is active. PWM output suppress is removed: ■ At next tc event after kill input is inactive.
IMM-Sync	PWM_IMM_KILL = 1; PWM_SYNC_KILL = 1; PWM_STOP_ON_KILL = 0; STOP_EDGE = RISING	PWM output is suppressed: ■ Immediately after kill input is active. PWM output suppress is removed: ■ At next tc event after kill input is inactive.
No-IMM-Stop	PWM_IMM_KILL = 0; PWM_SYNC_KILL = Don't care; PWM_STOP_ON_KILL = 1; STOP_EDGE = RISING_EDGE/FALLING_EDGE/ BOTH_EDGES	PWM output is suppressed: ■ At next active count clock after kill input is active. PWM output suppress is removed: ■ Counter restart after kill input is inactive.
IMM-Stop	PWM_IMM_KILL = 1; PWM_SYNC_KILL = Don't care; PWM_STOP_ON_KILL = 1; STOP_EDGE = RISING_EDGE/FALLING_EDGE/ BOTH_EDGES	PWM output is suppressed: ■ Immediately after kill input is active. PWM output suppress is removed: ■ Counter restart after kill input is inactive.

25.3.4.1 PWM Mode Functionalities

Note: One-shot mode and clock prescaling are the same as in timer mode.

Up/down Count Modes

Up/down count modes control the counting direction (increment or decrement) while counter is running.

Figure 25-40 illustrates a PWM in COUNT_UP mode. The counter is initialized (to 0) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 4, an overflow and tc event are generated.
- When the counter changes to a state in which COUNTER equals 2, a cc0_match event is generated.
- PERIOD is 4, resulting in an effective repeating counter pattern of $4+1 = 5$ counter clock periods.

MODE = PWM
UP_DOWN_MODE = COUNT_UP

COUNTER starts with 0

period is PERIOD+1

reload

COUNTER

PERIOD = 4

CC0 = 2

Underflow (UN)

Overflow (OV)

Terminal Count (TC)


cc0_match

no tc event

cc0_match event on entering the COUNTER value

Notes:


- Figure 25-41. PWM in Down Counting Mode



Notes:

- 429

Figure 25-42. Up/Down Counting PWM




PWM Output Lines

PWM provides two output lines

- A PWM LINE_OUT output signal.
- A complementary PWM LINE_COMPL_OUT output signal.

The generation of the PWM output signals is a multi-step process and is illustrated as follows.

Figure 25-43. Line Generation Logic



TR_PWM_CTRL is to control the line state change per four internal events.

Dead time insertion is not supported in PWM mode, only the PWM_DT mode has this feature.

Kill input will disable both LINE_OUT and LINE_COMPL_OUT. The kill mode is specified by PWM_IMM_KILL, PWM_STOP_ON_KILL, and PWM_SYNC_KILL.

The polarity of both LINE_OUT signals can be configured in the CTRL register. The QUAD_ENCODING_MODE [0] bit sets the polarity of LINE_OUT; QUAD_ENCODING_MODE [1] bit can be used to set the polarity of LINE_COMPL_OUT. The value '1' inverts the corresponding LINE_OUT signal. Note that the polarity configuration must be done only when the counter is not enabled or running.

CC0 and PERIOD Auto Reload with Switch Event

Auto CC reload and auto PERIOD reload will provide dynamic PWM duty cycle change and count period length change. The active switch event (capture0) is required for switch activity:

- At TC event, if switch event is active and AUTO_RELOAD_CC bit of the CTRL register is set to '1', CC0/1 and CC0/1_BUFF will exchange value.
- At TC event, if switch event is active and AUTO_RELOAD_PERIOD bit of CTRL register is set to '1', PERIOD and PERIOD_BUFF will exchange value.

The following figures illustrate the update of period value in COUNT_UP mode and COUNT_DOWN mode resulting in different period times after each switch event.

MODE = PWM
UP_DOWN_MODE = COUNT_UP

reload

5

4

3

2

1

0

COUNTER

5

4

5

4

5

4

5

PERIOD_BUFF

PERIOD

Underflow (UN)

Overflow (OV)

Terminal Count (TC)

MODE = PWM
UP_DOWN_MODE = COUNT_DOWN

reload

5
4
3
2
1
0

COUNTER

5 4 5 4 5 4 5 4

PERIOD_BUFFER
PERIOD

Underflow (UN)
Overflow (OV)
Terminal Count (TC)

Left/Right/Center Align PWM with CC0/CC0_BUFF Auto Reload

- Up/down count mode must be used to generate different phase aligned PWM.
- Line state is changed per underflow/overflow/cc0_match/cc1_match internal event and can be configured in TR_PWM_CTRL register.

- Left-align:

- ❑ Write the value '0' to UP_DOWN_MODE [17:16] field in the CTRL register to set the counter direction to COUNT_UP mode
- ❑ Write the value '0' (SET) to OVERFLOW_MODE [3:2] field of the TR_PWM_CTRL register to set the LINE_OUT signal to '1' when the COUNTER reaches PERIOD value.
- ❑ Write the value '1' (CLEAR) to CC0_MATCH_MODE [1:0] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '0' when the COUNTER equals CC0 value.

- Right-align:
 - Write the value '1' to UP_DOWN_MODE [17:16] field in the CTRL register to set the counter direction to COUNT_DOWN mode
 - Write the value '0' (SET) to CC0_MATCH_MODE [1:0] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '1' when the COUNTER equals CC0 value.
 - Write the value '1' (CLEAR) to UNDERFLOW_MODE [5:4] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '0' when the COUNTER reaches '0'.
- Center-align:
 - Write the value '2' to UP_DOWN_MODE [17:16] field in the CTRL register to set the counter direction to COUNT_UPDN1 mode
 - Write the value '0' (SET) to OVERFLOW_MODE [3:2] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '1' when the COUNTER reaches PERIOD value.
 - Write the value '1' (CLEAR) to UNDERFLOW_MODE [5:4] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '0' when the COUNTER reaches '0'.
 - Write the value '2' (INVERT) to CC0_MATCH_MODE [1:0] field in the TR_PWM_CTRL register to invert the LINE_OUT signal when the COUNTER equals CC0 value.

Figure 25-46 illustrates a PWM in COUNT_UP mode with different CC0 values. The figure also illustrates how a left-aligned and right-aligned PWM can be created using the PWM in COUNT_UP mode.

Note: CC0 is changed (to CC0_BUFF, which is not depicted) on a tc event. The switch event is a constant 1.

A special case of 0 or 100 percent duty cycle is realized using the following setting (for example, left-aligned):

- 0 percent → $CC0/1 = 0$
- 100 percent → $CC0/1 > PERIOD$ (a PERIOD value of 0xFFFF/0xFFFFFFFF is restricted)

Figure 25-46. Left- and Right-aligned PWM in COUNT_UP Mode




Figure 25-47 illustrates a PWM in COUNT_DOWN mode with different CC0 values. The figure also illustrates how a right-aligned PWM can be created using the PWM in COUNT_DOWN mode.

Note: CC0 is changed (to CC0_BUFF, which is not depicted) on a tc event. The switch event is a constant 1.

A special case of 0 or 100 percent duty cycle is realized using the following setting (for example, right-aligned):

- 0 percent → $CC0/1 > \text{PERIOD}$ (a PERIOD value of $0xFFFF/0xFFFFFFFF$ is restricted)
- 100 percent → $CC0/1 = \text{PERIOD}$

Figure 25-47. Right-aligned PWM in COUNT_DOWN Mode




Figure 25-48 illustrates a PWM in COUNT_UPDN1 with different CC0 values. The figure also illustrates how a center-aligned PWM can be creating using the PWM in COUNT_UPDN1 mode.

Note: The switch event is generated by hardware trigger 1, which is a constant '1' and therefore always active at the TC condition. CC0 is changed (to CC0_BUFF, which is not depicted) on a tc event.

A special case of 0 or 100 percent duty cycle is realized using the following setting (for example, center-aligned):

- 0 percent → $CC0/1 = \text{PERIOD}$ (there is no restriction for PERIOD)
- 100 percent → $CC0/1 = 0$

cc0_match will generate at the beginning of the count period when the CC0 switches to 0. In a special case, the cc0_match generates without COUNTER equals CC0 (CC0 changes to 0 and at the same time COUNTER changes to 1 from 0).

Figure 25-48. Center-align PWM in UPDN1 Mode





Figure 25-50. Asymmetric PWM




The previous waveform illustrated functionality when the CC values are neither 0 nor PERIOD. Corner case conditions in which the CC values equal 0 or PERIOD are illustrated in the following figures.

Figure 25-51 illustrates how the COUNT_UPDN2 mode is used to generate an asymmetric PWM.

Notes:

- When up counting and CC0 value at the underflow event is 0, a cc0_match event is generated.
- When down counting and CC0 value at the overflow event is PERIOD, a cc0_match event is generated.
- A tc event is generated for both an underflow and overflow event. The tc event is used to exchange the CC0 and CC0_BUFF values.
- Software updates CC0_BUFF and PERIOD_BUFF in an interrupt handler on the tc event (and overwrites the hardware updated values from the CC0/CC0_BUFF and PERIOD/PERIOD_BUFF exchanges).


Figure 25-51. Asymmetric PWM when Compare = 0 or Period



When the counter group includes also compare function 1 with registers CC1 and CC1_BUFF, which generate cc1_match event, the compare feature behaves the same as for compare 0 function.

The cc1_match event can also be used to generate the PWM output signals. Using both cc0_match and cc1_match events for PWM output signal generation provides another way to generate an asymmetric PWM as shown in [Figure 25-52](#).


Figure 25-52. Asymmetric PWM with cc0_match and cc1_match



Such asymmetric PWM generation is more flexible than using only one compare function in the COUNT_UPDN2 mode. However, if another (third) compare function is needed, For example, to trigger an ADC, another synchronously running counter should be used.

For Advanced Motor Control, the generation of compare match 0 and compare match 1 events can be enabled or disabled individually for up and down counting (during COUNT_UPDN1/2 mode). This allows asymmetric PWM generation in COUNT_UPDN1 mode where one compare match event modifies the PWM output only while counting up and the other compare match event modifies the PWM output only while counting down. This is illustrated in the following figure, which shows one of three center-aligned PWM phases for motor control when the duty cycle value is increased from one period to the next (rising part of sign wave modulated onto the PWM signal).

Figure 25-53. Asymmetric PWM by cc0_match and cc1_match in COUNT_UPDN1 Mode



Instead of an always center-aligned PWM, the phase of the PWM signal can be temporarily shifted to allow a single shunt current measurement (current measurement at two triggers with difference calculation in software) for motor control of a permanent-magnet synchronous motor (PMSM) when the current duty cycle values of the three phases do not allow that (too small window where one PWM channel is active and two others are not). Compared to the asymmetric PWM realized with only one compare function in the COUNT_UPDN2 mode this solution uses two independent buffered compare values and generates less CPU load (less interrupts needed). This means all updates can be done for example, in the 'ADC done' interrupt service routine calculating the new duty cycle values and introducing a temporary phase shift for single shunt current measurement.

The required settings for typical, asymmetric PWM output modes are:

- Asymmetric with CC0:
 - ❑ Write the '3' to the UP_DOWN_MODE [17:16] field in the CTRL register to set the counter direction to COUNT_UPDN2 mode.
 - ❑ Write '0' (SET) to the OVERFLOW_MODE [3:2] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '1' when the COUNTER reaches PERIOD value.
 - ❑ Write '1' (CLEAR) to the UNDERFLOW_MODE [5:4] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '0' when the COUNTER reaches '0'.
 - ❑ Write '2' (INVERT) to the CC0_MATCH_MODE [1:0] field in the TR_PWM_CTRL register to invert the LINE_OUT signal when the COUNTER equals CC0 value.
- Asymmetric with CC0 and CC1 (only for counter groups with a second compare function):
 - ❑ Write '0' to the UP_DOWN_MODE [17:16] field in the CTRL register to set the counter direction to COUNT_UP mode.

- ❑ Write '0' (SET) to the CC0_MATCH_MODE [1:0] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '1' when the COUNTER equals CC0 value.
- ❑ Write '1' (CLEAR) to the CC1_MATCH_MODE [7:6] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '0' when the COUNTER equals CC1 value.
- Center-align asymmetric with CC0 and CC1 (only for counter groups with a second compare function):
 - ❑ Write '2' to the UP_DOWN_MODE [17:16] field in the CTRL register to set the counter direction to COUNT_UPDN1 mode.
 - ❑ Write '0' (SET) to the OVERFLOW_MODE [3:2] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '1' when the COUNTER reaches PERIOD value.
 - ❑ Write '1' (CLEAR) to the UNDERFLOW_MODE [5:4] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '0' when the COUNTER reaches '0'.
 - ❑ Write '0' (SET) to the CC0_MATCH_MODE [1:0] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '1' when the COUNTER equals CC0 value.
 - ❑ Write '1' (CLEAR) to the CC1_MATCH_MODE [7:6] field in the TR_PWM_CTRL register to set the LINE_OUT signal to '0' when the COUNTER equals CC1 value.

Kill Mode

PWM mode has different stop/kill modes. The mode is specified by PWM_IMM_KILL, PWM_STOP_ON_KILL, and PWM_SYNC_KILL.

- PWM_IMM_KILL is '1'. The PWM output signals DT_LINE_OUT and DT_LINE_COMPL_OUT are immediately suppressed when a kill event is detected.
- PWM_IMM_KILL is '0'. The PWM output signals DT_LINE_OUT and DT_LINE_COMPL_OUT are suppressed synchronously with the next count clock after a kill event is detected.

Figure 25-54 and Figure 25-55 illustrate both configurations.

Figure 25-54. Kill Suppresses Line Output Immediately (PWM_IMM_KILL = 1)





Figure 25-55. Kill Suppresses Line Output by Count Clock (PWM_IMM_KILL = 0)



The PWM_STOP_ON_KILL and PWM_SYNC_KILL modes specifies the functionality of kill input. The following three modes are supported:

- PWM_STOP_ON_KILL is '1' (PWM_SYNC_KILL is don't care). This mode stops the counter on a stop/kill event.
- PWM_STOP_ON_KILL is '0' and PWM_SYNC_KILL is '0'. This mode keeps the counter running, but suppresses the PWM output signals synchronously with the next count clock (active count prescaled PCLK_TCPWM[x]_CLOCK[y]) and continues to do so for the duration of the stop/kill event.
- PWM_STOP_ON_KILL is '0' and PWM_SYNC_KILL is '1'. This mode keeps the counter running, but suppresses the PWM output signals synchronously with the next count clock (active count prescaled PCLK_TCPWM[x]_CLOCK[y]) and continues to do so until the next tc event without a stop/kill event.

Figure 25-56, Figure 25-57, and Figure 25-58 illustrate these three modes.

Figure 25-56. PWM Stop on Kill




Figure 25-57. PWM Async Kill





Figure 25-58. PWM Sync Kill



Different from asynchronous kill, synchronous kill will stop/disable line output at TC event immediately after kill.

For counter groups that support Capture1 event, a second kill input function is available, similar to a stop event. Both events are OR combined and share the same kill mode settings. Having two stop/kill events for a PWM allows selecting a common trigger for one stop/kill event from a PERI trigger multiplexer (allowing synchronous stop/kill operation of multiple PWMs) while selecting a dedicated ADC out-of-range trigger for the other stop/kill event (for example, allowing real-time hardware stop of a PWM when current of a PWM driven signal is out of range).

PWM Output Selection for Stepper Motor Control

In PWM mode, the dedicated TCPWM counter groups can be used for Stepper Motor Control (SMC) including micro stepping.

Therefore, the two PWM output signals LINE_OUT and LINE_COMPL_OUT (which are usually a pair of complementary PWM signals during normal PWM operation) can be set to the following options using an output select register (register LINE_SEL, fields OUT_SEL and COMPL_OUT_SEL):

- constant low ('0')
- constant high ('1')
- PWM signal "line"
- Inverted PWM signal "line"
- 'Z' (high impedance)

The TCPWM counters supporting SMC are intended to be connected to GPIO_SMC cells in the IOSS, which support a 30-mA drive strength for directly driving stepper motors. Constant low or high values are used for full stepping while a PWM signal is used for micro stepping. With the option Z the TCPWM output signal `line_out_en` or `line_compl_out_en` can be set to '0', which causes the connected I/O cell to not drive the output. This allows it to measure the induced voltage of a currently not driven coil via ADC, which is used for zero point detection.


Note that for full stepping the TCPWM counter needs to be set to PWM (or PWM_PR) mode. However, a PWM generation is not needed; that is, the compare function resources (CC0/CC1) are free to be used for other purposes, such as trigger generation.

For each coil of a stepper motor a separate TCPWM counter is used. These counters are running synchronously (sharing the same reload/start/stop/count events and the same period). To achieve a synchronous update of the output signals across multiple TCPWM counters, the same double buffering and switching method is used as for updating CC0/CC1 or PERIOD registers. Therefore, a `LINE_SEL_BUFF` register is used, which is exchanged with the `LINE_SEL` register on a terminal count (tc) event with an actively pending switch event (when specified by `AUTO_RELOAD_LINE_SEL` bit in the CTRL register). [Figure 25-59](#) shows an example driving the two coils of a stepper motor with two TCPWM counters.

It illustrates three full steps (180°, 270°, 0°(=360°) electrical angles), followed by three micro steps using a PWM (~30°, ~60°, ~120° electrical angles). The figure is only an illustration showing one step per PWM counter PERIOD. In reality, the steps for micro stepping are significantly slower compared to the PWM period and therefore the PWM duty cycles (CC0) and/or output select registers (`LINE_SEL`) are stable over many PWM periods before they are changed by a switch event.

The time base for the steps can be realized by an additional TPCWM counter in timer mode (running on a slower counter clock or with a higher PWM period or counting overflow events of the PWM counter). This counter can generate interrupts, which trigger the CPU to update PWM duty cycles and/or output select registers and to finally generate the switch event using the `TR_CMD` register in the PERI block.

Figure 25-59. PWM Output for Stepper Motor



25.3.4.2 Configuring Counter for PWM Mode

The steps to configure the counter for the PWM mode of operation and the affected register bits are as follows.


1. Disable the counter by writing '0' to the ENABLE bit of the CTRL register.
2. Select PWM mode by writing '100' to the MODE[26:24] field of the CTRL register.
3. Set clock prescaling by writing to the DT_LINE_OUT_L[7:0] field of the DT register.
4. Set the required 16-bit period in the PERIOD register and the buffer period value in the PERIOD_BUFF register to switch values, if required.
5. Set the 16-bit compare value in the CC0/1 register and buffer compare value in the CC0/1_BUFF register to switch values, if required.
6. Set the direction of counting by writing to the UP_DOWN_MODE[17:16] field of the CTRL register to configure left-aligned, right-aligned, or center-aligned PWM.
7. Set the PWM_IMM_KILL, PWM_STOP_ON_KILL, and PWM_SYNC_KILL fields of the CTRL register as required.
8. Set the TR_IN_SEL0 or TR_IN_SEL1 register to select the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
9. Set the TR_IN_EDGE_SEL register to select the edge of the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
10. LINE_OUT and LINE_COMPL_OUT can be controlled by the TR_PWM_CTRL register to set, reset, or invert upon cc0_match, cc1_match, overflow, and underflow conditions.
11. If required, set the interrupt upon TC or CC0/1 condition.
12. Enable the counter by writing '1' to ENABLED bit of the CTRL register. A start trigger must be provided through firmware (START bit in the TR_CMD register) to start the counter if the hardware start signal is not enabled.

25.3.5 Pulse Width Modulation with Dead Time Mode

The PWM-DT functionality is the same as PWM functionality, except for the following differences:

- PWM_DT supports dead time insertion; PWM does not support dead time insertion.
- PWM_DT does not support clock prescaling; PWM supports clock prescaling.

Figure 25-60. PWM with Dead Time Functionality



Dead time insertion is a step that operates on a preliminary PWM output signal line, as illustrated in [Figure 25-43](#).

The dead time insertion for two PWM complementary output lines ranges from 0 to 255 (8 bit) or from 0 to 65535 (16 bit, only for counter groups supporting Advanced Motor Control) counter clock cycles. The setup can be done in DT_LINE_OUT_L [7:0] field (low byte) and in DT_LINE_OUT_L [15:8] field in the DT register. For the Advanced Motor Control counter, DT_LINE_OUT[15:0] is for LINE_OUT, DT_LINE_COMPL_OUT[15:0] is for LINE_COMPL_OUT.

[Figure 25-61](#) illustrates dead time insertion for different dead times and different output signal polarity settings.

Figure 25-61. Dead-time Timing





Figure 25-62 illustrates how the polarity settings and stop/kill functionality combined control the PWM output signals LINE_OUT and LINE_COMPL_OUT.

Figure 25-62. Dead Time and Kill



25.3.5.1 Configuring Counter for PWM with Dead Time Mode

The steps to configure the counter for PWM with Dead Time mode of operation and the affected register bits are as follows:

1. Disable the counter by writing '0' to the ENABLE bit of the CTRL register.
2. Select PWM with Dead Time mode by writing '101' to the MODE[26:24] field of the CTRL register.
3. Set the required dead time by writing to the DT_LINE_OUT_L [7:0] and DT_LINE_OUT_H [15:8] fields of the DT register.
4. Set the required 16-bit period in the PERIOD register and the buffer period value in the PERIOD_BUFF register to switch values, if required.
5. Set the 16-bit compare value in the CC0/1 register and buffer compare value in the CC0/1_BUFF register to switch values, if required.
6. Set the direction of counting by writing to the UP_DOWN_MODE[17:16] field of the CTRL register to configure left-aligned, right-aligned, or center-aligned PWM.
7. Set the PWM_IMM_KILL, PWM_STOP_ON_KILL, and PWM_SYNC_KILL fields of the CTRL register as required.
8. Set the TR_IN_SEL0 or TR_IN_SEL1 register to select the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
9. Set the TR_IN_EDGE_SEL register to select the edge of the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
10. LINE_OUT and LINE_COMPL_OUT can be controlled by the TR_PWM_CTRL register to set, reset, or invert upon cc0_match, cc1_match, overflow, and underflow conditions.
11. If required, set the interrupt upon TC or CC0/1 condition.
12. Enable the counter by writing '1' to ENABLED bit of the CTRL register. A start trigger must be provided through firmware (START bit in the TR_CMD register) to start the counter if the hardware start signal is not enabled.

25.3.6 Pulse Width Modulation Pseudo-Random Mode (PWM PR)

The PWM PR functionality changes the counter value using the linear feedback shift register (LFSR). This results in a pseudo random number sequence. A signal similar to a PWM signal is created by comparing the counter value COUNTER with the CC0/1 register. The generated signal has different frequency/noise characteristics than a regular PWM signal.

Table 25-33. Input Events of PWM_PR

Generated Events	Usage
Reload	Same behavior as start event. Can only be used when the counter is not running.
Start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is independent on UP_DOWN_MODE. Note that when the counter is running, the start event has no effect. Can be used when the counter is running or not running.
stop/kill	Stops the counter. Different stop/kill modes exist.
count	Not used.
Capture0	This event acts as a switch event. When this event is active, the CC0/CC0_BUFF, CC1/CC1_BUFF, PERIOD/PERIOD_BUFF, and LINE_SEL/LINE_SEL_BUFF registers are exchanged on a tc event (when specified by AUTO_RELOAD_CC0, AUTO_RELOAD_PERIOD, and AUTO_RELOAD_LINE_SEL bits in the CTRL register). A switch event requires rising, falling, or rising/falling edge event detection mode. Pass-through mode is not supported, unless the selected event is a constant '0' or '1'. When a switch event is detected and the counter is running, the event is kept pending until the next tc event. When a switch event is detected and the counter is not running, the event is cleared by hardware.
Capture1	This event acts as a second stop/kill event. It has the same function as the stop0/kill0 event. Both events are OR combined.

Note: Event detection is on the peripheral clock CLK_PERI.

Table 25-34. Basic Features of PWM_PR

Supported Features	Description
Clock prescaling	Prescales the PCLK_TCPWM[x]_CLOCKS[y].
One shot	Counter is stopped by hardware, after a single period of the counter (counter value equals period value PERIOD).
Auto reload CC	CC0/1 and CC0/1_BUFF are exchanged on a switch event AND tc event (when specified by AUTO_RELOAD_CC bit in CTRL register).
Auto reload LINE_SEL	LINE_SEL and LINE_SEL_BUFF are exchanged on a switch event and tc event (when specified by AUTO_RELOAD_LINE_SEL bit in the CTRL register).
Auto reload PERIOD	PERIOD and PERIOD_BUFF are exchanged on a switch event and tc event (when specified by AUTO_RELOAD_PERIOD bit in CTRL register).
Kill modes	Specified by PWM_SYNC_KILL, PWM_STOP_ON_KILL, and PWM_IMM_KILL.

Note: The count event is not used. As a result, the PWM_PR functionality operates on the prescaled counter clock (PCLK_TCPWM[x]_CLOCKS[y]), rather than on an active count prescaled counter clock.


Table 25-35. Trigger Outputs of PWM_PR

Trigger Output	Description
cc0_match (CC)	Counter changes from a state in which COUNTER equals CC0.
cc1_match (CC)	Counter changes from a state in which COUNTER equals CC1.
Underflow (UN)	Not used.
Overflow (OV)	Not used.
TC	Counter changes from a state in which COUNTER equals PERIOD.

Table 25-36. PWM_PR PWM Outputs

PWM Outputs	Description
LINE_OUT	PWM line output.
LINE_COMPL_OUT	Complementary PWM line output.

Figure 25-63. PWM_PR Functionality



The PWM_PR functionality is described as follows:

- The counter value COUNTER is initialized by software (to a value different from 0).
- Programmable LFSR length

The COUNTER is changed based on an LFSR polynomial (http://en.wikipedia.org/wiki/Linear_feedback_shift_register) specified by taps in the PERIOD_BUFF (overloaded) register. Different period lengths are possible by different programmed polynomials.

Examples:

Maximum length 16-bit LFSR: $x^{16} + x^{14} + x^{13} + x^{11} + 1$

With counter groups including 16-bit counters:

- $\text{temp} = \text{COUNTER}[5] \wedge \text{COUNTER}[3] \wedge \text{COUNTER}[2] \wedge \text{COUNTER}[0];$
- $\text{COUNTER} = (\text{temp} \ll 15) | (\text{COUNTER} \gg 1)$

Maximum length 8-bit LFSR: $x^8 + x^6 + x^5 + x^4 + 1$

With counter groups including 16-bit counters, realized in 8 MSBs of 16-bit LFSR:

- $\text{temp} = \text{COUNTER}[12] \wedge \text{COUNTER}[11] \wedge \text{COUNTER}[10] \wedge \text{COUNTER}[8];$
- $\text{COUNTER} = (\text{temp} \ll 15) | (\text{COUNTER} \gg 1)$

Maximum length 32bit LFSR: $x^{32} + x^{30} + x^{26} + x^{25} + 1$

With counter groups including 32-bit counters:

- $\text{temp} = \text{COUNTER}[7] \wedge \text{COUNTER}[6] \wedge \text{COUNTER}[2] \wedge \text{COUNTER}[0];$
- $\text{COUNTER} = (\text{temp} \ll 31) | (\text{COUNTER} \gg 1)$

This will result in a pseudo random number sequence for COUNTER. For example, when COUNTER is initialized to 0xace1 and a 16-bit LFSR with taps 16, 14, 13, and 11 is used, the number sequence is: 0xace1, 0x5670, 0xab38, 0x559c, 0x2ace, 0x1567, 0x8ab3, ..., 0x59c3. This sequence will repeat after 65535 counter clock cycles.

The following tables show examples for maximum length LFSRs (from http://courses.cse.tamu.edu/csce680/walker/lfsr_table.pdf) and their equivalent bit positions in the MSBs of the 16-bit COUNTER register (for GRP_CNT_WIDTH = 16). This allows possible pseudo random sequences with a period of $2^n - 1$ with n in [2, 16]. The right column shows example values for the PERIOD register to generate a tc event when starting with an initialized COUNTER value of 0xace1 (taking the “unused” LSbs into account, which result from right-shifting of the “used” MSBs).

Table 25-37. Polynomial of Maximum Length LFSR

n	n-bit LFSR Taps	Equivalent Bit Positions in 16-bit COUNTER Register	TAP Value to be programmed to PERIOD_BUFF Register	Period of Sequence	Example for PERIOD Register
4	4,3	12,13	0x3000	15	0x591e
5	5,3	11,13	0x2800	31	0x5d8f
6	6,5	10,11	0x0c00	63	0x59bb
7	7,6	9,10	0x0600	127	0x5b24
8	8,6,5,4	8,10,11,12	0x1d00	255	0x5997
9	9,5	7,11	0x0880	511	0x593f
10	10,7	6,9	0x0240	1023	0x59eb
11	11,9	5,7	0x00a0	2047	0x59dc
12	12,11,8,6	4,5,8,10	0x0530	4095	0x59cb
13	13,12,10,9	3,4,6,7	0x00d8	8191	0x59ca
14	14,14,11,9	2,3,5,7	0x00ac	16383	0x59c2
15	15,14	1,2	0x0006	32767	0x59c3
16	16,14,13,11	0,2,3,5	0x002d	65535	0x59c3

The programmable taps allow LFSRs other than maximum cycle LFSRs in [Table 25-37](#), which can result in periods other than $2^n - 1$. The following table shows some examples.

Table 25-38. Polynomial of Non-maximum Length LFSR

n-bit LFSR Taps	Bit Positions in 16-bit Counter Register	Period of Sequence	n-bit LFSR Taps	Bit Positions in 16-bit Counter Register	Period of Sequence
16,15	0,1	255	16,15,11	0,1,5	4340
16,14	0,2	126	16,15,10	0,1,6	24573
16,13	0,3	8191	16,15,9	0,1,7	32766
16,12	0,4	60	16,15,8	0,1,8	4681
16,11	0,5	16383	16,15,7	0,1,9	504
16,10	0,6	434	16,15,6	0,1,10	10235
16,9	0,7	63457	16,15,5	0,1,11	3906
16,8	0,8	24	16,15,4	0,1,12	7161
16,15,14	0,1,2	32767	16,15,3	0,1,13	3276
16,15,13	0,1,3	11811	16,15,2	0,1,14	32767
16,15,12	0,1,4	63	16,15,1	0,1,15	30

However, it is not recommended to use such non-maximum cycle LFSRs to generate a pseudo-random PWM signal, even if they result in the same cycle length as shorter maximum cycle LFSRs (for example, 255 cycles for taps 16,15). This because the values occurring in such sequences are not equally distributed over the possible value space, which results in much bigger errors compared with the desired PWM duty cycle accumulated over a full pseudo random number sequence. For example, the 8-bit LFSR with taps 8,6,5,4 (realized in 8 MSBs of 16-bit LFSR) and the 16-bit LFSR with the taps 16,15 result both in a period of 255 cycles, but a CC0 value of 0x4000 (for a desired 50 percent “accumulated duty cycle”) results in an accumulated duty cycle of:

49.8% for the 8-bit LFSR with taps 8,6,5,4 (realized in 8 MSBs of 16-bit LFSR).

□ 0.39% error

46.7% for the 16-bit LFSR with the taps 16,15

□ 6.66% error

■ Asymmetric with CC0

□ Write ‘3’ to the UP_DOWN_MODE [17:16] field in the CTRL register to set the counter direction to COUNT_UPDN2 mode.

■ When COUNTER equals CC0 (CC1), a cc0_match (cc1_match) event is generated.

■ When COUNTER equals PERIOD, a tc event is generated. Note that the LFSR produces a deterministic number sequence (given a specific counter initialization value). Therefore, it is possible to calculate the COUNTER value after a certain number of LFSR iterations n. This calculated COUNTER value can be used as PERIOD value, and the tc event will be generated after precisely n counter clocks.

■ On a tc event, the CC0/CC0_BUFF and CC1/CC1_BUFF can be conditionally exchanged under control of the capture0/switch event and the AUTO_RELOAD_CC0 field of the CTRL register (see [25.3.4.1 PWM Mode Functionalities](#)).

Note: To generate a tc event, PERIOD must be set to a value which the LFSR (register COUNTER) reaches. When realizing a shorter maximum length LFSR ($n < \text{GRP\_CNT\_WIDTH}$) within the n MSBs of a GRP_CNT_WIDTH wide LFSR, the “unused” LSbs need to be set to a value which results from right-shifting of the “used” MSBs.

■ The output line reflects: $\text{COUNTER}[14:0] < \text{CC0}[15:0]$. Note that only the lower 15 bits of COUNTER are used. As a result, for CC0 greater or equal to 0x8000, “line” is always 1. The line polarity can be inverted (as specified by QUAD_ENCODING_MODE[0] of the CTRL register). For counter groups including 32-bit counters the output line reflects: $\text{COUNTER}[30:0] < \text{CC0}[31:0]$.

■ During PWM_PR operation:

□ When COUNTER equals CC0 (CC01), a cc0_match (cc1_match) event is generated.

□ When COUNTER equals PERIOD, a tc event is generated.

- ❑ On a tc event, the CC0/CC0_BUFF, CC1/CC1_BUFF and PERIOD/PERIOD_BUFF can be conditionally ex-changed under control of the capture/switch event and the AUTO_RELOAD_CC bit and AUTO_RELOAD_PERIOD bit in the CTRL register (see 25.3.4.1 PWM Mode Functionalities).
- ❑ The output line reflects: $COUNTER[14:0] < CC0[15:0]$. Note that only the lower 15 bits of COUNTER are used for comparison, while the COUNTER itself can run up to 16- or 32-bit values. As a result, for CC greater or equal to 0x8000, “line” is always 1. The line polarity can be inverted (as specified by QUAD_ENCODING_MODE[0] of the CTRL register). For counter groups including 32-bit counters the output line reflects: $COUNTER[30:0] < CC0[31:0]$.

As mentioned, different stop/kill modes exist. The mode is specified by PWM_STOP_ON_KILL (PWM_SYNC_KILL should be '0' - asynchronous kill mode). The memory map describes the modes and the desired settings for the stop/kill event. The following two modes are supported:

- PWM_STOP_ON_KILL is '1'. This mode stops the counter on a stop/kill event.
- PWM_STOP_ON_KILL is '0'. This mode keeps the counter running, but suppresses the PWM output signals immediately and continues to do so for the duration of the stop/kill event.


Note that the LFSR produces a deterministic number sequence (given a specific counter initialization value). Therefore, it is possible to calculate the COUNTER value after a certain number of LFSR iterations, n . This calculated COUNTER value can be used as PERIOD value, and the tc event will be generated after precisely n counter clocks.

Figure 25-64 illustrates PWM_PR functionality.

Notes:

- The shaded areas represent the counter region in which the line value is '1', for a CC0 value of 0x4000. There are two areas, because only the lower 15 bits of the counter value are used.
- When CC0 is set to 0x4000, roughly one-half of the counter clocks will result in a line value of '1' or in other words a 50 percent PWM duty cycle accumulated over a full pseudo random number sequence.
- When a shorter LFSR is realized using programmable taps (for example, an 8-bit LFSR is realized in 8 MSbs of the 16-bit COUNTER register) the compare is still done on the whole 16-bit COUNTER register. That means a CC0 set to 0x4000 still results into roughly half of the counter clocks with a “line” value of '1'.

Figure 25-64. PWM_PR Output



25.3.6.1 Configuring Counter for Pseudo-Random PWM Mode

The steps to configure the counter for pseudo-random PWM mode of operation and the affected register bits are as follows.


1. Disable the counter by writing '0' to the ENABLE bit of the CTRL register.
2. Select pseudo-random PWM mode by writing '110' to the MODE[26:24] field of the CTRL register.
3. Set the PERIOD register for tc event generation and the LFSR length (16-bit or 32-bit) in the PERIOD_BUFF register to define the LFSR polynomial.
4. Set the 16-bit compare value in the CC0/1 register and the buffer compare value in the CC0/1_BUFF register to switch values.
5. Set the PWM_IMM_KILL, PWM_STOP_ON_KILL, and PWM_SYNC_KILL fields of the CTRL register as required.
6. Set the TR_IN_SEL0 or TR_IN_SEL1 register to select the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
7. Set the TR_IN_EDGE_SEL register to select the edge of the trigger that causes the event (Reload, Start, Stop, Capture0/1, and Count).
8. LINE_OUT and LINE_COMPL_OUT can be controlled by the TR_PWM_CTRL register to set, reset, or invert upon cc0_match, cc1_match, overflow, and underflow conditions.
9. If required, set the interrupt upon TC or CC0/1 condition.
10. Enable the counter by writing '1' to the ENABLED bit of the CTRL register. A start trigger must be provided through firmware (START bit in the TR_CMD register) to start the counter if the hardware start signal is not enabled.

25.3.7 Shift Register (SR)

Shift Register functionality shifts the counter value in the right direction. The capture0 input is used to generate the MSB of the next counter value. The line output signal is driven from a programmable tab of the shift register (COUNTER).

This implements a signal delay function from the capture0 input to the line output, which can be used for functions such as detecting frequency shift keying (FSK) signals. It further allows parallel-in to serial-out data conversion (by shifting-out a preloaded counter value) as well as serial-in to parallel-out data conversion including compare match functionality (another synchronous TCPWM counter in timer mode to be used as time base for software).

Figure 25-65. SR Function Diagram



25.3.7.1 SR Mode Functionality Overview

Table 25-39. Input Events of SR

External Events	Usage
Reload	Sets the counter value to '0' and starts the counter shift operation. Can only be used when the counter is not running.
Start	Starts the counter shift operation. The counter is not initialized by hardware. The current counter value is used. Note that when the counter is running, the start event has no effect. Can be used when the counter is running or not running.
Stop	Stops the counter.
Count/Shift	Shifts the counter in the right direction.
Capture0/serial-in	This event input is used as serial input to the MSB of the counter.
Capture1	Stops the counter.

Note: Event detection is on the peripheral clock, CLK_PERI.

Count event works to generate active count prescaled counter clock same as other function modes. This is how the shift works. Shifting the counter value is controlled by the count event and the counter clock, PCLK_TCPWM[x]_CLOCKS[y]. A constant '1' as well as synchronized input trigger edges can be used as count events.

Table 25-40. Basic Features of SR

Supported Features	Description
Clock prescaling	Prescales the PCLK_TCPWM[x]_CLOCKS[y].
Auto reload CC	CC0 (CC1) and CC0_BUFF (CC1_BUFF) are exchanged on a cc0_match (cc1_match) event (when specified by AUTO_RELOAD_CC0/1 bit in CTRL register).

Table 25-41. Internal Events of SR

Internal Events	Description
cc0_match	Counter changes to a state in which COUNTER equals CC0.
cc1_match	Counter changes to a state in which COUNTER equals CC1.
Underflow	Not used.
Overflow	Not used.
TC	Not used.

Table 25-42. Line Output of SR

Supported Features	Description
LINE_OUT	PWM line output. In Shift Register mode it is generated from an XOR combination of all enabled counter taps (bit position) defined by PERIOD_BUFF. For a shift register function only one tap should be used; that is, a one-hot value must be written to PERIOD_BUFF. If multiple bits in PERIOD_BUFF are set then the taps are XOR combined.
LINE_COMPL_OUT	Complementary PWM line output.

25.3.7.2 Features of SR Mode

Clock Prescaling

Same function as in TIMER mode

One Shot Mode

One-shot mode is not supported

Input Event

- A hardware- or software-generated reload event sets the counter value COUNTER to 0. Alternatively, the counter value COUNTER is initialized by software.
- A reload or start event starts the Shift Register operation.
- A stop event will stop the Shift Register operation, with no shifting even if the count event is active.
- COUNTER shift in the right direction at active count counter clock.
- Capture0 event is the serial input of MSB of COUNTER.

COUNTER Shift

The counter value COUNTER is shifted in the right direction and shifts - in the serial input (capture0 event).

- 16-bit counter groups:
 - $\text{COUNTER} = (\text{serial-in} \ll 15) | (\text{COUNTER} \gg 1)$
- 32-bit counter groups
 - $\text{COUNTER} = (\text{serial-in} \ll 31) | (\text{COUNTER} \gg 1)$


This means that depending on counter bit length the COUNTER value is right-shifted by 1 and the capture event value is set in the MSB position.

LINE_OUT Output

The output line is generated from a programmable COUNTER tap to generate a shifted version of the serial input (capture0 event). For a shift register function, only one tap should be selected via PERIOD_BUFF register; that is, a one-hot value must be written into PERIOD_BUFF. For a delay of n cycles (from capture0 event to line output), the PERIOD_BUFF bit should be set to '1', and other bits should be set to '0'. If multiple bits are set in PERIOD_BUFF then the selected taps are XOR combined.

Figure 25-66 illustrates the Shift Register functionality.


Figure 25-66. SR Shift with PERIOD_BUFF = 0x0001



CC0 and CC0_BUFF Auto Reload

On a cc0_match (cc1_match) event, the CC0/CC0_BUFF (CC1/CC1_BUFF) can be conditionally exchanged under control of the AUTO_RELOAD_CC0 field in the CTRL register, no switch event is required. Figure 25-67 illustrates the behavior of CC0 and CC0_BUFF auto reload.

Figure 25-67. SR Shift when PERIOD_BUFF = 0x4000



25.4 Design Configuration Parameters

The TCPWM block provides different types of counter groups that include design time configurable parameters. The following parameters are supported:

- Number of TCPWM counter groups
- Number of counters
- Counter width in number of bits (16-bit and 32-bit counters)
- Second capture/compare unit
- Advanced Motor Control features
 - Dead time can be 16 bits
 - LINE_OUT and LINE_COMPL_OUT have different dead time
 - cc0_match and cc1_match generation can be enabled/disabled individually for up and down counting in PWM/PWMMDT UPDN1/2 mode
 - Select function for PWM output signals (LINE_OUT and LINE_COMPL_OUT) to drive '0', '1', PWM, inverted PWM, or 'Z' (high impedance) including buffer register and synchronous update across counters via switch event
- Number of input triggers per counter only routed to one counter (one-to-one input triggers)
- Number of input triggers routed to all counters (general-purpose input triggers)

25.5 Recovery

TCPWM can be recovered with any Active reset event, such as:

- Power-on reset (POR)
- External reset (XRES_L)
- Watchdog timer reset (MCWDT and WDT)
- Brownout detection reset
- Over-voltage and over-current detection reset

There is no unexpected state in which the TCPWM can enter.

25.6 Initialize

The initial state of TCPWM pins is Hi-Z. Some registers are reset on an Active reset; some of the MMIO registers are retained in DeepSleep. None of the registers are retained through Hibernate or other low-power modes. An Active reset will reset the pin state back to Hi-Z.

25.7 Pin Status

When TCPWM is unused, the status for TCPWM pins will be Hi-Z. To disable TCPWM, make sure the ENABLED and PWM_DISABLE_MODE bits in the CTRL register are set to '0'.

25.8 TCPWM Registers

Table 25-43. List of TCPWM Registers

Register	Name	Description
TCPWMx_GRPy_CNTz	Prefix of dedicated counter z register in counter group y for TCPWM instance x	More details are available in the <i>Traveo II Body Controller High Registers TRM</i>
TCPWMx_GRPy_CNTz_CTRL	Counter control register	Selects the counter mode and debug mode, and enables the counter
TCPWMx_GRPy_CNTz_STATUS	Counter status register	Reads the direction of counting and dead time duration, and indicates the actual level of trigger input and trigger outputs signals; checks if the counter is running
TCPWMx_GRPy_CNTz_COUNTER	Counter count register	Contains the 16- or 32-bit counter value
TCPWMx_GRPy_CNTz_CC0	Counter compare/capture 0 register	Captures the counter value 0 or compares the value with counter value 0
TCPWMx_GRPy_CNTz_CC0_BUFF	Counter buffered compare/capture 0 register	Buffer register for counter CC0 register
TCPWMx_GRPy_CNTz_CC1	Counter compare/capture 1 register	Captures the counter value 1 or compares the value with counter value 1
TCPWMx_GRPy_CNTz_CC1_BUFF	Counter buffered compare/capture 1 register	Buffer register for counter CC1 register
TCPWMx_GRPy_CNTz_PERIOD	Counter period register	Contains upper value of the counter
TCPWMx_GRPy_CNTz_PERIOD_BUFF	Counter buffered period register	Buffer register for counter period register
TCPWMx_GRPy_CNTz_LINE_SEL	Counter line selection register	Selects the source for the LINE_OUT and LINE_COMPL_OUT output signals
TCPWMx_GRPy_CNTz_LINE_SEL_BUFF	Counter buffered line selection register	Buffer register for the LINE_SEL register
TCPWMx_GRPy_CNTz_DT	Counter PWM dead time register	Configuration of PWM dead time affecting LINE_OUT and LINE_COMPL_OUT signals
TCPWMx_GRPy_CNTz_DT	Counter trigger command register	Enables software-controlled operation for this counter. It includes the software trigger for CAPTURE0, CAPTURE1, RELOAD, START, and STOP
TCPWMx_GRPy_CNTz_TR_IN_SEL0	Counter input trigger selection register 0	Selects triggers for specific counter events: CAPTURE0, COUNT, RELOAD, or STOP event
TCPWMx_GRPy_CNTz_TR_IN_SEL1	Counter input trigger selection register 1	Selects triggers for specific counter events: CAPTURE1 or START event
TCPWMx_GRPy_CNTz_TR_IN_EDGE_SEL	Counter input trigger edge selection register	Determines edge detection for specific counter triggers. Events will only take effect on enabled counters
TCPWMx_GRPy_CNTz_TR_PWM_CTRL	Counter trigger PWM control register	Controls counter LINE_OUT, DT_LINE_OUT, and DT_LINE_COMPL_OUT output signals
TCPWMx_GRPy_CNTz_TR_OUT_SEL	Counter output trigger selection register	Selects internal events for output trigger generation
TCPWMx_GRPy_CNTz_INTR	Interrupt request register	Sets the register bit when TC or CC0/1 condition is detected
TCPWMx_GRPy_CNTz_INTR_SET	Interrupt set request register	Sets the corresponding bits in interrupt request register
TCPWMx_GRPy_CNTz_INTR_MASK	Interrupt mask register	Mask for interrupt request register
TCPWMx_GRPy_CNTz_INTR_MASKED	Interrupt masked request register	Bitwise AND of interrupt request and mask registers

Note: In TCPWMx_GRPy_CNTz, 'x' signifies TCPWM instance number, 'y' is the group number and 'z' is the counter in the respective TCWPM group.

Note that overwriting the same value on each register has different effects and they are explained in the register map by the software access attributes. TCPWM registers have the following access restrictions:

- All status registers are not software-writable.
- TR_CMD is set in software and cleared in hardware.
- INTR is cleared in software and set in hardware (by writing '1' to INTR_SET).
- Read INTR_SET will return the value of INTR.
- Other registers are normal and can be overwritten with the same value.

26. Local Interconnect Network (LIN)



The LIN unit of Traveo II supports the serial interface protocols LIN and UART. It supports the autonomous transfer of the LIN frame to reduce CPU processing.

26.1 Features

26.1.1 LIN


- LIN protocol support in hardware according to ISO 17987 standard
- Master and slave functionality
- Master node
 - Autonomous header transmission and autonomous response transmission and reception
- Slave node
 - Autonomous header reception and autonomous response transmission and reception
- Message buffer for PID, data, and checksum fields
- Classic and enhanced checksum
- Timeout detection
- Error detection
- Test modes including hardware error injection
- Baud rate detection
- 16x bit time oversampling

26.1.2 UART

- Programmable 5/6/7/8-bit data fields
- Programmable number of STOP bits: $\frac{1}{2}$, 1, $1\frac{1}{2}$, or 2 bits
- Optional parity functionality with odd and even parity

26.2 Block Diagram

Figure 26-1. LIN Block Diagram



26.2.1 Internal Bus Interface

The LIN unit registers are connected via an AHB-Lite IF to the peripheral bus.

26.2.2 Test Registers

The test error injection and different LIN signal tests are controlled within this unit.

26.2.3 LIN Channel

The LIN channels are part of one common LIN unit. Each channel has its own control and status registers and its interrupts are routed to the external interrupt controller.

26.3 Clocking


Each LIN channel has its own LIN channel input clock, PCLK_LINx_CLOCK_CH_ENy. This LIN internal channel clock is derived from the peripheral interconnect (PERI) clock, CLK_PERI, and the peripheral clock divider settings in the clock tree, PCLK_LINx_CLOCK_CH_ENy. The register clock is derived from one of the group clocks (CLK_GRx) according to the clock tree description, to have a fast register access.

26.3.1 Baud Rate and Sample Point

One LIN bit length corresponds to 16 PCLK_LINx_CLOCK_CH_ENy cycles; that is, 16 oversample counters are executed. The LIN receiver starts counting after the detection of the falling edges on the synchronized rx_synced signal to identify START bits. A bit value is sampled when the oversample counter changes from '7' to '8'.

The LIN receiver can operate (detect and sample) on the internally rx_synced signal directly, or it can operate on a filtered version of this signal by setting the LINx_CHy_CTL0.FILTER_EN bit. The filter consists of a three-input median/majority filter that effectively performs a majority vote on a window of three consecutively rx_synced samples. For more details, see [Noise Filter on page 467](#).

Figure 26-2. LIN Bit Timing Diagram



The baud rate can be configured for each channel individually, which is derived from the PERI clock. As there is the fixed signal oversampling factor of 16 in the LIN channel, for the target baud rate the clock divider for the dedicated PCLK_LINx_CLOCK_CH_ENy in the PERI component must be calculated as follows. Thereby the baud rate calculation considered for the master resp. the slave with fixed clock and for the slave with required baud rate adjustment due to inaccurate system clock. For details about the possible the clock divider settings, see the [Clocking System chapter on page 182](#).

Depending on whether a fractional clock divider or an integer clock divider is applied for the LIN module input clock, check if the maximum permitted relative tolerance of the nominal LIN bit time according to the LIN ISO specification is exceeded or not. For example, the maximum master bit rate deviation from nominal bitrate (FTOL_RES_MASTER) is ± 0.5 percent.

CLK_PERI internal peripheral clock

PCLK_LINx_CLOCK_CH_ENy dedicated internal LIN channel clock derived from the internal peripheral clock

Tbit $16 \times T_{\text{PCLK\_LINx\_CLOCK\_CH\_ENy}}$

f_{bit} LIN baud rate

f_{CLK_PERI} peripheral interconnect (PERI) clock frequency

CLK_DIV clock divider for dedicated LIN channel

26.3.1.1 Baud Rate Calculation for LIN Master and Fixed LIN Slave Clock

$$CLK\_DIV = \frac{f_{CLK\_PERI}}{f_{PCLK\_LINx\_CLOCK\_CH\_ENy}} = \frac{f_{CLK\_PERI}}{16 \cdot f_{bit}} \quad \text{Equation 26-1}$$

26.3.1.2 Baud Rate Calculation Adjusted LIN Slave Clock

$$CLK\_DIV = \frac{f_{CLK\_PERI}}{f_{PCLK\_LINx\_CLOCK\_CH\_ENy}} \cdot SyncByteCorrection$$

$$CLK\_DIV = \frac{f_{CLK\_PERI}}{16 \cdot f_{bit}} \cdot \frac{"LIN\_CH\_TX\_RX\_STATUS.SYNC\_COUNTER" \text{ value}}{128} \quad \text{Equation 26-2}$$

26.3.1.3 Example: Master

f_{bit,nom} nominal bit rate 20 kBaud = 20 kHz

f_{bit,real} real bit rate

f_{CLK_PERI} 100 MHz

integer clock divider in use

$$CLK\_DIV = \frac{f_{CLK\_PERI}}{16 \cdot f_{bit,nom}} = \frac{100MHz}{16 \cdot 20kHz} = 312.5 \quad \text{Equation 26-3}$$

As there is no integer result and an integer clock divider is in use, the relative bit time tolerance is checked with $CLK\_DIV = 312$.

$$f_{bit,real} = \frac{f_{CLK\_PERI}}{16 \bullet CLK\_DIV} = \frac{100MHz}{16 \bullet 312} \approx 20.032kHz$$

The resulting relative bit time tolerance is +0.16% and within the $\pm 0.5\%$ of $FTOL\_RES\_MASTER$.


26.4 LIN Message Frame Format

A LIN message frame consists of two main elements, header and response (see Figure 26-3).

- A frame header, transmitted only by the master node, consists of a break field, followed by a synchronization (SYNC) field and a protected identifier (PID) field.
- A frame response consisting of a maximum of eight data fields and followed by a checksum field can be transmitted by the master node or by a slave node.

With exception of the LIN break field the LIN frame structure is based on byte fields, each with a START bit and a STOP bit. Due to frame support in the LIN module registers are provided for the PID field, data fields, and checksum field. The LIN break and SYNC field are processed in the LIN module and thus there is no message buffer required for the transmission as LIN master. The handling as master or slave is controlled implicitly by commands instead of a dedicated master or slave control bit.

The following sections describe the LIN protocol support by hardware.



26.4.1 Break and Synchronization Fields

The break field is generated by the master node with minimum 13-bit periods (on the master clock), whereas a slave node has to detect a break field after 11-bit periods (on the slave clock). For the master and slave, the break length must be configured in `LINx_CHy_CTL0.BREAK_WAKEUP_LENGTH` and the break delimiter length in `LINx_CHy_CTL0.BREAK_DELIMITER_LENGTH`.

The SYNC field with the signal pattern 0x55 is used to synchronize the slave clocks to the master clock. When the LIN module is configured as master (`LINx_CHy_CMD.TX_HEADER = 1` and `LINx_CHy_CMD.RX_HEADER = 0`), the SYNC field is generated autonomously. When the LIN channel is

configured as slave node (`LINx_CHy_CMD.TX_HEADER = 0` and `LINx_CHy_CMD.RX_HEADER = 1`), the detected baud rate is mirrored in the implicitly by the `LINx_CHy_TX_RX_STATUS.SYNC_COUNTER`.

Notes:

- Before the wakeup transmission start, the bus level of the LIN signal input `LINx_CHy_TX_RX_STATUS.RX_IN` must be on recessive level (logical 1). If the bus level is dominant level (logical 0), then the LIN module waits until the bus level is changing to the recessive level.
- The received signal pattern of the synchronization field is verified. When it is invalid, the error flag `LINx_CHy_INTR.RX_HEADER_SYNC_ERROR` is activated.

Baud rate adjustment

The baud rate detection is done by the 128 bit PCLK_LINx_CLOCK_CH_ENy synchronization field counter (see [Baud Rate and Sample Point on page 456](#)). The slave measures the duration of the 8-bit field, which starts from the falling edge of SYNC field START bit and stops counting with falling edge of the seventh data bit. One bit period corresponds to 16 PCLK_LINx_CLOCK_CH_ENy cycles

and 8-bit periods are finally 128 PCLK_LINx_CLOCK_CH_ENy cycles.

The following table lists the synchronization cases with the resulting SYNC byte correction factor for the new clock divider calculation. The clock divider calculation for the synchronized slave is shown in [Baud Rate and Sample Point on page 456](#).

Table 26-1. Baud Rate Adjustment Correction Factor

Clock Ratio: master to slave	Slave Value LINx_CHy_TX_RX_STATUS.SYNC_COUNTER	Counter Action	SYNC Byte Correction Factor for LIN ch. Clock Divider
$f_{\text{master}} = f_{\text{slave}}$	$x = 128$	No change	$(128 / 128) = 1$
$f_{\text{master}} < f_{\text{slave}}$	$x > 128$	Decrease the slave clock (increase the LIN ch. clock divider)	$(x / 128) > 1$
$f_{\text{master}} > f_{\text{slave}}$	$x < 128$	Increase the slave clock (decrease the LIN ch. clock divider)	$(x / 128) < 1$

26.4.2 PID Field

The 8-bit PID field consists of a 6-bit frame identifier and a 2-bit parity over the frame identifier, for which the LINx_CHy_PID_CHECKSUM register is provided exclusively.

- Master operation: Before the transmission start of the message frame the PID field will be written.
- Slave operation: After the reception of the STOP bit from the PID field the LINx_CHy_PID_CHECKSUM register is updated. The confirmation of a finished and valid LIN header reception is flagged by LINx_CHy_INTR.RX_HEADER_DONE.
- The parity of the received PID field is verified. In case of verification failure, the error flag LINx_CHy_INTR.RX_HEADER_PARITY_ERROR is activated.

26.4.3 Data Fields

As well the master as a slave can transmit a response field including maximum eight data fields. As message buffer for the data fields LINx_CHy_DATA0 and LINx_CHy_DATA1 are provided. The target number of data fields is processed in the register bit field LINx_CHy_CTL1.DATA_NR. The status of transferred numbers of data bytes including the checksum field within a response is given in LINx_CHy_STATUS.DATA_IDX. Additionally the status of an ongoing frame transfer is represented, when LINx_CHy_STATUS.HEADER_RESPONSE is '1'. All these registers are used for response transmission and response reception.

The response transfer can be aborted by disabling the LIN channel (clear LINx_CHy_CTL0.ENABLED to '0').

26.4.3.1 Response Transmission (LINx_CHy_CMD.TX_RESPONSE)

Before the transmission response is started by the command LINx_CHy_CMD.TX_RESPONSE, it must be ensured, that the data is written into the message buffer and the data length is stored.

Master operation: The response transmission can be prepared either after the reception of the PID or before the LIN frame transmission, to reduce the CPU load.

Slave operation: no additional note.

26.4.3.2 Response Reception (LINx_CHy_CMD.RX_RESPONSE)

The response reception is enabled by the command LINx_CHy_CMD.RX_RESPONSE. It is strongly recommended, to enable it before each LIN frame start. Otherwise there is the risk of losing the response data, when the response reception is enabled after another node has already started to transmit the response. The configuration of data response length in LINx_CHy_CTL1.DATA_NR and the checksum type selection must be configured at latest before the reception of the STOP bit in the first data byte.

Master operation: To reduce the CPU load, the data length can be stored before the LIN frame, as it is already known to the master.

Slave operation: The correct data length can be stored after the reception of the PID field. Therefore it is recommended, to configure the maximum data length for the response reception before the LIN frame transmission, to avoid timing constraints in the PID processing.

Notes: When the LIN response transmission and reception are active, both the transmission and reception error flags

occur simultaneously. The transmitted data fields in the LINx_CHy_DATA0/1 registers are not overwritten by the received data fields.

26.4.4 Checksum Field

The checksum field provides an integrity check over the response data fields and optionally over the header PID field, which is controlled by the LINx_CHy_CTL1.CHECKSUM_ENHANCED register field. The checksum field is supported through a message buffer register in LINx_CHy_PID_CHECKSUM.CHECKSUM.

26.4.4.1 Response Transmission (LINx_CHy_CMD.TX_RESPONSE)

For the completion of the response transmission the checksum value is calculated by hardware and is transmitted automatically after the last data field. For an invalid checksum read back the LINx_CHy_INTR.TX_RESPONSE_BIT_ERROR is set. The checksum type selection can be done already before the LIN frame start.

26.4.4.2 Response Reception (LINx_CHy_CMD.RX_RESPONSE)

When receiving, the checksum over the received PID field and data fields is calculated to verify the received checksum field. In case of verification failure a LINx_CHy_STATUS.RX_RESPONSE_CHECKSUM_ERROR is activated. The checksum type should be selected before the reception of the first data byte STOP bit reception.

26.5 Timeout Operation

For development purposes a timeout functionality is provided to determine an incomplete LIN message frame operation. The timeout detection mode can be selected between a complete frame (header and response), header, and response transfer by the LINx_CHy_CTL1.FRAME_TIMEOUT_SEL field and the timeout value is specified by the LINx_CHy_CTL1.FRAME_TIMEOUT field in number of bit periods. The LINx_CHy_INTR.TIMEOUT flag is set, when either the timeout detected or the stop condition is reached.

Note: An ongoing frame transfer is not aborted due to a timeout.

Table 26-2. Timeout Selection

FRAME_TIMEOUT_SEL Bit Field Value	Timeout Selection	Timer Start	Timer Stop
0	Timeout disabled	None	None
1	Frame mode	Falling edge of START bit in break field	Checksum field STOP bit OR timeout
2	Frame header mode	Falling edge of START bit in break field	PID field STOP bit OR timeout
3	Frame response mode	End of STOP bit	Checksum field STOP bit OR timeout

26.6 Wakeup

When a LIN cluster is in sleep state, a wakeup signal can initiate a transfer to operational state. Both the dominant wake up signal generation and detection are supported in hardware.

26.6.1 Wakeup Signal Transmission

Before the generation of the dominant wake up signal, its dominant pulse length should be defined in the register field LINx_CHy_CTL0.BREAK_WAKEUP_LENGTH in bit periods, which corresponds to the specified wake up pulse length range according to the LIN specification. The transmission starts by setting LINx_CHy_CMD.TX_WAKEUP. The flag LINx_CHy_INTR.TX_WAKEUP_DONE confirms the completed dominant wakeup pulse, except when the received signal is different to the generated one, then the error is LINx_CHy_INTR.TX_BIT_ERROR is set.

Note: Before the wakeup transmission start, the bus level of the LIN signal input LINx_CHy_TX_RX_STATUS.RX_IN must be on recessive level (logical 1). If the bus level is dominant level (logical 0), then the LIN module waits until the bus level is changing to the recessive level.

26.6.2 Wakeup Signal Reception

To activate the wakeup reception, the commands LINx_CHy_CMD.TX_HEADER and LINx_CHy_CMD.RX_HEADER should be disabled.

Typically, external transceivers support remote wakeup detection. The generated 'low' level signal can be detected by polling of the receiver input LINx_CHy_TX_RX_STATUS.RX_IN within the LIN unit. Other opportunities such as an input capture detection of the falling edge need to be checked for the dedicated port pin.

The coding information of the TX and RX transceiver pins about the wake up source can be captured directly with the internal LIN module signals

LINx_CHy_TX_RX_STATUS.RX_IN and LINx_CHy_TX_RX_STATUS.TX_IN. For this case the LIN_TX GPIO input function must be enabled (see the [I/O System chapter on page 246](#)).

When the external LIN transceiver is in operational mode, the dominant wake up pulse is passed on. To detect it, the minimum expected pulse length must be configured in the form of bit periods in the register bit field LINx_CHy_CTL0.BREAK_WAKEUP_LENGTH. When the rising edge of the dominant pulse is detected, then the flag LINx_CHy_INTR.RX_BREAK_WAKEUP_DONE is set.

26.6.3 Wake up in Low Power Mode

The LIN unit cannot detect a wakeup condition, when the device is DeepSleep or Hibernate power mode. To support a CPU wakeup, refer to the interrupt on falling edge support for the LIN_RX port pin of the LIN channel.

26.7 External Transceiver Control

Discrete LIN transceiver devices may consume a significant amount of power when enabled. Fortunately, most transceivers support the Sleep power mode in which power consumption is reduced. To this end, most transceivers have an enable “en” input signal to control the power mode.

Each LIN channel has an “en” line that is used to control the transceiver enable input signal. Before a message transfer, the en line should be activated, and after the message transfer the en line can be deactivated. The en line can be controlled by either software or hardware.

- Software control requires setting LINx_CHy_TX_RX_STATUS.EN_OUT to '1' before a message transfer and clearing LINx_CHy_TX_RX_STATUS.EN_OUT to '0' after a message transfer.
- Hardware control ensures setting LINx_CHy_TX_RX_STATUS.EN_OUT to '1' before a message transfer and clearing

LINx_CHy_TX_RX_STATUS.EN_OUT to '0' after message transfer.

The LINx_CHy_CTL0.AUTO_EN field enables the hardware control of the en signal line.

26.8 Test Modes

26.8.1 Interrupt Test

To test the internal interrupt signals line within the LIN module regarding functionality, an interrupt set function is provided by the LINx_CHy_INTR_SET register.

26.8.2 Loop-back Mode

A self-test circuit allows the channels to be connected to each other, to test the LIN functionality without an external transceiver or without affecting an operational LIN cluster by enabling the register bit LINx_TEST_CTL.ENABLED. The LIN operation configuration of the two selected channels, to operate as LIN master and LIN slave, is done as usual.

Following channel loop back connections are permitted:

- Channel [0, CH_NR-2], which is identified by the LINx_TEST_CTL.CH_IDX register field and
- the last channel [CH_NR-1].

Note: CH_NR refers to the maximum LIN channel number.

26.8.2.1 Partial Disconnect Mode

In this mode both channels to be tested the loop back is done via the port pins. In this case the GPIO input function of TX port pin from channel [i] has to be enabled (see the [I/O System chapter on page 246](#)).

26.8.2.2 Full Disconnect Mode

In this mode the LIN channels under test are routed with each other completely inside the LIN unit (see [Figure 26-6](#)). There is no connection to existing port pins and thereby no impact to the LIN bus.

Figure 26-4. Functional Mode




Figure 26-5. Partial Disconnect Mode





Figure 26-6. Full Disconnect Mode



26.8.3 Error Injection Mode

For test purposes, hardware injected transmitter errors can be generated, which result in the activation of the corresponding error flag on the reception line.

The error injection type is selected by the LINx_ERROR_CTL register. The LINx_ERROR_CTL.CH_IDX field specifies the channel to which the errors are applied. Table 26-3 shows the error injection types.

Table 26-3. Error Injection Support in LIN/UART Unit

Error Injection	Error Injection Description	Mode support	
		LIN	UART
TX_SYNC_ERROR	The transmitted synchronization field is changed from 0x55 to 0x00.	Yes	No
TX_SYNC_STOP_ERROR	The synchronization field STOP bits are inverted to '0'.	Yes	No
TX_PARITY_ERROR	LIN: The highest parity bit of the PID field is inverted. UART: parity bit in data field is inverted.	Yes	Yes
TX_PID_STOP_ERROR	The PID field STOP bits are inverted to '0'.	Yes	No
TX_DATA_STOP_ERROR	The data field STOP bits are inverted to '0'.	Yes	Yes
TX_CHECKSUM_ERROR	The checksum field is inverted.	Yes	No
TX_CHECKSUM_STOP_ERROR	The checksum field STOP bits are inverted to '0'.	Yes	No

26.9 Operation

26.9.1 LIN Operation

26.9.1.1 LIN Message Transfer

The LIN protocol supports three types of message transfers:

- Master response: The master node transmits the header and transmits the response. This type can be used to control slave nodes.
- Slave response: The master node transmits the header. A slave node transmits the response and the master node receives the response. This type can be used to observe slave node status.
- Slave to slave: The master node transmits the header. A slave node transmits the response and another slave receives the response.

To support these different message types, the handling of the LIN master or LIN slave operation mode is implicitly done by command sequences.

- LINx_CHy_CMD.TX_HEADER: This command is used exclusively by the master node to transmit a complete header such as, LIN break, SYNC field, PID field.
- LINx_CHy_CMD.RX_HEADER: This command is used exclusively by a slave node to receive a header. After a slave node receives the header, LINx_CHy_INTR.RX_HEADER_DONE is activated and slave node application may use the received PID field to decide to either:
 - Continue with receipt of a response (LINx_CHy_CMD.RX_RESPONSE command).
 - Continue with transmission of a response (LINx_CHy_CMD.TX_RESPONSE command).
 - Ignore the incoming response by disabling the channel and re-enabling for the next frame
- LINx_CHy_CMD.TX_RESPONSE: This command is used by the master node or a slave node to transmit a response; that is, the hardware sends the data field and the autonomously generated checksum.
- LINx_CHy_CMD.RX_RESPONSE: This command is used by the master node or a slave node to receive a response; that is, the hardware receives the data field in one buffer and verifies the checksum.

In Table 26-4 and Table 26-5 the command sequences for master and slave for the different message types are shown.

Table 26-4. LIN Master Command Sequences

Message type	Command Sequence in register CMDi ^a			
	CMDi.TX_HEADER	CMDi.RX_HEADER	CMDi.TX_RESPONSE	CMDi.RX_RESPONSE
Master Response	1	0	1	0
Slave Response	1	0	0	1
Slave-to-Slave Response	1	0	0	0

a. Command sequence can be done before frame start.

Table 26-5. LIN Slave Command Sequences

Message types	Command Sequence in register CMDi ^a			
	CMDi.TX_HEADER	CMDi.RX_HEADER	CMDi.TX_RESPONSE	CMDi.RX_RESPONSE
Master Response	0	1	0	1
Slave Response	0	1	1	1 ^b
Slave-to-Slave Response (transmitting node)	0	1	1	1
Slave-to-Slave Response (receiving node)	0	1	0	1
Ignore Response	0	1	0	0

a. LINx_CHy_CMD.RX_HEADER and LINx_CHy_CMD.RX_RESPONSE are enabled before break detection to avoid break loss and loss of data bytes in response. Disabling of LINx_CHy_CMD.RX_RESPONSE after PID reception is permitted.

b. When both LINx_CHy_CMD.TX_RESPONSE and LINx_CHy_CMD.RX_RESPONSE is set, then a bus collision can be detected by LINx_CHy_INTR.RX_RESPONSE_DONE.

Master

The master node needs to enable one interrupt cause (LINx_CHy_INTR.TX_HEADER_DONE, LINx_CHy_INTR.TX_RESPONSE_DONE, LINx_CHy_INTR.RX_RESPONSE_DONE) and only enters the associated interrupt handler once.

Slave

The slave nodes will always set both LINx_CHy_CMD.RX_HEADER and LINx_CHy_CMD.RX_RESPONSE commands to '1'. The received header PID field will specify if a slave node:

- Has to receive a response.
- Has to transmit a response.
- Abort the transfer and ignore the response.

By setting LINx_CHy_CMD.RX_HEADER and LINx_CHy_CMD.RX_RESPONSE simultaneously, the slave node anticipates response reception, to avoid loss of data bytes in the response.

Master and slave

When a message transfer is successful, the commands are cleared to '0' and must be enabled again for the next transfer. On a detected error, the transmission commands are cleared to '0', but the reception commands are not. This behavior is essential to support break-while-receive functionality on a slave node.

Both the response commands LINx_CHy_CMD.TX_RESPONSE and LINx_CHy_CMD.RX_RESPONSE can be enabled in parallel, a command order is processed in following priority:

- Highest priority: LINx_CHy_CMD.TX_RESPONSE command.
- Middle priority: LINx_CHy_CMD.RX_RESPONSE command.
- Lowest priority: No response as indicated by the absence of BOTH the LINx_CHy_CMD.TX_RESPONSE and LINx_CHy_CMD.RX_RESPONSE commands.

26.9.1.2 LIN Software Flow Chart

This section shows software flow charts for the LIN master and slave operation.

Figure 26-7. LIN Master Software Flow Chart





Figure 26-8. LIN Slave Software Flow Chart



26.9.2 UART Operation

The LIN unit supports limited UART functionality:

- Programmable 5/6/7/8-bit data fields (LINx_CHy_CTL0.BREAK_DELIMITER_LENGTH[1:0]).
- Programmable number of STOP bits: ½, 1, 1½, or 2 bits (LINx_CHy_CTL0.STOP_BITS[1:0]).
- Optional parity functionality (LINx_CHy_CTL0.PARITY_EN) with odd and even parity (LINx_CHy_CTL0.PARITY).

The UART operation mode is enabled, when LINx_CHy_CTL0.MODE is set to '1'.

A single UART frame consists of a single START bit, a data field (transferred least significant bit first), an optional parity bit, and a programmable number of STOP bits.

26.9.2.1 Transmission

The TX_HEADER command is used to transmit a single data field as specified by LINx_CHy_DATA0.DATA1[7:0]. The LINx_CHy_INTR.TX_HEADER_DONE interrupt cause is activated, when the transfer is completed. The LINx_CHy_INTR.TX_HEADER_BIT_ERROR interrupt cause is activated when a bit error is detected. If the parity function is enabled, then hardware executes the parity bit calculation.

26.9.2.2 Reception

The RX_HEADER command is used to receive a single data field in LINx_CHy_DATA0.DATA1[7:0]. The LINx_CHy_INTR.RX_HEADER_DONE interrupt cause is activated when the transfer is completed. The LINx_CHy_INTR.RX_HEADER_FRAME_ERROR interrupt cause is activated when a frame error is detected (unexpected START or STOP bit value). The LINx_CHy_INTR.RX_HEADER_PARITY_ERROR interrupt cause is activated, when a parity error is detected in case of enabled parity function.

When the noise detection is enabled and noise is seen, the LINx_CHy_INTR.RX_NOISE_DETECT error is set.

26.9.2.3 Extended Features

The UART operation mode supports following features, which are described in the previous sections:

- LINx_CHy_CTL0.AUTO_EN
- LINx_CHy_CTL0.BIT_ERROR_IGNORE
- LINx_CHy_CTL0.FILTER_EN

26.9.2.4 Multiple Transfer

To transfer multiple UART frames, multiple TX/RX_HEADER commands are required; that is, the UART operation mode data length counter LINx_CHy_CTL1.DATA_NR is not supported.

26.10 Noise Filter


The LIN receiver operates on the synchronized rx_synced input signal, as shown in Figure 26-9.

- When LINx_CHy_CTL0.FILTER_EN is '0', the receiver operates on rx_synced directly.
- When LINx_CHy_CTL0.FILTER_EN is '1', the receiver operates on the majority of the last three rx_synced signal values based on the internal module clock PCLK_LINx_CLOCK_CH_ENy. This filter suppresses noise on the rx_in input. Note that the filter adds a delay of one cycle to the receiver. Figure 26-9 shows the block diagram of the noise filter and Figure 26-10 shows the noise filtering timing behavior including the sample point position.

26.10.1 Example

When a '0', '1', '0' sequence is synchronized, the '1' is effectively filtered out due to majority decision for '0'.

Figure 26-9. LIN Signal Line Synchronization Block Diagram



Even when the median filter effectively eliminates the rx_in noise, it is of interest to be notified of this noise, as the noise can be an indication of a malfunctioning LIN cluster. Therefore, the receiver verifies the rx_in signal by investigating the last three rx_synced signal values, which

are the same values as used by the median filter. The verification consists of two types:

- **Sampling verification**
 When a START bit, a data bit or STOP bit value is sampled (in the middle of a bit period), all three rx_synced signal values should be the same (a '0', '0', '0' sequence or a '1', '1', '1', sequence).
- **Generic verification**
 The isolated '0' or '1' values may not occur (a '1', '0', '1' sequence or a '0', '1', '0' sequence)

When the noise filter is enabled (LINx_CHy_CTL0.FILTER_EN is '1'), the error flag LINx_CHy_INTR.RX_NOISE_DETECT is set in case of a verification failure. An ongoing frame is not aborted by the noise detection.

Figure 26-10. LIN Noise Filter Block Diagram





Figure 26-11. LIN Noise Filtering Timing Diagram



26.11 Interrupts

26.11.1 Overview

The LIN module supports multiple LIN channels; each LIN channel has its dedicated interrupt line and accordingly its own set of interrupt registers LINx_CHy_INTR, LINx_CHy_INTR_SET, LINx_CHy_INTR_MASK, and LINx_CHy_INTR_MASKED.

To reduce interrupt load of the interrupt source flags listed in the LINx_CHy_INTR register, an AND masking is done with the LINx_CHy_INTR_MASK. The masked interrupts, which cause interrupt on the interrupt controller, are shown in the LINx_CHy_INTR_MASKED register

Data	Register
00000111	LIN_CH_INTR
AND 00000111	LIN_CH_INTR_MASK
00000111	LIN_CH_INTR_MASKED

The following tables give an overview of the interrupt events in the module in different modes.

Table 26-6. Interrupt Events in LIN Master Mode

Event Type	Event	Event Detection Condition	Clear Event Flag	Transfer Abort	Enable Interrupt	Register Flag Bit
TX	Header Transmission done	Header transmission succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. TX_HEADER_DONE
TX	Response Transmission done	Response transmission succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. TX_RESPONSE_DONE
TX	Wakeup Transmission done	Wake up signal successfully transmitted	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. TX_WAKEUP_DONE
RX	Response Reception done	Response reception succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. RX_RESPONSE_DONE
RX	Wakeup Reception done	Wake up signal received, after wake up reception detection was enabled.	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. RX_BREAK_WAKEUP_DONE
Error	Time out	A frame, header or response does not finish within a specified time	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	no	yes	LINx_CHy_INTR. TIMEOUT
TX Error	Transmitter Header Bit Error	The incoming bus level does not match with the transmitted value during: - header transmission - wake up transmission	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes ^a	yes	LINx_CHy_INTR. TX_HEADER_BIT_ERROR.

Table 26-6. Interrupt Events in LIN Master Mode

Event Type	Event	Event Detection Condition	Clear Event Flag	Transfer Abort	Enable Interrupt	Register Flag Bit
TX Error	Transmitter Response Bit Error	During the response transmission the received bus value does not match with the transmitted value	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes ^a	yes	LINx_CHy_INTR. TX_RESPONSE_BIT_ERROR
RX Error	Noise Detection	Noise on RX input detected, when LINx_CHy_CTL0. FILTER_EN is '1'	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	no	yes	LINx_CHy_INTR. RX_NOISE_DETECT
RX Error	Receiver Response Frame Error	An invalid start bit or stop bit occurs during response reception (data field, checksum)	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_RESPONSE_FRAME_ERROR
RX Error	Receiver Response Checksum Error	The calculated checksum over the data bytes and optionally the PID field does match with the received checksum.	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_RESPONSE_CHECKSUM_ERROR

a. When LINx_CHy_CTL0.BIT_ERROR_IGNORE is '1', then bit errors are still reported, but do not abort an ongoing transfer.

Table 26-7. Interrupt Events in LIN Slave Mode

Event Type	Event	Event Detection Condition	Clear Event Flag	Transfer Abort	Enable Interrupt	Register Flag Bit
TX	Response Transmission done	Response transmission succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. TX_RESPONSE_DONE
TX	Wakeup Transmission done	Wake up signal successfully transmitted	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. TX_WAKEUP_DONE
RX	Header Reception done	Header reception succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. RX_HEADER_DONE
RX	Response Reception done	Response reception succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. RX_RESPONSE_DONE
RX	Wakeup Reception done	Wake up signal received, after wake up reception detection was enabled.	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. RX_BREAK_WAKEUP_DONE
RX	Synchronization Field Reception done	Synchronization field successfully received	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. RX_HEADER_SYNC_DONE
Error	Time out	A frame, header or response does not finish within a specified time	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	no	yes	LINx_CHy_INTR. TIMEOUT
TX Error	Transmitter Response Bit Error	The incoming bus level does not match with the transmitted value during the response	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes ^a	yes	LINx_CHy_INTR. TX_RESPONSE_BIT_ERROR
RX Error	Noise Detection	noise on RX input detected, when LINx_CHy_CTL0. FILTER_EN is '1'	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	no	yes	LINx_CHy_INTR. RX_NOISE_DETECT

Table 26-7. Interrupt Events in LIN Slave Mode

Event Type	Event	Event Detection Condition	Clear Event Flag	Transfer Abort	Enable Interrupt	Register Flag Bit
RX Error	Receiver Header Frame Error	- An invalid start bit occurs during PID field. - An invalid stop bit occurs during SYNC or PID field.	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_HEADER_FRAME_ERROR
RX Error	Receiver Synchronization Error	An invalid data field pattern is detected during the reception of the SYNC field	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_HEADER_SYNC_ERROR
RX Error	Receiver PID Parity Error	The received PID field has a parity error	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_HEADER_PARITY_ERROR
RX Error	Receiver Response Frame Error	An invalid stop bit occurs during response reception (data field, checksum)	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_RESPONSE_FRAME_ERROR
RX Error	Receiver Response Checksum Error	The calculated checksum over the data bytes and optionally the PID field does match with the received checksum.	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_RESPONSE_CHECKSUM_ERROR

a. When LINx_CHy_CTL0.BIT_ERROR_IGNORE is '1', then bit errors are still reported, but do not abort an ongoing transfer.

Table 26-8. Interrupt Events in UART Mode

Event Type	Event	Event Detection Condition	Clear Event Flag	Transfer Abort	Enable Interrupt	Register Flag Bit
TX	Transmission done	Transmission succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. TX_HEADER_DONE
RX	Reception done	Reception succeeded	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	-	yes	LINx_CHy_INTR. RX_HEADER_DONE
TX Error	Transmitter Bit Error	The incoming bus level does not match with the transmitted value during transmission	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes ^a	yes	LINx_CHy_INTR. TX_HEADER_BIT_ERROR
RX Error	Noise Detection	noise on RX input detected, when LINx_CHy_CTL0. FILTER_EN is '1'	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	no	yes	LINx_CHy_INTR. RX_NOISE_DETECT
RX Error	Receiver Frame Error	An invalid start bit resp. stop bit occurs during reception	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_HEADER_FRAME_ERROR
RX Error	Receiver Parity Error	The received PID field has a parity error	- Write '1' to flag - LINx_CHy_CTL. ENABLED to '0'	yes	yes	LINx_CHy_INTR. RX_HEADER_PARITY_ERROR

a. When LINx_CHy_CTL0.BIT_ERROR_IGNORE is '1', then bit errors are still reported, but do not abort an ongoing transfer.

26.11.2 Transmission Interrupts

26.11.2.1 TX Header Done


After a successful header transmission as master the flag LINx_CHy_INTR.TX_HEADER_DONE is activated. This means, the flag is set after the valid PID STOP bit verification. The enabled command bits such as LINx_CHy_CMD.TX_HEADER within this frame session are not cleared, as long as a selected legal command sequence

(see [LIN Operation on page 463](#)) is not successfully completed.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

Figure 26-12. TX Header Done Flag Timing Diagram




26.11.2.2 TX Response Done

After a valid completion of a frame including the CHECKSUM STOP bit, the LINx_CHy_INTR.TX_RESPONSE_DONE flag is activated; that is, the flag is set after the valid CHECKSUM STOP bit verification. The enabled commands such as LINx_CHy_CMD.TX_RESPONSE within this frame session are not cleared, as long as a selected legal command sequence (see [LIN Operation on page 463](#)) is not successfully completed.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

Figure 26-13. TX Response Done Flag Timing Diagram



26.11.2.3 TX Wakeup Done

To support remote wakeup detection, the header reception commands LINx_CHy_CMD.RX_HEADER and LINx_CHy_CMD.TX_HEADER are in cleared state. At the end of the successfully transmitted dominant wake up pulse the flag LINx_CHy_INTR.TX_WAKEUP_DONE is set to '1'.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

Note: The flag is not set when LINx_CHy_INTR.TX_HEADER_BIT_ERROR is set due to transmission error.

26.11.3 Reception Interrupts

26.11.3.1 RX Break Wakeup Done

After transition from the break low pulse to the break delimiter bit, a break detection interrupt is set by the LINx_CHy_INTR.BREAK_WAKEUP_DONE flag. This interrupt flag does not need to be enabled for the regular header processing.

As the wakeup function is shared with the break function the end of the wakeup pulse detection is represented by the same flag.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

26.11.3.2 RX Header SYNC Done

After reception of a valid SYNC byte pattern and valid SYNC STOP bit the LINx_CHy_INTR.RX_HEADER_DONE flag is set.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).


26.11.3.3 RX Header Done

After reception of a valid LIN header including a valid PID STOP bit and PID parity check, the LINx_CHy_INTR.RX_HEADER_DONE flag is set. The command bit LINx_CHy_CMD.RX_HEADER is not cleared, as long as a legal command sequence (see [LIN Operation on page 463](#)) is not successfully completed.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

Figure 26-14. "RX Header Done" Flag Timing Diagram




26.11.3.4 RX Response Done

After a valid completion of a frame including CHECKSUM STOP bit and the checksum verification the LINx_CHy_INTR.RX_RESPONSE_DONE flag is set. The enabled commands such as LINx_CHy_CMD.TX_RESPONSE within this frame session are not cleared, as long as a selected legal command sequence (see [LIN Operation on page 463](#)) is not successfully completed.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

Figure 26-15. "RX Response Done" Flag Timing Diagram



26.11.4 Error and Status Interrupts

To ensure robust behavior, several types of errors are detected. When an error is detected, the associated interrupt cause in the INTR register is activated. [Figure 26-16](#) and [Figure 26-17](#) give an overview about the appearance of error events for the LIN master and LIN slave.

Figure 26-16. LIN Master Error Events Timing Diagram





Figure 26-17. LIN Slave Error Events Timing Diagram



Notes:

- When the LINx_CHy_CTL0.BIT_ERROR_IGNORE is '1', a bit error (the timeout error is not included) does not abort an ongoing transfer, although the bit errors are always reported.
- As the transmission commands (such as TX_REPONSE) have higher priority than the reception commands (such as RX_RESPONSE) in the processing order the transmission errors are only reported, when both commands are activated.

26.11.4.1 Transmitter Bit Error

During transmission the transmitted value on the RX line is also received over the TX line. The transmitted and received values should be the same. If this verification detects a failure, an LINx_CHy_INTR.TX_HEADER_BIT_ERROR or LINx_CHy_INTR.TX_RESPONSE_BIT_ERROR is activated and the transmission is automatically aborted by

the hardware. This also includes the detection of an invalid START and STOP bit.

The error flag LINx_CHy_INTR.TX_HEADER_BIT_ERROR is valid for:

- Break field
- Synchronization field
- PID field
- Wake up low pulse


The error flag LINx_CHy_INTR.TX_RESPONSE_BIT_ERROR is valid for:

- Data fields
- Checksum field

Clearing the flag

Both flags can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

Figure 26-18. Transmitter Bit Error Timing Diagram



26.11.4.2 Receive Synchronization Error

The slave experiences a synchronization error, when SYNC byte pattern is either incorrect or the synchronization range is exceeded. The error is shown by the LINx_CHy_INTR.RX_HEADER_SYNC_ERROR flag.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

26.11.4.3 Receiver Frame Error

A START bit should be received as a '0' on the RX line and a STOP bit should be received as a '1' on the RX line. A START bit occurs at specific moments in the frame after a falling edge on the RX line and a STOP bit occurs after every 8-bit field. The error is detected after the sample time of the RX line, which is in the center of bit period (see [Baud Rate and Sample Point on page 456](#)).

Header Reception

When a frame error is detected during the header the LINx_CHy_INTR.RX_HEADER_FRAME_ERROR flag is set. The ongoing transfer is aborted automatically.

Response Reception

During the response, the LINx_CHy_INTR.RX_RESPONSE_FRAME_ERROR flag is activated, when the frame error occurs in the data bytes 2 to 8 or in the checksum. Additionally the ongoing response reception is aborted by the hardware.

Exception: Framing Error in Data Byte 1

In the LIN "No response" error, the response part is missing and followed by a LIN break. The event flag LINx_CHy_INTR.RX_RESPONSE_DONE and error flag

LINx_CHy_INTR.RX_RESPONSE_FRAME_ERROR stay '0'. Whereas when a detected invalid STOP bit is followed by a START bit detection, then LINx_CHy_INTR.RX_RESPONSE_FRAME_ERROR is set to '1'.

Clearing the flag

Both flags can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0). The LINx_CHy_STATUS.RX_DATA0_FRAME_ERROR is cleared automatically after start of a new response.

26.11.4.4 Receiver PID Parity Error

The receiver calculates the parity bits over the received frame identifier in the PID field. The calculated parity bits are verified against the received parity bits in the PID field. In case of verification failure, the LINx_CHy_INTR.RX_HEADER_PARITY_ERROR flag is set.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

26.11.4.5 Response Checksum Error

The receiver calculates the checksum over the received PID field (optionally as specified by the LINx_CHy_CTL0.CHECKSUM_ENHANCED register field) and the received data fields. The calculated checksum is verified against the received checksum field. In case of verification failure, the LINx_CHy_INTR.RX_RESPONSE_CHECKSUM_ERROR is activated.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

26.11.4.6 Receiver Noise Detection

When the noise filter is enabled (LINx_CHy_CTL0.FILTER_EN is '1'), the error flag LINx_CHy_INTR.RX_NOISE_DETECT is set in case of a verification failure. But a going transfer is not aborted. See [Noise Filter on page 467](#) for more details.

Clearing the flag

The flag can be cleared either by a write access to the flag with '1' within the LINx_CHy_INTR register or disabling the LIN channel (LINx_CHy_CTL0.ENABLED = 0).

Note: An ongoing frame is not aborted by the noise detection.

26.11.4.7 Timeout Detection

As described in [Timeout Operation on page 460](#), the timer operation inside the LIN module is supported. When one of the selected timeouts is detected, the LINx_CHy_INTR.TIMEOUT flag is activated.

Note: The timeout detection does not abort an ongoing frame.

26.12 Registers

Table 26-9. LIN Global Unit Registers

Register	Name	Description
LINx_ERROR_CTL	Error Control Register	Error injection control for the full LIN unit.
LINx_TEST_CTL	Test Control Register	Test control is done for all channels.

Table 26-10. LIN Channel Registers

Register	Name	Description
LINx_CHy_CTL0	Control 0 Register	In this register the channel can be enabled. Furthermore the communication mode selection and mode configurations are provided.
LINx_CHy_CTL1	Control 1 Register	Beside the LIN data length and the checksum the timeout counter is processed in the register.
LINx_CHy_STATUS	Status Register	The communication state flags and the error flags, which are mirrored from the INTR register, are listed.
LINx_CHy_CMD	Command Register	The communication protocol is controlled.
LINx_CHy_TX_RX_STATUS	TX/RX Status Register	An input and output status of the LIN transceiver control is reported. Additionally the LIN synchronization counter provides a counter value, which needs to be processed for the synchronization procedure in software.
LINx_CHy_PID_CHECKSUM	PID Checksum Register	PID and checksum buffer.
LINx_CHy_DATA0	Data 0 Register	The response buffer for the data byte fields 0 to 3 is covered.
LINx_CHy_DATA1	Data 1 Register	The response buffer for the data byte fields 4 to 7 is covered.
LINx_CHy_INTR	Interrupt Register	The status of communication and error flags is shown.
LINx_CHy_INTR_SET	Interrupt Set Register	Communication and error flags in the INTR register can be set for test purposes.
LINx_CHy_INTR_MASK	Interrupt Mask Register	A bit mask over the communication and error flags can be defined.
LINx_CHy_INTR_MASKED	Interrupt Masked Register	Masked communication and error flags are listed.

Note: In LINx_CHy, 'x' signifies the LIN instance and 'y' is the channel number under the LIN instance.

27. Cryptography Block



The Cryptography block (Crypto) provides hardware implementation and acceleration of cryptographic functions. Implementation in hardware takes less time and energy than the equivalent firmware implementation. In addition, the block provides true random number generation functionality in silicon, which is not available in firmware. It supports symmetric key encryption and decryption, hashing, message authentication, random number generation (pseudo and true), cyclic redundancy checking, and utility functions such as enable/disable, interrupt settings, and flags.

27.1 Features Overview

The cryptography function block of Traveo II supports the following features:

- **Advanced Encryption Standard (AES) functionality according to FIPS 197:**
The AES component can be used to encrypt/decrypt data blocks of 128-bit length and supports programmable key length (128/192/256-bit key).
- **CHACHA20 functionality according to RFC 7539:**
CHACHA20 is a stream cipher, which produces output consisting of 512-bit random-looking bits. These random-bits can be XORed with plain-text to produce cipher-text.
- **Triple Data Encryption Standard (TDES):**
The TDES component can be used to encrypt/decrypt data blocks of 64-bit length using a 64-bit key.
- **Secure Hash Algorithm (SHA) functionality according to FIPS 180-4/FIPS-202:**
This component can be used to produce a fixed-length hash (also called “message digest”) of up to 512 bits from a variable-length input data (called “message”). SHA1, SHA2, SHA3 hashes are supported.
- **Cyclic Redundancy Check (CRC) functionality:**
This component performs a cyclic redundancy check with a programmable polynomial of up to 32-bits.
- **String (STR) functionality:**
This component can be used to efficiently copy, set, and compare memory data.
- **Pseudo Random Number Generator (PR):**
This component generates pseudo random numbers in a fixed range. This generator is based on three Linear Feedback Shift Registers (LFSRs).
- **True Random Number Generator (TR):**
This component generates true random numbers of up to 32 bits using ring oscillators.
- **Vector Unit (VU):**
This component act as coprocessors to offload asymmetric key operations from the main processor.
- **AHB master-interface:**
This allows to fetch operands directly from the system memory.
- **Device Key functionality:**
The device key has its usage restricted to specific functionality; they cannot be accessed by the software that implements that functionality. Two independent device keys are supported.


27.2 System Diagram

The Cryptography block provides the cryptography functionality on Traveo II MCU. The complete cryptography implementation is done in conjunction with third-party software. In a secure system implementation, the cryptography block can be accessed only by the secure master (CM0+). For other masters to avail any cryptography services, they need to request CM0+ via system calls using IPC. For details, see the [Inter-Processor Communication chapter on page 48](#).

27.3 Block Diagram

This section explains the major components within the cryptography block.

Figure 27-1. High-level Block Diagram



The cryptography block provides cryptographic functionality:

- DES, Triple DES, AES, and Chacha20 symmetric key ciphers.
- SHA1, SHA2, and SHA3 hashes.
- Pseudo and true random number generators.
- Vector unit for asymmetric key cryptography.
- CRC functionality.

- A single interrupt signal is used to indicate the completion of an operation.
- A clock and reset signal interface connects to the System Resources subsystem (SRSS). The cryptography block operates a gated version of "clk" and uses both Active and DeepSleep reset signals.

The cryptography block is connected to the AHB-Lite bus infrastructure through a slave bus interface and a master bus interface. The block has the following interfaces:

- An AHB-Lite slave interface connects the cryptography block to the AHB-Lite infrastructure. This interface supports 8/16/32-bit AHB-Lite transfers. MMIO registers accesses are 32-bit accesses only (8/16-bit accesses to MMIO registers results in an AHB-Lite bus error). Memory buffer accesses can be 8/16/32-bit accesses.
- An AHB-Lite master interface connects the cryptography block to the AHB-Lite infrastructure. This interface supports 8/16/32-bit AHB-Lite transfers. The interface enables the Crypto block to access operation operand data from system memory (for example, flash or SRAM).

27.4 Function Description

The basic functions of the cryptography block are described here.

27.4.1 Operating Mode

The cryptography block operates only in Active/Sleep/LPActive/LPSleep power modes. In DeepSleep mode, the block retains only the contents of its retention registers with optional retention of internal SRAM contents.

27.4.2 Memory Map and Register Definitions

The memory map and register definitions for the cryptography block are located in the product register map.

27.4.3 Instruction Set

Most operations in the cryptography block are initiated through an instruction by CM0+ via IPC.

28. Event Generator (EVTGEN)



The Event Generator (EVTGEN) in Traveo II implements event generation for interrupts and triggers in Active mode and only interrupts in DeepSleep mode. The Active functionality interrupt is connected to the CPU interrupt controller. Active trigger events can be used to trigger a specific device functionality mode (for example, execution of an interrupt handler, a SAR ADC conversion, and so on) in Active power mode. The DeepSleep functionality interrupts can be used to wake up the CPU from the DeepSleep power mode. The event generator includes a single counter and a maximum of 16 comparator structures for each Active and DeepSleep mode. EVTGEN reduces CPU involvement and thus overall power consumption and jitters.


This chapter explains the features, implementation, and operational modes of the Event Generator block

28.1 Features

- CPU-free triggers for device functions
- Reduces CPU involvement in triggering device functions, thus reduces overall power consumption and CPU bandwidth
- 16 comparators for each DeepSleep and Active mode to generate interrupts and triggers
- 32-bit counter, one each for DeepSleep and Active mode for comparators
- Individual configurable thresholds for each comparator
- DeepSleep and Active mode clock sources for counters
- Jitter-free initiation of specific device functionality
- One DeepSleep and one Active mode interrupt for CPU
- Supported in Active, Sleep, LPActive, LPSleep, and DeepSleep power modes

28.2 Block Diagram

Figure 28-1. EVTGEN Block Diagram



The EVTGEN consists two blocks: Active and DeepSleep mode blocks. There are 16 comparator structures and one 32-bit counter for each of the modes. The EVTGEN block has these interfaces:

- Bus interface – Connects the block to the CPU subsystem.
- Trigger Interface – Provides one trigger signal from each Active mode comparators.
- System interface – Consists of control signals such as clock and reset from the system resources subsystem.
- Interrupts – Provides one interrupt signal from Active and DeepSleep mode blocks, based on the comparator outputs.

This EVTGEN block can be configured by writing to the EVTGEN registers. See [28.2.7 Register List](#) for more information on all registers required for this block.

28.2.1 Enabling and Disabling EVTGEN Block

The EVTGEN block can be enabled by setting the Enable bit of the EVTGENx_CTL register. All non-retention (not retained in Sleep mode) registers (command and status registers) are reset to their default value when this is disabled. All retention (retained in Sleep mode) registers retain their value when this is enabled.

28.2.2 Counters

There is one 32-bit counter for each of the Active and the DeepSleep modes. These counters keep track of time; the

time measured is referenced with respect to the CLK_REF clock.

28.2.2.1 Clock and Prescaling


The counter working is based on the following two clocks from CLK_SYS.

- **CLK_REF:** Time is measured with respect to a divided version of this clock – CLK_REF_DIV. The divider value is provided by EVTGENx_REF_CLOCK_CTL.INT. The CLK_REF_DIV clock is assumed to have a higher frequency and a higher precision than CLK_LF. The clock is available only in Active power mode. Typically, CLK_REF is connected to a high-precision SRSS clock source (for example, a PLL).
- **CLK_LF:** This is a low-frequency clock (typically around 16 kHz to 32 kHz). The clock is assumed to have a lower precision than CLK_REF. It is available in both Active and DeepSleep power modes. Typically, CLK_LF is connected to an SRSS 32-kHz low-frequency clock.

Comparator components are used to compare time with a programmed value and generate control signals when the counter exceeds the programmed value. The clock CLK_REF_DIV provides fine resolution (high frequency) and high precision. This clock is not available in DeepSleep power mode. DeepSleep control signals are generate based only on CLK_LF.

- Clock CLK_REF used to generate Active control signals.
- Clock CLK_LF is used to generate DeepSleep control signals.

Figure 28-2. Counter Block Diagram



28.2.2.2 Ratio

All the count registers and comparator count thresholds are expressed with respect to CLK_REF_DIV domain.

The number of CLK_REF_DIV cycles per CLK_LF cycle can be controlled in either software or hardware.

$$\text{Ratio} = \frac{\text{CLK\_REF\_DIV}}{\text{CLK\_LF}}$$

28.2.2.3 Software Control

The software control is provided through EVTGENx_RATIO. This register contains a ratio value expressing the relative frequency of CLK_LF with respect to CLK_REF_DIV. Specifically, this registers contains the average number of CLK_REF_DIV cycles per CLK_LF cycle. The RATIO value has a 16-bit integer component (EVTGENx_RATIO.INT) and an 8-bit fractional component (EVTGENx_RATIO.FRAC). This register is retained in DeepSleep mode.

28.2.2.4 Hardware Control

Hardware control auto-calibrates and makes the New Ratio value available through the EVTGENx_RATIO register. Auto-calibration continuously measures the ratio and combines this new measurement with the current RATIO value to calculate the new RATIO value. This calculation is based on a weighted average operation. The weights of the new measurement and the current RATIO value are controlled through EVTGENx_RATIO_CTL register.

The weighted average operation addresses jitter in the low-frequency clock CLK_LF. The weights of the weighted average calculation try to trade off clock CLK_LF jitter sensitivity and speed of adaptability to a new clock CLK_LF frequency. Note that gradual changes to CLK_LF may occur to differences in operating conditions (such as temperature).

Auto calibration is Active functionality logic; that is, the RATIO value is not updated in DeepSleep power mode. However, the RATIO value is retained in DeepSleep mode.

Hardware control is enabled through DYNAMIC bit in EVTGENx_RATIO_CTL register. The **weight** for average calculation is the 3-bit value, which can be set using DYNAMIC_Mode bits in EVTGENx_RATIO_CTL.

The EVTGENx_RATIO register fields INT16 and FRAC8 are only valid when the VALID bit in EVTGENx_RATIO_CTL is one. This register is retained in DeepSleep mode.

The RATIO value is required in the EVTGENx_RATIO.INT16 and EVTGENx_RATIO.FRAC8 register fields. Either:

- Hardware establishes the RATIO value. This process takes a maximum of two CLK_LF cycles.
- Software provides the RATIO value in the EVTGENx_RATIO.INT16 and

EVTGENx_RATIO.FRAC8 register fields. This process is immediate.

The RATIO value is used to initialize the DeepSleep counter. This process takes one CLK_LF cycle.

28.2.3 Counter Status

The Active counter functionality is available only in Active power mode. This is an UP counter and auto reloads itself. The Active counter is not retained in DeepSleep power mode. The status of active counter can be read through EVTGENx_COUNTER register. This register is not retained in DeepSleep mode.


The Active COUNTER register value is only valid when the EVTGENx_COUNTER_STATUS has valid bit set.

- After a DeepSleep to Active power mode transition, the Active counter is not immediately valid. On the first CLK_LF clock after the power mode transition, the DeepSleep counter value is used to initialize the Active counter.
- On every CLK_REF_DIV clock, the Active counter is incremented by '1'.

The DeepSleep counter functionality is available in both the Active and DeepSleep power modes.

On every CLK_LF, the RATIO value is added to the DeepSleep counter status. On every CLK_LF cycle, the status of DeepSleep counter will be the same as that of the Active Counter. The status of DeepSleep counter is not accessible in Active or DeepSleep mode.

Figure 28-3 illustrates an example where CLK_REF_DIV is five times as fast as CLK_LF (RATIO = 5).

Figure 28-3. Active and DeepSleep Counter Status with $RATIO = 5$


After a wakeup from DeepSleep power mode, the Active counter is re-initialized. The DeepSleep counter will initialize the Active counter and `EVTGENx_COUNTER_STATUS.VALID` is set to '1'. This process will take at most two `CLK_LF` cycle.

If the `RATIO` value (the average number of `CLK_REF_DIV` cycles per `CLK_LF` cycle) is not established, the DeepSleep counter is not valid.

If the `RATIO` value is established:

- On the first `CLK_LF` clock, the current `RATIO` value is used to initialize the DeepSleep counter.
- On every other `CLK_LF` clock (DeepSleep counter is initialized), the current `RATIO` value is added to the DeepSleep counter. Note that the DeepSleep counter has a fractional component.

The EVTGEN block has the following clocking conditions that needs to be adhered. F_{EVTGEN} is the clock frequency of the EVTGEN block.

1. $F_{EVTGEN} \geq F_{CLK\_REF\_DIV}$ and $F_{EVTGEN} \geq F_{CLK\_REF}$
2. $F_{CLK\_REF\_DIV} \geq 4 \times F_{CLK\_LF}$

28.2.4 Comparator Structures


The EVTGEN block supports 16 comparator structures in Active mode and 16 comparator structures in DeepSleep mode.

One set of Active and DeepSleep mode comparators have one individual control register. Each comparator structure compares the Active `COUNTER_INT` and DeepSleep `COUNTER_INT_LF` with an Active and DeepSleep compare value respectively.

The Active functionality is enabled through enable `EVTGENx_COMP_STRUCTy_COMP_CTL.COMP0_EN` bit in the comparator control register. Similarly, `EVTGENx_COMP_STRUCTy_COMP_CTL.COMP1_EN` bit in the comparator control register enables/disables DeepSleep comparator. There is one of this register for every pair consisting of an Active and a DeepSleep comparator. This register is retained in DeepSleep mode.

Trigger method for the Active comparator can be selected through `TR_OUT_EDGE`.

Figure 28-4. Comparator Structure



The Active counter `EVTGENx_COUNTER` is compared to the `EVTGENx_COMP_STRUCTy_COMP0` register value.

- The Active comparator `COMP0_OUT` output is activated when the Active counter becomes greater than or equal to the Active compare value.
- The Active comparator `COMP0_OUT` output is deactivated, when the comparator is disabled (`EVTGENx_COMP_STRUCTy_COMP_CTL.COMP0_EN`).

The DeepSleep compare functionality is available in both Active and DeepSleep power modes. The functionality is enabled through `EVTGENx_COMP_STRUCTy_COMP_CTL.COMP1_EN`. The DeepSleep counter status is compared to the `EVTGENx_COMP_STRUCTy_COMP1` register value.

- The DeepSleep comparator `COMP1_OUT_LF` output is activated, when the DeepSleep counter becomes greater than or equal to the DeepSleep compare value.
- The DeepSleep comparator `COMP1_OUT_LF` output is deactivated, when the comparator is disabled (`EVTGENx_COMP_STRUCTy_COMP_CTL.COMP1_EN`).

`EVTGENx_COMP_STRUCTy_COMP0` and `COMP` registers hold the compare values for the Active and DeepSleep comparators respectively. There is one register for every one pair consisting of one Active and one DeepSleep comparator. These registers are retained in DeepSleep mode.

When the Active counter is initialized (`EVTGENx_COUNTER_STATUS.VALID` is '1'), `COUNTER` and future `EVTGENx_COMP_STRUCT_COMP0` and `EVTGENx_COMP_STRUCT_COMP1` comparator values can be read and programmed. These future comparator values should have a minimum delay with respect to the Active counter value `COUNTER`. This is to ensure that the counter value has not passed these future values before the comparators are enabled.

The value written to `EVTGENx_COMP_STRUCT_COMP0` or `EVTGENx_COMP_STRUCT_COMP1` should be at least four `CLK_LF` cycles ahead of the current `COUNTER` value. A comparator structure "i" produces a `EVTGEN_TR_OUT[i]` trigger and interrupt cause signals.

The Active functionality `EVTGEN_TR_OUT[i]` trigger is available only in Active power mode.

- The `EVTGEN_TR_OUT[i]` trigger is activated, when the Active comparator status `COMP0_OUT[i]` is activated.
- The `EVTGEN_TR_OUT[i]` trigger is deactivated, when the comparator is disabled (`EVTGENx_COMP_STRUCTy_COMP_CTL.COMP0_EN`).

The trigger `EVTGEN_TR_OUT[i]` can be used to trigger specific device functions such as execution of an interrupt

handler, a SAR ADC conversion, and a LIN message transfer.

The status an Active comparator can be read from corresponding bit in `EVTGENx_COMP0_STATUS` register. The status an DeepSleep comparator can be read from corresponding bit in `EVTGENx_COMP1_STATUS` register. These register are retained in DeepSleep mode.

The Active interrupt cause signal is available only in Active power mode.

- The cause is activated, when the Active comparator is activated.
- The cause is activated by software through `EVTGENx_INTR_SET.COMP0[i]`.

The DeepSleep interrupt cause signal is available in Active and DeepSleep power modes.

- The cause is activated, when the DeepSleep comparator is activated.
- The cause is activated by software through `EVTGENx_INTR_DPSP.LP.COMP1[i]`.

Typically, the Active and DeepSleep comparators are used as follows:


- The DeepSleep comparator value (`EVTGENx_COMP_STRUCT_COMP1`) is set to a time `X` on `CLK_LF`. The intent is to wake up the device (SRSS wakeup signal) and to ensure that the device is in Active power mode at time `X`+wakeup time.
- The Active comparator value (`EVTGENx_COMP_STRUCT_COMP0`) is set to a time `X`+wakeup time. The intent is to use the associated output trigger `EVTGEN_TR_OUT[i]` to initiate a specific device functionality in Active power mode.
- On completion of the specific device functionality, the functions interrupt signal is activated. The CPU interrupt handler may process the result of the specific functionality. The CPU may also setup the Event generator and may turn the device to DeepSleep power mode through a WFI instruction (resulting in activation of the SRSS DeepSleep signal).

28.2.5 Interrupts

The EVTGEN block uses two interrupts:

- An Active interrupt signal. Each Active comparator has a dedicated interrupt cause field. This interrupt notifies the CPU, when an output trigger is activated.
- A DeepSleep INT_DPSLP signal. Each DeepSleep comparator has a dedicated interrupt cause field. This interrupt is connected to CPUs' WICs and allows for wakeup from DeepSleep power mode.
- The Active and DeepSleep interrupts are available in the system interrupt sources.

Figure 28-5. Interrupt Process



The Active interrupt cause field register bit is set `EVTGENx_INTR` when a corresponding comparator 0 (`COMP0_OUT`) event is generated. Each bit in this corresponds to one comparator. Writing one to this register will clear it. The `EVTGENx_INTR` register is not retained in DeepSleep mode.

The DeepSleep interrupt cause field register bit is set `EVTGENx_INTR_DPSLP` when a corresponding comparator 1 (`COMP1_OUT_LF`) event is generated. Each bit in this corresponds to one comparator. The `EVTGENx_INTR_DPSLP` register is retained in DeepSleep mode.

The `EVTGENx_INTR_SET` register when read, reflects the `EVTGENx_INTR` register. For debug purposes, Software can write a '1' to activate a specific interrupt cause (this allows for debug of the software ISR, without relying on hardware to activate the interrupt cause). Each bit in this corresponds to one comparator. The `EVTGENx_INTR_SET` register is not retained in DeepSleep mode.

`EVTGENx_INTR_DPSLP_SET` register when read reflects the `EVTGENx_INTR_DPSLP` register. For debug purposes, software can write a '1' to activate a specific interrupt cause (this allows for debug of the software ISR, without relying on hardware to activate the interrupt cause). Each bit in this corresponds to one comparator. The `EVTGENx_INTR_DPSLP_SET` register is retained in DeepSleep mode.

The mask for corresponding bit field of active comparator can be set using `EVTGENx_INTR_MASK` register. The `EVTGENx_INTR_MASK` register is retained in DeepSleep mode.

The mask for corresponding bit field of DeepSleep comparator can be set using `EVTGENx_INTR_DPSLP_MASK` register. The `EVTGENx_INTR_MASK` register is retained in DeepSleep mode.

`EVTGENx_INTR_MASKED` register reflect the logical AND of corresponding `EVTGENx_INTR` with `EVTGENx_INTR_MASK` fields. The `EVTGENx_INTR_MASKED` register is not retained in DeepSleep mode.

`EVTGENx_INTR_DPSLP_MASKED` register reflects logical AND of corresponding `EVTGENx_INTR_DPSLP` with `EVTGENx_INTR_DPSLP_MASK` fields. The `EVTGENx_INTR_DPSLP_MASKED` register is retained in DeepSleep mode.


Logical OR operation is applied to all the bit field of `EVTGENx_INTR_MASKED` register to generate the active interrupt signal that is connected to the CPU interrupt controller

Logical OR operation is applied to all the bit field of `EVTGENx_INTR_DPSLP_MASKED` register to generate the DeepSleep interrupt signal that is connected to the CPUs' wakeup interrupt controllers (WICs). When enabled in WIC, this signal activation will wake up the CPU from DeepSleep power mode to Active power mode.

28.2.6 Use Case

The following waveform illustrates a single Wakeup/DeepSleep sequence (the wakeup and DeepSleep signals initiate SRSS power mode transitions, the EVTGEN_TR_OUT[i] trigger initiates Active functionality).

Figure 28-6. Use Case Waveform



The waveforms are explained as follows.

1. The CPU reads the Active counter and sets the DeepSleep comparator to wake up the device at time 3 and the Active comparator to generate a trigger at time 4.
2. The CPU enters DeepSleep power mode. For example, a WFI instruction with the DeepSleep bit field set to '1'. When the CPU is in DeepSleep power mode, its DeepSleep signal is activated '1'.
3. The event generator activates its DeepSleep interrupt and the WIC activates its SRSS wakeup request.
4. The event generator activates the trigger and device-specific functionality is initiated. The CPU may set up the event generator block before the transition to DeepSleep power mode.
5. The device-specific functionality completes as indicated by the active interrupt. The CPU clears the function's interrupt causes. The CPU may process the results of the device-specific functionality.
6. The CPU enters DeepSleep power mode.

28.2.7 Register List

Register	Name	Description
EVTGENx_CTL	Event Generator Control register	ENABLED - EVTGEN enable: 0: Disabled. All non-retention registers (command and status registers) are reset to their default value when the EVTGEN is disabled. All retention registers retain their value when the EVTGEN is disabled. 1: Enabled.
EVTGENx_REF_CLOCK_CTL	Event Generator Clock divider register	INT_DIV - Divider control for CLK_REF_DIV:"0": Divide by 1, ... "255": Divide by "256". $CLK\_REF\_DIV = \frac{CLK\_REF}{INT\_DIV + 1}$
EVTGENx_RATIO	Event Generator ratio register	INT16 - Integer component of ratio value. FRAC8 - Fractional component of ratio value.
EVTGENx_RATIO_CTL	Event Generator ratio control register	DYNAMIC_MODE - Weighted average calculation (only used when DYNAMIC is '1'): 0: new RATIO value = measurement (no averaging). 1: new RATIO value = (RATIO + measurement + 1) / 2. 2: new RATIO value = (3 × RATIO + measurement + 2) / 4. 3: new RATIO value = (7 × RATIO + measurement + 4) / 8. 4: new RATIO value = (15 × RATIO + measurement + 8) / 16. 5: new RATIO value = (31 × RATIO + measurement + 16) / 32. 6: new RATIO value = (63 × RATIO + measurement + 32) / 64. 7: new RATIO value = (127 × RATIO + measurement + 64) / 128. Note: "measurement" (integer component only) is defined as: 256 × number of measured CLK_REF_DIV cycles per CLK_LF cycle. The RATIO value (integer and fractional component) is defined as: 256 × EVTGENx_RATIO.INT16 + EVTGENx_RATIO.FRAC8 (EVTGENx_RATIO.INT16 = RATIO >> 8 and EVTGENx_RATIO.FRAC8 = RATIO% 256). DYNAMIC - Specifies if EVTGENx_RATIO_CTL.VALID and RATIO are under software or hardware control: 0: Software control. 1: Hardware control. Auto calibration is used to derive the RATIO value. Hardware measures the number of CLK_REF_DIV cycles per CLK_REF cycle. This measurement is combined with the current ratio value to calculate a new ratio value. VALID - Ratio value valid: 0: Invalid. 1: Valid. The RATIO register fields INT16 and FRAC8 are only valid when VALID is '1'. Typically, when the RATIO value is under software control this field is set to '1' and when it is under hardware control this field is set to '0'. However, when the RATIO value is under hardware control, it is possible for this field to be set to '1' and to have software provide an initial RATIO.
EVTGENx_COUNTER	Event Generator counter register	INT32 - Active counter COUNTER_INT on CLK_REF_DIV.
EVTGENx_COUNTER_STATUS	Event Generator Control status register	VALID - Active counter validity: 0: Invalid. 1: Valid. The EVTGENx_COUNTER register field INT32 is only valid when VALID is '1'.

Register	Name	Description
EVTGENx_COMP_STRUCTy_COMP_CTL	Event Generator comparator control register	COMP0_EN - Active comparator (COMP0) enable: 0: Enabled. 1: Disabled. The comparator output COMP0_OUT is '0'. COMP1_EN - DeepSleep comparator (COMP1) enable: 1: Enabled. 0: Disabled. The comparator output COMP1_OUT_LF is '0'. TR_OUT_EDGE - Specifies the TR_OUT output trigger: 0: The trigger is a level-sensitive trigger. The Active comparator output (COMP0_OUT) is reflected on TR_OUT. 1: The trigger is an edge-sensitive trigger. Activation of the Active comparator output (rising edge on COMP0_OUT) results in a two cycle '1'/high pulse on TR_OUT. ENABLED - Comparator structure enable: 0: Disabled. 1: Enabled.
EVTGENx_COMP_STRUCTy_COMP0	Event Generator active comparator compare value register	INT32 - This value is a 32-bit unsigned integer in the range $[0, 2^{32} - 1]$. The comparator COMP0_OUT output is activated when the Active counter EVTGENx_COUNTER becomes greater than or equal to COMP0. Note: Software must ensure that EVTGENx_COMP_STRUCTy_COMP_CTL.COMP_EN [0] is '0' when COMP0 is written.
EVTGENx_COMP_STRUCTy_COMP1	Event Generator DeepSleep comparator compare value register	INT32 - This value is a 32-bit unsigned integer in the range $[0, 2^{32} - 1]$. The comparator COMP1_OUT_LF output is activated when the DeepSleep counter COUNTER_INT_LF becomes greater than or equal to COMP1. Note: Software must ensure that EVTGENx_COMP_STRUCTy_COMP_CTL.COMP_EN [1] is '0' when COMP1 is written.
EVTGENx_COMP0_STATUS	Event Generator active comparator status register	COMP0_OUT - Active comparator COMP0_OUT[i] outputs.
EVTGENx_COMP1_STATUS	Event Generator DeepSleep comparator status register	COMP1_OUT - DeepSleep comparator COMP_OUT_LF outputs (synchronized from CLK_LF).
EVTGENx_INTR	Event Generator Interrupt register	COMP0 - This interrupt cause field is activated when a comparator 0 event is generated (Active counter EVTGENx_COUNTER becomes greater than or equal to EVTGENx_COMP0). COUNTER - This interrupt cause field is activated (hardware sets the field to '1') when the Active counter becomes valid (COUNTER_STATUS.VALID is set to '1').
EVTGENx_INTR_DPSLP	Event Generator DeepSleep Interrupt register	COMP1 - This interrupt cause field is activated when a comparator 1 event is generated (DeepSleep counter COUNTER_LF becomes greater than or equal to EVTGENx_COMP1)

Register	Name	Description
EVTGENx_INTR_SET	Event Generator active Interrupt set register	COMP0 - These bits when read reflects bits [15:0] of EVTGENx_INTR register. For debug purposes, software can write a '1' to activate a specific interrupt cause (this allows debug of the software ISR without relying on hardware to activate the interrupt cause). COUNTER - This bit when read reflects bit 16 of EVTGENx_INTR register. For debug purposes, software can write a '1' to activate a specific interrupt cause (this allows debug of the software ISR without relying on hardware to activate the interrupt cause).
EVTGENx_INTR_DPSLP_SET	Event Generator DeepSleep Interrupt set register	COMP1 - These bits when read reflects bits [15:0] of EVTGENx_INTR_DPSLP register. For debug purposes, software can write a '1' to activate a specific interrupt cause (this allows for debug of the software ISR without relying on hardware to activate the interrupt cause).
EVTGENx_INTR_MASK	Event Generator Active Interrupt mask register	COMP0 - Mask bit for corresponding field in the EVTGENx_INTR register.
EVTGENx_INTR_DPSLP_MASK	Event Generator DeepSleep Interrupt mask register	COMP1 - Mask bit for corresponding field in the EVTGENx_INTR_DPSLP register.
EVTGENx_INTR_MASKED	Event Generator Active Interrupt masked register	COMP0 - Logical AND of corresponding EVTGENx_INTR and EVTGENx_INTR_MASK fields.
EVTGENx_INTR_DPSLP_MASKED	Event Generator DeepSleep Interrupt masked register	COMP1 - Logical AND of corresponding EVTGENx_INTR_DPSLP and EVTGENx_INTR_DPSLP_MASK fields.

Note: 'x' signifies the EVTGEN instance, and 'y' signifies the structure number.

29. Trigger Multiplexer



Every peripheral in the TVII-B-H device is interconnected using trigger signals. Trigger signals are means by which peripherals inform the occurrence of an event or transit to a state. These triggers are used to effect or initiate an action in other peripherals. They help the user to route triggers from a source peripheral to a destination.

Triggers are produced by a peripheral and consumed by another. Unlike interrupts, triggers are used to synchronize between peripherals, rather than between a peripheral and the Arm CPU Core.

29.1 Features


Triggers are functional in the active power mode.

- Ability to connect any trigger signal from one peripheral to another,
- Provides up to 16 multiplexer-based trigger groups and up to 16 one-to-one trigger groups
- Supports a software trigger, which can trigger any signal in the block
- Ability to configure a trigger multiplexer with trigger manipulation features in hardware such as inversion and edge/level detection

29.2 Description

Triggers are used to synchronize the functionality of peripherals in hardware (as opposed to software-based synchronization). Peripheral DMA uses triggers to initiate the transfer of a data element from one address location to another. For example, an “ADC conversion done” event may initiate the transfer of an ADC sample from an ADC module to an SRAM memory location.

Triggers are digital signals generated by peripheral blocks to denote a state such as FIFO level, or an event such as completion of an action. These trigger signals typically serve as an initiator of other actions in other peripheral blocks. An example is an ADC peripheral block sampling three channels. After the conversion is complete, a trigger signal will be generated, which in turn triggers a DMA channel that transfers the ADC data to a memory buffer. This example is shown in Figure 29-1.



A TVII-B-H device has multiple peripheral blocks; each of these blocks can be connected to other blocks through trigger signals, based on the system implementation. To support this, the device has hardware, which is a series of multiplexers used to route the trigger signals from potential sources to destinations. This hardware is called the trigger multiplexer block. The trigger multiplexer can connect to any trigger signal emanating out of any peripheral block in the device and route it to any other peripheral to initiate or affect an operation at the destination peripheral block.

Triggers come in two types:

- **High active, level-sensitive triggers.** This type typically reflects a peripheral state. For example, `tr_fifo_empty` indicates that a FIFO is empty. The trigger remains '1' as long as the FIFO is empty. This type requires an action by the consumer of the trigger for the producer to deactivate the trigger. For example, `tr_fifo_empty` is deactivated by writing a data element

to the associated FIFO. This trigger type can be produced on any clock.

- **Rising edge, pulse triggers.** This type typically reflects the occurrence of an event. For example, `tr_adc_done` indicates that a SAR ADC has converted a sample. This trigger type is produced on the peripheral system interface clock: the trigger remains '1' for a two-cycle pulse on the peripheral system interface clock, and returns to '0' by itself.

If the consumer of the trigger cannot immediately react to the trigger, it needs to be able to remember that the trigger occurred. If the consumer of the trigger is confronted with multiple triggers in short succession, it may need to remember multiple triggers or decide to miss triggers.

Triggers should be treated as asynchronous signals between producer and consumer peripheral: the consumer peripheral should synchronize input triggers. In addition, for pulse triggers, the consumer peripheral may perform rising edge detection and memorize occurrence of the trigger. Treating triggers as asynchronous signals, eases timing closure (similar to DSI signals, triggers may travel a large distance).

At a high-level abstraction, a trigger is a wire connection between a producer and a consumer peripheral. However, at a more detailed level, several processing steps are distinguished. From a platform perspective, it is important these processing steps are implemented consistently. The following text elaborates on these processing steps for input triggers: trigger multiplexing, synchronization, edge detection, and storage. It also elaborates the processing step for output triggers.

29.3 Trigger Multiplexing

The trigger component multiplexes trigger signals. The trigger input signals are typically peripheral output signals. The trigger output signals are typically peripheral input signals. Examples of trigger input signals are:

- TCPWM output signals; for example, counter reaches a pre-programmed limit.
- ADC output signals; for example, an ADC conversion has completed.
- I/O input signals.
- P-/M-DMA controller output signals that indicate the completion of a transfer.

Examples of trigger output signals are:

- TCPWM input signals; for example, "start a counter".
- ADC input signals; for example, "start an ADC conversion".
- I/O output signals.
- P-/M-DMA controller input signals to start a transfer.

In general, a trigger input signal indicates the completion of a peripheral action or a peripheral event. In general, a trigger output signal initiates a peripheral action.

The decision for standard multiplexer components (as part of peripheral groups), rather than multiplexer components integrated as part of peripherals, has the following rationale:

- A standard multiplexer component enforces a consistent user interface. Selection of a trigger input signal for a specific trigger output signal (peripheral input signal) is the same for all components.
- Any additional functionality provided by the multiplexer components, such as software control over output trigger signal activation, benefits all components.

The trigger module provides multiplexing functionality. It may be required to perform any of the following functions in a peripheral that uses the trigger output signals:


- Synchronization of the trigger signal to the peripheral clock domain.
- Edge detection on the trigger signal.
- Storing/remembering the trigger signal.

A trigger component consists of multiple trigger groups. A trigger group can be of two types:

- A one-to-one-based connectivity group. This group type connects a peripheral input trigger to one specific peripheral output trigger.
- A multiplexer-based connectivity group. This type connects a peripheral input trigger to multiple peripheral output triggers. The selection is under software control: `PERI_TR_GR_TR_CTL.TR_SEL`.

The trigger component can provide up to 16 multiplexer-based trigger groups and up to 16 one-to-one trigger groups.

Figure 29-2. Trigger Configuration Parameters



Each group is associated with the trigger inputs of a specific peripheral. Figure 29-2 gives an overview.

Peripheral output triggers are connected to trigger component input triggers TRIG_IN_MUX_x and TRIG_IN_1TO1_x. Peripheral input triggers are connected to trigger component output triggers TRIG_OUT_MUX_x and TRIG_OUT_1TO1_x. It is important to distinguish the functionality that is provided by the trigger component in PERI and the functionality that is provided by the peripheral. The trigger component provides the following functionality (on TRIG_IN_MUX_x and TRIG_IN_1TO1_x):

- For a multiplexer-based connectivity group: an input trigger TRIG_IN_MUX_x is selected for each output trigger TRIG_OUT_MUX_x. Note that all output triggers in a group i share the same input triggers. For a one-to-one based group, an input trigger TRIG_IN_1TO1_x is connected to output trigger TRIG_OUT_1TO1_x.
- Software control is provided for trigger activation. This control allows for activation of a specific trigger. For level-sensitive triggers, the trigger activation is completely under software control. For edge-sensitive triggers, the trigger activation is two high/'1' cycles on CLK_PERI.
- Trigger propagation can be blocked in debug mode (typically when a CPU is halted). This allows it to isolate the trigger consumer peripheral from getting input triggers. The debug mode is indicated by the level trigger input tr_debug_freeze, which is connected to a CPUSS CTI trigger output, sys.tr_cti_out.
- Hardware edge-detection is provided to allow pulse triggers that transfer to the output trigger clock domain. This hardware performs asynchronous edge detection to support input triggers that operate on a higher clock frequency than the output trigger clock domain. This functionality is intended for pulse triggers (level triggers typically bypass the edge detection functionality).
- Hardware trigger signal level inversion.

- Most trigger multiplexers have all output signals connected to a common peripheral. The manipulation logic is tied to the clock for that peripheral. However, the debug multiplexer has outputs in many clock domains. This is valid because most destinations are level-sensitive signals, I/Os, or some other debug destinations that might not need any clock manipulations.

A peripheral may provide the following functionality:

- Synchronization of the output triggers TRIG_OUT_MUX_x and TRIG_OUT_1TO1_x.
- Edge detection of the synchronized output triggers.
- Storing/remembering the trigger signal.

Figure 29-3. Trigger Group





Figure 29-4. Trigger One-to-One Group




Note that a one-to-one group has AND-gate functionality to disable an input trigger.

29.4 Trigger Functionality

The following figure gives an overview of a multiplexer group.

Figure 29-5. Multiplexer Trigger Group



Group trigger configuration:

PERI_TR_GR[Group Number].PERI_TR_GR_TR_CTL[Trigger Number].TR_SEL = Input trigger to be connected to

PERI_TR_GR[Group Number].PERI_TR_GR_TR_CTL[Trigger Number].TR_INV = Invert the trigger or not

PERI_TR_GR[Group Number].PERI_TR_GR_TR_CTL[Trigger Number].TR_EDGE = Edge or level-sensitive trigger

Activating the group trigger via software:

PERI_TR_CMD.TR_SEL = Input of Trigger Number


PERI_TR_CMD.GROUP_SEL = Trigger Number Group_Nr

PERI_TR_CMD.TR_EDGE = Edge or level-sensitive trigger

PERI_TR_CMD.OUT_SEL = 1 for output trigger and 0 for input trigger

PERI_TR_CMD.ACTIVATE = 1 signifies configured trigger activation


Figure 29-6. Example of Group Trigger



The following figure gives an overview of a one-to-one group.

Figure 29-7. One-to-One Trigger Group

Single trigger within a 1-to-1 group



1TO1 trigger configuration:

PERI_TR_1TO1_GR[Group Number].PERI_TR_1TO1_GR_TR_CTL[Trigger Number].TR_SEL = True (input trigger) or False (constant signal level '0')

PERI_TR_1TO1_GR[Group Number].PERI_TR_1TO1_GR_TR_CTL[Trigger Number].TR_INV = Invert the output trigger or not

PERI_TR_1TO1_GR[Group Number].PERI_TR_1TO1_GR_TR_CTL[Trigger Number].TR_EDGE = Edge or level-sensitive trigger


A trigger group consists of multiple trigger multiplexers with associated trigger manipulation logic. All trigger multiplexers in a trigger group share the same number of input triggers.

Figure 29-8. Example of 1-to-1 Trigger



As mentioned, the trigger manipulation logic provides asynchronous edge detection logic. The trigger manipulation is detailed in the following figure.

Figure 29-9. Trigger Manipulation



In addition, there are I/Os (TRIG_IN[0:x]) that can be used to generate triggers. These inputs can be used to trigger TCPWM, SAR ADC, PERI, and CPU-CTI. For device-specific trigger mux assignments; refer to the device datasheet.

29.5 Registers

Symbol	Name	Description
PERI_TR_CMD	Trigger Command Register	This register provides software control over trigger activation. This is useful for software-initiated triggers (such as P-/M-DMA transfers) or for debugging purposes. The control enables software activation of one specific input trigger or output trigger of the trigger multiplexer structure.
PERI_TR_GR_TR_CTL	Trigger Group Control Register	Controls group trigger actions and connects a peripheral input trigger to multiple peripheral output triggers
PERI_TR_1TO1_GR_TR_CTL	Trigger 1-to-1 Group Control Register	Controls the one-to-one trigger actions and connects a peripheral input trigger to a specific peripheral output trigger.

30. FlexRay Controller



30.1 Overview

The Traveo II FlexRay controller supports communication based on the FlexRay protocol V2.1 Rev. A. The maximum bit rate is 10 Mbit/s. Message buffers with a length of up to 254 data bytes are configurable. The message storage area consists of a single-port message RAM that holds up to 128 message buffers.

The message handler supports the following message processing functions:

- Acceptance filtering
- Transfer of messages between message RAM and channel protocol controllers
- Transmission schedule management
- Providing message status information

Features of the FlexRay Controller


- Based on the E-Ray block
- Complies with the FlexRay Protocol Specifications V2.1 Rev. A
- Up to 10 Mbit/s bit rate on each channel
- Configures up to 128 message buffers
- 8-Kbyte message RAM (equivalent to the following storage capacity)
 - 128 message buffers with maximum 48-byte data section
 - 30 message buffers with maximum 254-byte data section
- Variable-length message buffer configuration
- One configurable reception FIFO
- Each message buffer is configurable as a reception buffer, as a transmission buffer, or as part of the reception FIFO
- Host access to message buffers via input and output buffers
 - Input buffer: Holds messages to be transferred to the message RAM
 - Output buffer: Holds messages read from the message RAM
- Slot counter, cycle counter, and channel filtering
- Maskable interrupts
- Stop watch trigger input
- Timer trigger output
- Trigger interface for input/output buffer access by DMA

30.2 Configuration

30.2.1 Block Diagram

Figure 30-1 shows the block diagram of the Traveo II FlexRay controller.

Figure 30-1. FlexRay Controller Block Diagram



CPU I/F(CIF)

This interface connects the host CPU to the FlexRay controller.

Input Buffer (IBF)

This buffer is used to write to the message buffers configured in the message RAM. The host CPU can write header and data sections for a specific message buffer to the input buffer. The message handler transfers data from the input buffer to the selected message buffer in message RAM.

Output Buffer (OBF)

This buffer is used to read the message buffers configured in the message RAM. The message handler transfers data from the selected message buffer to the output buffer. When the data transfer is complete, the host CPU can read the header and data sections of the transferred message buffer from the output buffer.

Message Handler (MHD)

The message handler controls the data transfers between the following components:

- Input/output buffer and message RAM
- Transient buffer RAMs of the two FlexRay protocol controllers and message RAM

Message RAM (MRAM)

The message RAM consists of a single-port RAM that stores configuration data (header and data) for up to 128 FlexRay message buffers.

Transient Buffer RAM (TBF A/B)

This RAM stores the data sections of two messages.

FlexRay Channel Protocol Controller (PRT A/B)

The FlexRay channel protocol controller consists of a shift register and the FlexRay protocol finite state machine (FSM).

The protocol controller provides the following functions:

- Checking and controlling bit timings
- Receiving and transmitting FlexRay frames and symbols
- Checking the header CRC
- Generating and checking frame CRC
- Connecting to the bus driver

In addition, this block connects to the following blocks:

- Physical layer (bus driver)
- Transient buffer RAM
- Message handler
- Global time unit

- System universal control
- Frame and symbol processing
- Network management
- Interrupt control

Global Time Unit (GTU)

The global time unit provides the following functions:

- Microtick generation
- Macrotick generation
- Fault-tolerant clock synchronization using the Fault Tolerant Midpoint (FTM) algorithm
 - Rate correction
 - Offset correction
- Cycle counter
- Static segment timing control
- Dynamic segment (minislot) timing control
- Support for external clock correction

System Universal Control (SUC)

The system universal control controls the following functions:

- Configuration
- Wakeup
- Startup
- Normal operation
- Passive operation
- Monitor mode

Frame and Symbol Processing (FSP)

Frame and symbol processing controls the following functions:

- Ensuring the timing of frames and symbols is correct
- Testing the syntactic and semantic validity of received frames
- Setting the slot status flags.

Network Management (NEM)

The network management provides the following function:

- Handling the network management vector

Interrupt Control (INT)

The interrupt control provides the following functions:

- Provision of error and status interrupt flags
- Enabling and disabling interrupt factors
- Controlling the allocation of interrupt factors to the two module interrupt lines
- Enabling and disabling module interrupt lines
- Managing two interrupt timers
- Stop watch time capturing

30.3 Operations

This section describes the function of the FlexRay protocol. For more details, see the FlexRay Protocol Specifications V2.1 Rev. A.

30.3.1 Communication Cycle

The FlexRay communication cycle consists of the following elements and they are structured as in [Figure 30-2](#):


- Static segment
- Dynamic segment (optional)
- Symbol window (optional)
- Network idle time (NIT)

The network communication time (NCT) consists of a static segment, dynamic segment, and symbol window.

Starting at 1, the slot counter for each communication channel is incremented until it reaches the end of the dynamic segment. Also, both channels use the same synchronized macrotick.

[Figure 30-2](#) shows the structure of a communication cycle.

Figure 30-2. Communication Cycle Structure



The structure of the communication cycle can be configured by the FlexRay parameters. Static segment and NIT are always required.

30.3.1.1 Static Segment

The static segment has the following features.

- The bus guardian (if available) protects slots.
- Frame transmission begins at the action point of each static slot.
- The payload length is the same in all the frames of both channels.

The relevant parameters are as follows:

- Number of static slots, GTUC7:NSS[9:0]
- Static slot length, GTUC7:SSL[9:0]
- Static frame data length, MHDC:SFDL[6:0]
- Action point offset, GTUC9:APO[5:0]

30.3.1.2 Dynamic Segment

The dynamic segment has the following features.

- The bus guardian (even if available) is disabled, and all controllers have bus access.
- The slot length is variable and different in both channels.
- Transmission begins at a minislot action point.

The relevant parameters are as follows:

- Number of minislots, GTUC8:NMS[12:0]
- Minislot length, GTUC8:MSL[5:0]
- Minislot action point offset, GTUC9:MAPO[4:0]
- Start of latest transmission (last minislot), MHDC:SLT[12:0]

30.3.1.3 Symbol Window

The FlexRay Protocol Specifications V2.1 Rev. A defines several symbols.

Media access test symbol (MTS) is transmitted in the NORMAL_ACTIVE state to test the bus guardian.

MTS is one of such symbols and one MTS can be transmitted per channel during the symbol window process. The symbol window has the following features.

- One symbol is transmitted.
- MTS symbol transmission begins at a symbol window action point.

The relevant parameters are as follows:

- Action point offset, GTUC9:APO[5:0]
- Network idle time start, GTUC4:NIT[13:0]

30.3.1.4 Network Idle Time (NIT)

During the network idle time (NIT), the FlexRay controller performs the following tasks:

- Calculating the clock correction time (offset and rate)
- Performing offset correction over multiple macroticks after the start of offset correction
- Performing cluster cycle-related tasks

The relevant parameters are as follows:

- Network idle time start, GTUC4:NIT[13:0]
- Offset correction start, GTUC4:OCS[13:0]

30.3.1.5 Communication Cycle Configuration

The timing of FlexRay communication is managed by the time units microtick, macrotick (MT), and so on, as shown in [Figure 30-3](#) and [Figure 30-4](#) and described later. The italic names refer to FlexRay protocol parameters.

Figure 30-3. Hierarchy of the FlexRay Timing #1





Figure 30-4. Hierarchy of the FlexRay Timing #2



The length of one communication cycle is defined by the FlexRay parameter $gMacroPerCycle$ and assumed to be m . This value is set to GTUC2:MPC.

The static/dynamic segment starts at macrotick 0 and ends with macrotick n , which is defined as follows:

$$n = \text{static segment length} + \text{dynamic segment offset} + \text{dynamic segment length} - 1 \text{ MT}$$

$$n = gNumberOfStaticSlots * gdStaticSlot + \text{dynamic segment offset} + gNumberOfMinislots * gdMinislot - 1 \text{ MT}$$

The static segment length is set in GTUC7:SSL and GTUC7:NSS. The dynamic segment length is set in GTUC8:MSL and GTUC8:NMS.

The dynamic segment offset is calculated as follows:

if $gdActionPointOffset \leq gdMinislotActionPointOffset$:

$$\text{Dynamic segment offset} = 0 \text{ MT}$$

else if $gdActionPointOffset > gdMinislotActionPointOffset$:

$$\text{Dynamic segment offset} = gdActionPointOffset - gdMinislotActionPointOffset$$

The NIT starts at macrotick $k+1$ and ends with the last macrotick $m-1$. It should be configured as follows:


$$\text{GTUC4:NIT} = k$$

Also, the offset correction start needs to be configured to satisfy the following condition:

$$\text{GTUC4:OCS} \geq \text{GTUC4:NIT} + 1 = k + 1$$

The symbol window length results from the number of macroticks between the end of the static/dynamic segment; the NIT start and can be calculated as $k-n$.

Figure 30-5. Configuration of NIT Start and Offset Correction Start



30.3.2 Clock Synchronization

FlexRay communication is managed with distributed clock synchronization in the network.

Each FlexRay node synchronizes to the cluster by measuring the reception timing of synchronization frames from other nodes.

30.3.2.1 Global Time

Each node operates according to the concept of global time, although the node has its own clock. Global time consists of a vector with two values: cycle (cycle counter) and cycle time (macrotock counter).

- Macrotick (MT) = Base unit of FlexRay network time measurement, a macrotick is an integral multiple of microticks (μT).
- Cycle Length = the duration of one communication cycle expressed in macroticks (MT).

30.3.2.2 Local Time

Internally, nodes time their behavior with microtick resolution. Microticks are time units derived from the oscillator clock tick of the specific node. Therefore, microticks are controller-specific units. They may have different duration in different controllers. The precision of a node's local time difference measurements is a microtick (μT).

- Microtick generation: System clock 'prescaler' Microtick (μT)
- μT = Base unit of time measurement in the FlexRay controller, clock is corrected in units of μT .
- Cycle counter + macrotick counter = Node's local view of the global time

30.3.2.3 Synchronization Process

Synchronization frames are used as a means of clock synchronization. Only the synchronization nodes configured in advance can transmit synchronization frames. A synchronization node in a two-channel cluster will need to transmit a synchronization frame on both channels.

FlexRay has the following restrictions for synchronization.

- There is a maximum of one synchronization frame per node in one communication cycle.
- There is a maximum of 15 synchronization frames per cluster in one communication cycle.
- Every node must use a preconfigured number of synchronization frames (GTUC2:SNM[3:0]) for clock synchronization.
- A minimum of two synchronization nodes are required for clock synchronization and startup.

The time deviation between the expected reception time and observed reception time of a synchronization frame received during the static segment period is measured for clock synchronization. The correction time is calculated with the FTM algorithm during the NIT period (offset correction: all cycles; rate correction: odd-numbered cycles).

For details, see the FlexRay Protocol Specifications V2.1 Rev. A.

- Offset (Phase) correction
 - Only the deviation values measured and stored in the current cycle are used
 - For a node with two channels, the smaller value will be taken
 - This correction is calculated during the NIT of all communication cycles.

- The offset correction value calculated in an even-numbered cycle is used only to check for errors.
- The correction value is checked against limit values.
- The correction value is a signed integer number of μT .
- The correction is performed in odd-numbered cycles. The offset correction is distributed over each macrotick from the offset correction start to the end of the cycle (end of NIT), to shift the next start of the node's communication cycle (MTs are shortened or lengthened)

■ Rate (Frequency) correction

- Pairs of deviation values measured and stored in even-numbered and odd-numbered cycles are used.
- For a node with two channels, the average of the differences between the two channels is used.
- This correction is calculated during the NIT of odd-numbered communication cycles.
- Cluster drift damping is performed using the global damping value.
- The correction value is checked against limit values.
- The correction value is a signed integer number of μT .
- The correction is distributed over the macroticks comprising the next even/odd-numbered cycle pair (MTs are shortened or lengthened).

■ Sync Frame Transmission

Sync frames can be transmitted only from buffer 0 and 1. Message buffer 1 is used to transmit a sync frame when the sync frame should have different payloads on the two channels. In this case, MRC:SPLM bit must be set to '1'.

The message buffer used to transmit sync frames must be configured with the key slot ID, which can be set only in the DEFAULT_CONFIG or CONFIG state.

Nodes that transmit sync frames have the SUCC1:TXSY setting of '1'

30.3.2.4 External Clock Synchronization

There may be significant drifting in independent clusters during normal operation. For a required synchronization operation within an independent cluster, external clock synchronization is necessary even though the nodes within each cluster are synchronized. The offset correction time and rate correction time for the cluster are inferred by the host, enabling the operation to be accomplished.

- The external offset/rate correction value is a signed integer.
- The external offset/rate correction value is added to the calculated offset/rate correction value.
- The total offset/rate correction time (external and internal) is not checked against the set limit value.

30.3.3 Error handling

The error handling implemented in FlexRay is intended to ensure that communication between unaffected nodes is guaranteed during periods where nodes have a lower layer protocol error. In some cases, an operation to restart normal operation of the FlexRay controller must be implemented in an application program.

EIR:PEMC is set to '1' during a transition of the error handling state. Then, an interrupt is generated if the interrupt is enabled. CCEV:ERRM[7:6] indicates the actual error mode. Table 30-1 shows the error handling state during communication.

Table 30-1. POC Error Mode ^a

Error Mode	Description
ACTIVE	Full operation State: NORMAL_ACTIVE The FlexRay controller is completely synchronized and supports clock synchronization in the entire cluster. The error status and status change information can be obtained by reading the error interrupt flag and status interrupt flag from the EIR register and SIR register, respectively. An interrupt is generated if the interrupt is enabled.
PASSIVE	Limited operation State: NORMAL_PASSIVE FlexRay controller self-recovery is available. The FlexRay controller stops transmitting frames and symbols but can process received frames. Clock synchronization continues based on the received frames, but no active contribution to the clock synchronization for the entire cluster. The error status and status change information can be obtained from reading the error interrupt flag and status interrupt flag from the EIR register and SIR register, respectively. An interrupt is generated if the interrupt is enabled.
COMM_HALT	Operation stopped State: HALT FlexRay controller self-recovery is unavailable. The FlexRay controller stops frame and symbol processing, clock synchronization processing, and macrotick generation. The error status and status change information can be obtained by reading the error interrupt flag and status interrupt flag from the EIR register and SIR register, respectively. The bus driver is stopped.

a. POC: Protocol Operation Control

30.3.3.1 Clock Correction Failure Counter

The NORMAL_ACTIVE state transitions to the NORMAL_PASSIVE state when the clock correction failure counter reaches the maximum PASSIVE transition time without clock correction, SUCC3:WCP[3:0]. The NORMAL_ACTIVE/NORMAL_PASSIVE state transitions to the HALT state when the counter reaches the maximum HALT transition time without clock correction, SUCC3:WCF[7:4].

After passing the startup phase, the clock correction failure counter, CCEV:CCFC[3:0], can be used to monitor the periods during which the node clock correction time cannot be calculated. If either the missing offset correction signal, SFS:MOCS, or the missing rate correction signal, SFS:MRCS, is set to '1', the clock correction failure counter is incremented at the end of an odd-numbered communication cycle. If neither the missing offset correction signal, SFS:MOCS, nor the missing rate correction signal, SFS:MRCS, is set to '1', the clock correction failure counter is set to '0' at the end of an odd-numbered communication cycle.

The clock correction failure counter stops incrementing when it reaches the maximum HALT transition time without clock correction, SUCC3:WCF[7:4]. (In other words, the counter does not return to 0 even when incremented at its maximum value.) The clock correction failure counter is set to '0' when the CONFIG state transitions to the READY or NORMAL_ACTIVE state.

Note: There is no transition to the HALT state when SUCC1:HCSE has not been set.

30.3.3.2 Counter for State Transition from Passive to Active

The passive to active counter, CCEV:PTAC[12:8], controls the transition of POC from the NORMAL_PASSIVE state to the NORMAL_ACTIVE state. SUCC1:PTA[20:16] defines the number of consecutive even/odd-numbered cycle pairs that must have the valid clock correction time before transition from the NORMAL_PASSIVE state to the NORMAL_ACTIVE state is allowed. If SUCC1:PTA[20:16] is set to '0', the transition from the NORMAL_PASSIVE state to the NORMAL_ACTIVE state is not possible.

30.3.3.3 HALT Command

The FlexRay communication of the local node can be stopped by setting SUCC1:CMD[3:0] to 0110 (CHI command HALT).

If the command is executed in the NORMAL_ACTIVE or NORMAL_PASSIVE state, POC transitions to the HALT state at the end of the current cycle. If it is executed in any other state, SUCC1:CMD[3:0] becomes 0000, which is command_not_accepted, and EIR:CNA is set to '1'. An interrupt is generated if the interrupt is enabled.

30.3.3.4 FREEZE Command

The setting where SUCC1:CMD[3:0] is 0111 (CHI command FREEZE) enables a transition to the HALT state when the

host detects a severe error state. This command triggers a transition to the HALT state regardless of the current POC state.


The POC state from which the transition to the HALT state happened can be read from CCSV:PSL[5:0].

30.4 Communication Controller States

30.4.1 Communication Controller State

Figure 30-6 and Table 30-2 show the state transitions of the FlexRay communication controller.

Figure 30-6. State Diagram of the Communication Controller (CC)



State transitions are controlled by software resets by the host, external pins (RXDA/RXDB), the POC state machine, and the CHI command vector, SUCC1:CMD[3:0].

If the SUCC1:CMD[3:0] setting is 0111 (CHI command FREEZE), all states transition to the HALT state.

Table 30-2. FlexRay Controller State Transitions

Tn	State	From	To
1	■ Hard reset	All states	DEFAULT_CONFIG
2	■ CONFIG command Set SUCC1:CMD[3:0] = 0001 (CHI command CONFIG)	DEFAULT_CONFIG	CONFIG
3	■ Unlock sequence (followed by MONITOR_MODE command) Set SUCC1:CMD[3:0] = 1011 (CHI command MONITOR_MODE)	CONFIG	MONITOR_MODE
4	■ CONFIG command Set SUCC1:CMD[3:0] = 0001 (CHI command CONFIG)	MONITOR_MODE	CONFIG
5	■ Unlock sequence (followed by READY command) Set SUCC1:CMD[3:0] = 0010 (CHI command READY)	CONFIG	READY
6	■ CONFIG command Set SUCC1:CMD[3:0] = 0001 (CHI command CONFIG)	READY	CONFIG
7	■ WAKEUP command Set SUCC1:CMD[3:0] = 0011 (CHI command WAKEUP)	READY	WAKEUP
8	■ Wakeup pattern transmitted ■ WUP received ■ Frame header received ■ Wakeup collision occurred ■ READY command Set SUCC1:CMD[3:0] = 0010 (CHI command READY)	WAKEUP	READY
9	■ RUN command Set SUCC1:CMD[3:0] = 0100 (CHI command RUN)	READY	STARTUP
10	■ Startup successful	STARTUP	NORMAL_ACTIVE
11	■ Clock correction failure counter reached set value of SUCC3:WCP[3:0]	NORMAL_ACTIVE	NORMAL_PASSIVE
12	■ Number of valid cycle pairs in terms of clock correction time reached set value of SUCC1:PTA[4:0]	NORMAL_PASSIVE	NORMAL_ACTIVE
13	■ READY command Set SUCC1:CMD[3:0] = 0010 (CHI command READY)	STARTUP NORMAL_ACTIVE NORMAL_PASSIVE	READY
14	■ SUCC1:HCSE set to '1' after clock correction failure counter reached set value of SUCC3:WCF[3:0] ■ HALT command Set SUCC1:CMD[3:0] = 0110 (command HALT)	NORMAL_ACTIVE	HALT
15	■ SUCC1:HCSE set to '1' after clock correction failure counter reached set value of SUCC3:WCF[3:0] ■ HALT command Set SUCC1:CMD[3:0] = 0110 (command HALT)	NORMAL_PASSIVE	HALT
16	■ FREEZE command Set SUCC1:CMD[3:0] = 0111 (CHI command FREEZE)	All states	HALT
17	■ CONFIG command Set SUCC1:CMD[3:0] = 0001 (CHI command CONFIG)	HALT	DEFAULT_CONFIG

30.4.2 DEFAULT_CONFIG State

The FlexRay controller is stopped in the DEFAULT_CONFIG state. All configuration registers are accessible (if CTL.ENABLE is set to '1'; see [30.15.1 Enable/Disable FlexRay Controller](#)). The pins to the physical layer are inactive.

A transition to this state occurs in the following cases:

- When a hard reset is performed
- When there is a transition from the HALT state

To transition from the DEFAULT_CONFIG state to the CONFIG state, write '0001' in SUCC1:CMD[3:0].

30.4.3 CONFIG state

The FlexRay controller is stopped in the CONFIG state. All configuration registers are accessible, but the pins to the physical layer are inactive. This state is used to initialize the FlexRay controller settings.

A transition to this state occurs in the following cases:

- When there is a transition from the DEFAULT_CONFIG state
- When there is a transition from the MONITOR_MODE or READY state

If there is a transition to this state via the HALT and DEFAULT_CONFIG states, status information and settings can be analyzed. Confirm that all settings are correct before a transition from the CONFIG state.

A transition from the CONFIG state requires execution of the unlock sequence with Lock register (LCK register). After unlocking the CONFIG state, write to SUCC1:CMD for a transition to the next state.

30.4.4 MONITOR_MODE

After unlocking the CONFIG state and writing SUCC1:CMD to '1011', there is a transition to MONITOR_MODE. FlexRay frames and wakeup patterns can be received in this mode. The integrity of the time of the received frames is not checked. Consequently, cycle counter filtering is not supported. This mode can be used for debugging purposes. After writing SUCC1:CMD to '0001', there is a transition to the CONFIG state.

In MONITOR_MODE, the "pick first" mechanism is disabled. This means that a receiving message buffer may only be created to receive on one channel. A received frame is stored in the message buffer according to the frame ID and receive channel. Null frames are handled in the same way as data frames. Only the MBS:VFRA, MBS:VFRB, MBS:MLST, MBS:RCIS, MBS:SFIS, MBS:SYNS, MBS:NFIS, MBS:PPIS, and MBS:RESS status bit values become valid after a frame is received.

In MONITOR_MODE, the CAS and MTS symbols cannot be distinguished from each other. If either of these symbols are detected on the channel, either SIR:MTSA or SIR:MTSB is set. SIR:CAS does not function in MONITOR_MODE.

30.4.5 READY State

After unlocking the CONFIG state and writing SUCC1:CMD to '0010', there is a transition to the READY state. The following is possible: cluster wakeup through a transition from this state to the WAKEUP state, a coldstart through a transition from this state to the STARTUP state, or integration into the running cluster from this state.

Each of the following states transition to the READY state by writing SUCC1:CMD to '0010' (CHI command READY):

- CONFIG state
- WAKEUP state
- STARTUP state
- NORMAL_ACTIVE state
- NORMAL_PASSIVE state

The READY state transitions to the following states:

- To the CONFIG state by writing SUCC1:CMD to 0001 (CHI command CONFIG)
- To the WAKEUP state by writing SUCC1:CMD to 0011 (CHI command WAKEUP)
- To the STARTUP state by writing SUCC1:CMD to 0100 (CHI command RUN)

Note: The transition of POC from the READY state to the STARTUP state does not affect status bits (MHDS[14:0]), registers (TXRQ1/2/3/4), and status data stored in the message RAM.

30.4.6 WAKEUP State

This section describes the wakeup settings for the FlexRay controller. [Table 30-3](#) and [Figure 30-7](#) describe the state transition for WAKEUP.

This state is entered from READY state by writing SUCC1:CMD[3:0] to '0011' (CHI command WAKEUP).

The WAKEUP state exits to the READY state under the following conditions:

- Writing SUCC1:CMD[3:0] to '0010' (CHI command READY).
- After transmission of a wakeup pattern is completed
- After a WUP is received
- After a WUP collision is detected
- After a frame header is received
- When SUCC1:CMD[3:0] of "0010" is written (CHI command READY)

WAKEUP must be executed in a cluster before communication startup to ensure that all nodes are awake. When receiving a wakeup pattern on a channel, the bus driver wakes up the other components of the node. At least one node in a cluster needs an external wakeup source.

The host controls the wakeup procedure. It is informed about the cluster state by the bus driver and FlexRay controller. The host configures the FlexRay controller (and bus guardian if available) and wakes up the cluster. Configuring the FlexRay controller enables transmission of a special wakeup pattern on each available channel separately. The FlexRay controller needs to recognize the wakeup pattern only during the WAKEUP state.

WAKEUP can be executed on only one channel at a time. The wakeup channel must be configured by writing SUCC1:WUCS during the CONFIG state. The FlexRay controller ensures that the communication on this channel is not disturbed. It is impossible to confirm the wakeup of all

nodes connected to this channel before the startup phase has completed. The described procedure allows one channel node in a two-channel system to trigger the wakeup. Any coldstart node can then initiate the wakeup on the other channel.

The wakeup procedure supports any number of nodes trying to send a wakeup pattern simultaneously and resolves this situation so that only one node sends the pattern. Even in the case of a fault causing two nodes to send the wakeup pattern, the pattern is collision-resilient and the resulting signal is capable to wake other nodes.

After a transition to the READY state following WAKEUP, the FlexRay controller reports the change in the WAKEUP status by setting the SIR:WST flag to '1'. The WAKEUP status vector can be read from CCSV:WSV[2:0]. If the received wakeup pattern is valid, either the SIR:WUPA or SIR:WUPB flag is set to '1'.

Figure 30-7. POC Configuration in WAKEUP State




Table 30-3. WAKEUP State Transition

Tn	State	From	To
Enter	- WAKEUP command - Set SUCC1:CMD[3:0] = 0011 (CHI command WAKEUP)	READY	WAKEUP
1	- WAKEUP command - Set SUCC1:CMD[3:0] = 0011 (CHI command WAKEUP)	WAKEUP_STANDBY	WAKEUP_LISTEN
2	- A WUP is received on the wakeup channel selected by SUCC1:WUCS. - A frame header is received on either of the valid channels.	WAKEUP_LISTEN	WAKEUP_STANDBY
3	A timer event has occurred.	WAKEUP_LISTEN	WAKEUP_SEND
4	Transmission of a wakeup pattern has completed normally.	WAKEUP_SEND	WAKEUP_STANDBY
5	A collision is detected.	WAKEUP_SEND	WAKEUP_DETECT
6	- The wakeup timer has timed out. - A WUP is received on the wakeup channel selected by SUCC1:WUCS. - A frame header is received on either of the valid channels.	WAKEUP_DETECT	WAKEUP_STANDBY
Exit	- Wakeup has completed (after T2, T4, and T6). - READY command - SetSUCC1:CMD[3:0] = 0010 (CHI command READY) (There is a simultaneous reset to the WAKEUP_STANDBY state by this CHI command.)	WAKEUP	READY

The wakeup timer and wakeup noise timer control the WAKEUP_LISTEN state. These timers are controlled using the following parameters: listen timeout value, SUCC2:LT[20:0] and listen timeout noise value, SUCC2:LTN[3:0]. Listen timeout allows a quick cluster wakeup in a noise-free environment. Listen timeout noise allows wakeup in an environment with a high level of noise interference.

In the WAKEUP_SEND state, wakeup patterns are transmitted on the configured channel while checking for collisions. After a return from the WAKEUP state to the READY state, CMD[3:0] must be 0100 (CHI command RUN) for a transition to the STARTUP state.

In the WAKEUP_DETECT state, the cause of a wakeup collision detected in the WAKEUP_SEND state can be identified. The identification is stopped when the listen timeout setting in SUCC2:LT[20:0] is exceeded.

There is a direct transition to the READY state upon either detection of a wakeup pattern from another node or reception of a frame header. If neither detection nor reception happens, there is a transition from the WAKEUP_DETECT state after listen timeout elapses. In such cases, the cause of the wakeup collision is unknown.

The host must recognize possible failures during wakeup and take the necessary actions.

It is recommended that the startup is delayed for the minimum time until another coldstart node wakes up and is initialized. The FlexRay Protocol Specifications V2.1 Rev. A recommends waking up the two channels using two different FlexRay controllers.

30.4.6.1 Host Operations

The host must coordinate wakeup of the two channels and determine whether to wake up a specific channel. The host starts transmission of wakeup patterns. The bus drivers of the other nodes detect the wakeup patterns and notify their hosts.

The host controls the single channel wakeup procedure as follows:

- Configure the FlexRay controller in the CONFIG state.
 - Select the wakeup channel that is set by the SUCC1:WUCS bit.
- Check the bus driver to see whether a wakeup pattern (WUP) is received.
- Start the bus driver on the selected wakeup channel.
- Write SUCC1:CMD[3:0] as '0010' for a transition to the READY state.
- Write SUCC1:CMD[3:0] as '0011' to start waking up the set channel.
 - The FlexRay controller transitions to the WAKEUP state.
 - After completing a wakeup, the FlexRay controller transitions to the READY state and displays the wakeup status (CCSV:POCS[5:0]).
- Wait for a predetermined duration, until another node is woken up and configured.
- Perform the following procedure for a coldstart node.
 - Wait until a WUP has occurred on the other channel in a two-channel cluster configuration.

- Write SUCC1:CMD[3:0] as '1001' (CHI command ALLOW_COLDSTART) to reset the coldstart inhibit flag, CCSV:CSI
- Write SUCC1:CMD[3:0] as '0100' (CHI command RUN) for a transition to the STARTUP state.

The bus driver triggers the following wakeup procedure.


- The bus driver identifies the wakeup.
- The bus driver notifies the host of the wakeup event.
- The host configures the FlexRay controller.
- If necessary, the host may send a wakeup command for the second channel and wait for the predefined duration so that other nodes can complete wake-up and configuration
- Write SUCC1:CMD[3:0] as '0100' (CHI command RUN) for a transition to the STARTUP state.

30.4.6.2 Wakeup Pattern (WUP)

A wakeup pattern (WUP) consists of at least two wakeup symbols (WUS). The wakeup symbol and wakeup pattern are set by the PRTC1 and PRTC2 registers, respectively.

- Single-channel wakeup, that is, wakeup symbol cannot be transmitted to both channels at the same time.
- Wakeup symbol is collision resilient for at least two sending nodes (that is, two overlapping symbols are always recognizable)
- The wakeup symbol configuration must be identical in all nodes of the cluster.
- The low time of the wakeup symbol is set by PRTC2:TXL[5:0].
- The wakeup symbol idle time, which is used to listen to activity on the bus, is set by PRTC2:TXI[7:0].
- The wakeup pattern consists of at least two transmission wakeup symbols necessary for wakeup.
- The number of repetitions (between 2 and 63 repetitions) can be set by PRTC1:RWP[5:0].
- The length of the wakeup symbol reception window is set by PRTC1:RXW[8:0].
- The low time of wakeup reception is set by PRTC2:RXL[5:0].
- The wakeup reception idle phase time is set by PRTC2:RXI[5:0].

Figure 30-8. Wakeup Pattern



30.4.7 STARTUP State

Any coldstart node that enters STARTUP state should ensure that both channels are woken up before initiating the coldstart attempt.

The time required to complete wakeup and configuration cannot be assumed to be the same for all nodes and stars of a clusters. At least two nodes are required for the startup of cluster communications. Therefore, it is recommended to postpone the startup of the node on which a wakeup had been initiated for the duration of the minimum time required for another coldstart node to complete the wakeup, initial setting, and startup. The time delay due to the completion of wakeup and configuration of all nodes and clusters can exceed several 100 ms (depending on the hardware).

Startup is executed on all channels simultaneously. During startup, nodes only transmit the startup frames that are both sync and null frames.

A fault-tolerant, distributed startup strategy is specified for initial synchronization of all nodes. Generally, nodes transition to the NORMAL_ACTIVE state through the following procedures (see [Figure 30-9](#)):

- Coldstart procedure to start schedule synchronization (leading coldstart node)
- Coldstart procedure for participation of other coldstart nodes (following coldstart node)
- Integration procedure for integration into the existing communication schedule (all other nodes)

A coldstart attempt starts with transmission of a collision avoidance symbol (CAS). Only the coldstart node that transmitted the CAS transmits frames during the first four cycles after the CAS. Afterwards, it is joined by the other coldstart nodes, and then all other nodes participate in the cluster.

The coldstart node transmits synchronization frames in the key slot by setting '1' in SUCC1:TXST and SUCC1:TXSY. (See SUCC1:TXST/TXSY bit). Message buffer 0 contains the key slot ID that defines the slot number in which the startup frame is transmitted. The startup frame indicator is set to '1' in the frame header of the startup frame.


A cluster consisting of three or more nodes are configured such that at least three nodes are coldstart nodes. A cluster consisting of two nodes must have both nodes configured as coldstart nodes. At least two fault-free coldstart nodes are required for startup of the cluster.

Each startup frame should also be a synchronization frame. Therefore, all coldstart nodes are also synchronization nodes. The number of coldstart attempts is set by SUCC1:CSA[4:0].

To integrate non-coldstart nodes into a cluster, at least two startup frames from other nodes are required. Integration of the non-coldstart nodes may begin before startup of the coldstart nodes is completed. However, startup of those non-coldstart nodes is never completed until startup of at least two coldstart nodes has completed.

Both non-coldstart nodes and coldstart nodes start passive integration when they receive sync frames from which they can obtain TDMA (time division multiple access) schedule information. During the integration, the nodes adjust their clock to the global clock (rate and offset) and their cycle time to the network global cycle. Later, these settings are checked for consistency with all the available network nodes. Nodes can actively participate in communication only when they have passed these checks.

Figure 30-9. State Diagram of the Time-triggered Startup



30.4.7.1 Coldstart Inhibit Mode

Nodes cannot initialize cluster communication in coldstart inhibit mode (CCSV:CSI is 1). This means that startup using the coldstart procedure is prohibited. Such nodes can either be integrated into the running cluster or transmit startup frames after another coldstart node has started initializing cluster communication.

The coldstart inhibit bit, CCSV:CSI, is set whenever POC enters the READY state. Clear this bit by setting SUCC1:CMD[3:0] to 1001 (CHI command ALLOW_COLDSTART).

30.4.7.2 Startup Timeouts

The FlexRay controller provides two μ T timers that support two timeout values (startup timeout and startup noise timeout). These timers start during transition to the COLDSTART_LISTEN state. When either of these timers ends, the node terminates the detection phase of the other node (a transition from the COLDSTART_LISTEN state to another state) to start communication.

Note: The startup timer and startup noise timer are the same as the wakeup timer and wakeup noise timer and use SUCC2:LT[20:0] and SUCC2:LTN[3:0], respectively.

■ Startup Timeout

The startup timeout limits the listening time used by a node to determine whether communication has been established between other nodes, or whether at least one coldstart node is requesting integration with other nodes. The startup timer is set by SUCC2:LT[20:0] (*pdListenTimeout*).

The startup timer is restarted under the following conditions.

- A transition to the COLDSTART_LISTEN state.
- Both channels reach the idle state in the COLDSTART_LISTEN state.

The startup timer is stopped under the following conditions.

- Communication is detected on one of the configured channels while in the COLDSTART_LISTEN state.
- A transition from the COLDSTART_LISTEN state to another state.

When the startup timer elapses, neither a timer overflow nor a periodic restart occurs. The timer status is retained for further processing by the startup state machine.

■ Startup Noise Timeout

The startup timer and startup noise timer start during a transition from the STARTUP_PREPARE state to the COLDSTART_LISTEN state. The startup noise timeout is used to improve the reliability of the startup procedure in an environment with noise.

The startup noise timeout is determined by SUCC2:LTN[3:0] as follows:

$$pdListenTimeout \times gListenNoise = SUCC2:LT[20:0] \times (SUCC2:LTN[3:0] + 1)$$

The startup noise timer is restarted under the following conditions.

- A transition to the COLDSTART_LISTEN state.
- A header or CAS symbol is correctly decoded while in the COLDSTART_LISTEN state.

The startup noise timer stops at a transition from the COLDSTART_LISTEN state to another state.

After the startup noise timer elapses, neither a timer overflow nor a periodic restart occurs. The timer status is retained for further processing by the startup state machine. Because this timer will not be restarted during random channel activity, it acts as a fall-back solution, which guarantees that a node will try to start up the cluster even in a noisy environment.

30.4.7.3 Startup Process of Leading Coldstart Node

A coldstart node in the COLDSTART_LISTEN state monitors the states of the connected channels.

If no communication is detected, the node transitions to the COLDSTART_COLLISION_RESOLUTION state and starts a coldstart. The CAS symbol is transmitted first and succeeded by the first regular cycle with number 0.

The node transmits its startup frame from cycle 0. Multiple coldstart nodes may attempt to initiate a coldstart and transmit CAS symbols simultaneously. This situation is resolved in the first four cycles after CAS transmission.

If any node that has started a coldstart receives a CAS symbol or a frame header within these four cycles, the node transitions again to the COLDSTART_LISTEN state. As a result, only one node within the cluster continues the coldstart procedure. The other coldstart nodes start transmitting their own startup frames in cycle 4.

After four cycles in the COLDSTART_COLLISION_RESOLUTION state the node that initiated the coldstart transitions to the COLDSTART_CONSISTENCY_CHECK state. This node collects all startup frames from cycle 4 and cycle 5, and corrects the clock. If the clock is corrected without an error and at least one valid startup frame pair is received, it transitions from the COLDSTART_CONSISTENCY_CHECK state to the NORMAL_ACTIVE state.

The number of coldstart attempts is set by SUCC1:CSA[4:0]. The remaining number of coldstart attempts can be read from CCSV:RCA[4:0]. The remaining number of coldstart attempts is decremented each time that an attempt is made. A transition to the COLDSTART_LISTEN state is possible when the remaining number of attempts is greater than 1. A transition to the

COLDSTART_COLLISION_RESOLUTION state is possible when the remaining number of attempts is greater than 0. Integration into a cluster is possible when the remaining number of coldstart attempts is 1, but a coldstart is prohibited.

30.4.7.4 Startup Process of Following Coldstart Node

When transitioning to the COLDSTART_LISTEN state, a coldstart node attempts to receive a valid pair of startup frames to obtain the cycle schedule and clock correction from the leading coldstart node.

Upon receiving the first valid startup frame, it immediately transitions to the INITIALIZE_SCHEDULE state. Upon receiving the second valid startup frame and obtaining the cycle schedule, it transitions to the INTEGRATION_COLDSTART_CHECK state.

In the INTEGRATION_COLDSTART_CHECK state, the following is guaranteed: clock correction can be done correctly and the leading coldstart node is still available for use. (The following coldstart node initializes its schedule according to this leading coldstart node.)

The following coldstart node collects all synchronization frames and corrects the clock in the following cycle pair. The node transitions to the COLDSTART_JOIN state when the clock correction does not show an error and the node continues to receive sufficient frames from the same node.

The following coldstart node starts transmitting its own startup frames in the COLDSTART_JOIN state and continues transmitting these frames in the next cycles. This enables the leading coldstart node and participating nodes to check whether their cycle schedules match with each other. If an error is detected in clock correction, the node aborts the integration attempt. If the node in this state receives at least one valid startup frame in an even-numbered cycle and at least one valid pair of startup frames in all cycle pairs, the node transitions from the COLDSTART_JOIN state to the NORMAL_ACTIVE state. Therefore, the following coldstart node transitions from the STARTUP state to NORMAL_ACTIVE state at least one cycle later than the leading coldstart node.

30.4.7.5 Startup Process of Non-coldstart Node

A non-coldstart node in the INTEGRATION_LISTEN state monitors the states of the connected channels.

Upon receiving the first valid startup frame, it immediately transitions to the INITIALIZE_SCHEDULE state. Upon receiving the second valid startup frame and obtaining the cycle schedule, it transitions to the INTEGRATION_CONSISTENCY_CHECK state.

In the INTEGRATION_CONSISTENCY_CHECK state, the non-coldstart node checks whether clock correction is

working correctly and whether enough coldstart nodes (at least two) are transmitting startup frames matching the node's cycle schedule. The integration is stopped when any error is detected while clock correction is in operation. This non-coldstart node needs to receive either two valid startup frames or a valid startup frame from the node that this has integrated on within the first even-numbered cycle in this state. Otherwise, the node aborts the integration. This non-coldstart node needs to receive either two valid pairs of startup frames or a valid pair of startup frames from the node that this has integrated on within the first cycle pair in this state. Otherwise, the node aborts the integration.

If after the first double-cycle less than two valid startup frames are received within an even cycle, or less than two valid startup frame pairs are received within a double-cycle, the startup attempt is aborted.

Nodes in this state need to see two valid startup frame pairs for two consecutive double-cycles each to be allowed to leave STARTUP and enter NORMAL_OPERATION. Consequently, they leave startup at least one double-cycle after the node that initiated the coldstart and only at the end of a cycle with an odd cycle number.

30.4.8 NORMAL_ACTIVE State

The startup phase of the entire cluster ends immediately after the transition of the node that transmitted the first CAS symbol and an additional node to the NORMAL_ACTIVE state. The transmission timing of all transmission messages is scheduled in the NORMAL_ACTIVE state. This includes all data frames in the same way as synchronization frames. Rate and offset measurement begins in all even-numbered cycles. (Even/odd-numbered cycle pairs are required.)

The FlexRay controller supports the normal communication functions in the NORMAL_ACTIVE state.

- Transmission and reception on the FlexRay bus are performed according to settings.
- Clock synchronization is in operation.

The FlexRay controller transitions from the NORMAL_ACTIVE state to the following states:

- To the HALT state after the end of the current cycle by writing SUCC1:CMD[3:0] as 0110 (CHI command HALT)
- Immediately to the HALT state by writing SUCC1:CMD[3:0] as 0111 (CHI command FREEZE)
- To the HALT state through an error state change from ACTIVE to COMM_HALT
- To the NORMAL_PASSIVE state through an error state change from ACTIVE to PASSIVE
- To the READY state by writing SUCC1:CMD[3:0] as 0010 (CHI command READY)

30.4.9 NORMAL_PASSIVE State

The NORMAL_ACTIVE state transitions to the NORMAL_PASSIVE state when the error state changes from ACTIVE to PASSIVE.

Nodes can receive all frames in the NORMAL_PASSIVE state. (Nodes are completely synchronized, and clock synchronization is possible.) However, nodes do not actively participate in communication, as compared with the NORMAL_ACTIVE state. This means that neither symbols nor frames are transmitted.

The following operations are performed in the NORMAL_PASSIVE state.

- Frames are received on the FlexRay bus.
- Neither frames nor symbols are transmitted to the FlexRay bus.
- Clock synchronization is in operation.

The FlexRay controller transitions from the NORMAL_PASSIVE state to the following states:

- To the HALT state after the end of the current cycle by writing SUCC1:CMD[3:0] as 0110 (CHI command HALT)
- Immediately to the HALT state by writing SUCC1:CMD[3:0] as 0111 (CHI command FREEZE)
- To the HALT state through an error state change from PASSIVE to COMM_HALT
- To NORMAL_PASSIVE through an error state change from PASSIVE to ACTIVE (This error state change occurs when CCEV:PTAC[4:0] equals SUCC1:PTA[4:0]-1.)
- To the READY state by writing SUCC1:CMD[3:0] as 0010 (CHI command READY)

30.4.10 HALT State

All communications (transmission and reception) are halted in this state.

The FlexRay controller transitions to the HALT state in the following cases:

- Transition from the NORMAL_ACTIVE or NORMAL_PASSIVE state when SUCC1:CMD[3:0] of 0110 (CHI command HALT) is written
- Transition from all states when SUCC1:CMD[3:0] of 0111 (CHI command FREEZE) is written
- Transition from the NORMAL_ACTIVE state when the clock correction fatal counter reaches the maximum HALT transition time without clock correction, WCF[3:0]. SUCC1:HCSE is set to '1'.
- Transition from the NORMAL_PASSIVE state when the clock correction fatal counter reaches the maximum HALT transition time without clock correction, WCF[3:0]. SUCC1:HCSE is set to '1'.

The FlexRay controller transitions from this state to the DEFAULT_CONFIG state in the following case:

- When SUCC1:CMD[3:0] of 0001 (CHI command CONFIG) is written
- When SUCC1:CMD[3:0] of 0110 (CHI command HALT) is written, the FlexRay controller sets the CCSV:HRQ bit to '1' and transitions to the HALT state after the next cycle ends.
- When SUCC1:CMD[3:0] of 0111 (CHI command FREEZE) is written, the controller immediately transitions to the HALT state and sets the CCSV:FSI bit to '1'.

The state from which the transition to the HALT state took place can be read from CCSV:PSL[5:0].

30.5 Network Management

The accrued network management (NM) vector can be read from the NMV1 to NMV3 registers. The FlexRay controller performs a bit OR operation on all NM vectors in all valid received NM frames where the payload preamble indicator (PPI) is set. Only static frames can be configured as NM frames. After each cycle is completed, NM vectors are updated.

An NM vector length between 0 and 12 bytes can be set by NEMC:NML[3:0]. The NM vector length must be the same in all nodes of the cluster.

The PPIT bit in the header section of a transmission buffer must be set through WRHS1:PPIT to set the PPI bit in the corresponding frame. NM information must also be written in the data section of the transmission buffer. A mechanism for evaluating the NM vector must be implemented by the application.

Note: If the message buffer is configured to transmit/receive network management frames, the payload length configured in header 2 of the message buffer must be equal to or greater than the NM vector length configured in NEMC:NML[3:0].

The cycle count does not increase when the HALT state is passed, so the NM vectors are not updated. In this case, the NMV1 to NMV3 registers retain the values from the previous cycle.

30.6 Filtering and Masking

Filtering is done through a comparison of the settings of the current slot, cycle counter value, channel ID (channel A and channel B), and message buffer. If this comparison produces a match, the message buffer is updated or transmitted.

Filtering is applied to the following:

- Slot counter
- Cycle counter
- Channel ID

The following filter combinations can be used to filter the transmission or reception time:

- Slot counter + channel ID
- Slot counter + channel ID + cycle counter

To store received messages in the message buffer, all the set filters must be matched against information on the received messages.

These filtering and masking conditions are configured in the message header and set to each message buffer (see [30.11.1 Header Partition](#) for the details of the message header information).

Note: The FIFO rejection filter and FIFO rejection filter mask configure the FIFO acceptance filter.

A message is transmitted in the time slot corresponding to the set frame ID in the set channel. If cycle counter filtering is enabled, the configured cycle filter value must also match.

30.6.1 Slot Counter Filtering

All the transmission and reception buffers contain a frame ID in the header section. The frame ID is compared with the current slot counter value to assign it to the slot corresponding to the transmission or reception buffer.

If multiple buffers are configured with the same frame ID and same channel ID, and if the buffers have cycle counter filter values corresponding to the same slot, the message buffer with the smallest buffer number is used.

30.6.2 Cycle Counter Filtering

Cycle counter filtering is based on the concept of a cycle set. A match to a filter is detected when any one element of the cycle set matches. The cycle set is defined by the cycle code field in header section 1 in each message buffer.

When configuring message buffer 0 to store startup/synchronization frames or single-slot frames by setting the respective SUCC1:TXST, SUCC1:TXSY, and SUCC1:TSM bits, cycle counter filtering should be disabled in message buffer 0.

Note: Static time slot sharing by cycle counter filtering between different nodes in the FlexRay network is not permitted.

[Table 30-4](#) shows settings of the number of cycles belonging to a cycle set, and [Table 30-5](#) shows some examples of valid cycle sets used for cycle counter filtering.

Table 30-4. Cycle Set Definitions

Cycle Code	Match with Cycle Counter Values
0b000000x	All cycles
0b000001c	Every second cycle at (Cycle Count) mod2 = c
0b00001cc	Every fourth cycle at (Cycle Count) mod4 = cc
0b0001ccc	Every eighth cycle at (Cycle Count) mod8 = ccc
0b001cccc	Every sixteenth cycle at (Cycle Count) mod16 = cccc
0b01ccccc	Every thirty-second cycle at (Cycle Count) mod32 = ccccc
0b1cccccc	Every sixty-fourth cycle at (Cycle Count) mod64 = ccccccc

Table 30-5. Examples for Valid Cycle Set

Cycle Code	Match with Cycle Counter Values
0b0000011	1-3-5-7-...-63
0b0000100	0-4-8-12-...-60
0b0001110	6-14-22-30-...-62
0b0011000	8-24-40-56
0b0100011	3-35
0b1001001	9

A received message is stored only if the cycle counter value during the cycle in which the message was received matches an element of the cycle set of the reception buffer. Other filter criteria must also be satisfied.

The contents of the transmission buffer are transmitted to the set channel when an element of the cycle set matches the current cycle counter value. Other filter criteria must also be satisfied.

30.6.3 Channel ID Filtering

The header section of each message buffer in the message RAM has 2-bit channel filtering fields (CHA and CHB). They act as filters for the reception buffer and as control fields for the transmission buffer (see Table 30-6).

Table 30-6. Channel Filtering Settings

CHA	CHB	Transmission Buffer (Transmission Frame)	Reception Buffer (Storage of Received Frame)
1	1	On both channels (static segment only)	Received on channel A or B (store first semantically valid frame, static segment only)
1	0	On channel A	Received on channel A
0	1	On channel B	Received on channel B
0	0	No transmission	Ignore frame

The contents of the transmission buffer are transmitted to the channel specified by the channel filtering field when slot counter filtering and cycle counter filtering criteria are satisfied. However, setup (CHA and CHB settings) for transmission to both channels is permitted only in the static segment.

If the channel specified by the channel filtering field receives valid frames when slot counter filtering and cycle counter filtering criteria are satisfied, the frames are stored. However, setup (CHA and CHB settings) for frame reception by both channels is permitted only in the static segment.

Note: If the message buffer is configured for dynamic segments and both channel filtering field bits (CHA and CHB) are set to '1', frames are not transmitted, and received frames are ignored. (The function is the same as when CHA and CHB are '0'.)

30.6.4 FIFO Filtering

One rejection filter and one rejection filter mask are available for FIFO filtering. The rejection filter consists of the channel filter FRF:CH[1:0], frame ID filter FRF:FID[10:0], and cycle counter filter FRF:CYF[6:0]. The FRF and FRFM registers can be configured only in the DEFAULT_CONFIG or CONFIG state. The filter settings in the header section of a message buffer belonging to a FIFO group are ignored.

FRF:CYF[6:0] is a 7-bit cycle counter filter that specifies a cycle set and determines the communication cycle to which to apply the frame ID filter and channel filter. All frames are rejected in cycles not belonging to the configured cycle set.

A valid received frame is stored in the FIFO if channel ID, frame ID, and cycle counter are not rejected by the configured rejection filter and rejection filter mask, and if there is no matching dedicated receive buffer.

30.7 Transmission Procedure

30.7.1 Static Segment

If several messages are pending for transmission in the static segment, the message with the frame ID corresponding to the next transmission slot is selected as the next message to transmit. The data section of the transmission buffer assigned to the static segment can be updated by the end of the previous time slot. This means that message transmission from the input buffer must be started by writing to the input buffer command request register at this time.

30.7.2 Dynamic Segment

If several messages are pending for transmission in the dynamic segment, the message with the highest priority (the smallest frame ID) is selected as the next message to transmit. Also, different slot counter sequences may be possible for channel A and channel B in the dynamic segment (simultaneous transmission with different frame IDs on both channels).

The data section of the transmission buffer assigned to the dynamic segment can be updated by the end of the previous slot. This means that message transmission from the input buffer must be started by writing to the input buffer command request register at this time.

The start of latest transmission is configured by MHDC:SLT[12:0], which defines the maximum minislot value where a transmission can be started before frame transmission in the dynamic segment of the current cycle is prohibited.

30.7.3 Transmission Buffer

Each message buffer can be used as a transmission buffer when the CFG bit in the header section of the message buffer is set to '1' through WRHS1.

A transmission buffer can be assigned to a FlexRay controller channel in the following ways:

- Static segment:
 - Channel A or channel B
 - Channel A and channel B
- Dynamic segment:
 - Channel A or channel B

Message buffer 0 resp. 1 is used as a dedicated buffer to store startup frames and synchronization frames or as a dedicated buffer for the specified single-slot frames, as set by SUCC1:TXST, SUCC1:TXSY, and SUCC1:TSM, respectively.

In these cases, message buffer 0 can be reconfigured only in the DEFAULT_CONFIG or CONFIG state. This ensures that any node transmits at most one startup or sync frame per communication cycle. Transmission of startup or sync frames from other message buffers is not possible. Except buffer 0, all message buffers configured for static segment or dynamic segment transmission can be reconfigured during runtime according to the MRC:SEC[1:0] setting. However, the data pointer in the header partition references the data partition in the message RAM. Therefore, if the payload length and data pointer in the header section of a message buffer are set again, the message buffer structure may be incorrect.

If a message buffer is reconfigured during runtime (the header section is updated), the message buffer may not be transmitted in this communication cycle.

The header CRC must be provided to all transmission buffers because the FlexRay controller has no function for calculating the header CRC.

If network management is required, the host must set the PPIT bit in the header section of each message buffer to '1' and write network management information to the data section of the message buffer. (See [30.5 Network Management](#)).

The payload length field stores the payload length in 2-byte units. If the payload length of the set static transmission buffer is shorter than the setting in MHDC:SFDL[6:0], padding bytes are inserted to guarantee the payload length of the static frame. The padding bytes have value '0'.

Note: In cases of an odd payload length (where PLC is 1, 3, 5, and so on), a 16-bit zero must be written at the end of the message buffer's data section to ensure that padding is all '0'.

The transmission mode can be set on each transmission buffer by the transmission mode flag, TXM. If this bit is set to '1', the message is transmitted in single-shot mode. If the bit is set to '0', it is transmitted in continuous mode.

In single-shot mode, the TXR flag of each message buffer is cleared to '0' after transmission is completed. Then, the transmission buffer can be overwritten by the next message to be transmitted.

In continuous mode, the TXR flag of each message buffer is not cleared to '0' after transmission is completed. In this case, a frame is transmitted each time the filter criteria match. The TXR flag can be reset by the host by writing the respective message buffer number to the IBCR register while bit IBCM.STXRH is set to '0'.

If multiple transmission buffers satisfy the filter criteria, the transmission buffer with the lowest buffer number is used for transmission in each slot.

30.7.4 Frame Transmission

The following procedure is required for preparing a message buffer for transmission.

- Configure the transmission buffer in the message RAM through WRHS1, WRHS2, and WRHS3.
- Write data to the data section of the transmission buffer through WRDSn.
- Write the target buffer number to the IBCR register to transfer the configuration and message data from the input buffer to the message RAM.
- If the IBCM register is configured for message transmission, the transmission request flag TXR for the message buffer is set to '1' as soon as the transfer from the input buffer is completed, and the message buffer is ready for transmission.
- A check of each TXR bit (TXR is 0) in the TXRQ1/2/3/4 register (only in single-shot mode) can verify whether transmission of the message buffer has completed.

After the transmission is completed (in single-shot mode), each TXR flag of the TXRQ1/2/3/4 register is cleared to '0'. Also, if the MBI bit in the header section of the message buffer is set to '1', SIR:TXI is set to '1'. An interrupt is generated if the interrupt is enabled.

30.7.5 Null Frame Transmission

If the transmission request flag is not set to '1' in the static segment before the transmission time and none of the other transmission buffers satisfies the filter criteria, the FlexRay controller transmits a null frame with the null frame indicator set to '0' and payload data cleared to '0'.

A null frame is transmitted in the following cases.

- The transmission request flag of the TXRQ1/2/3/4 registers is not set (TXR is '0') on the message buffer that matches the filter criteria and has the smallest buffer number.
- All the transmission buffers have cycle counter filters, but none of them matches the current cycle. In this case, the message buffer status (MBS) is not updated.

No null frame is transmitted in the dynamic segment.

30.8 Reception Procedure

30.8.1 Reception Buffer

Message buffers can be used as dedicated reception buffers when the CFG bit in the header section of the message buffer is set to '0' through WRHS1.

A reception buffer can be assigned to a FlexRay controller channel in the following ways:

- Static segment:
 - Channel A or channel B

- ❑ Channel A and channel B (first semantically valid frame is stored)
- Dynamic segment:
 - ❑ Channel A or channel B

The FlexRay controller stores all the elements (except the frame CRC) of the frame that matches the filter criteria to the reception buffer.

All reception message buffers configured for static segments or dynamic segments can be reconfigured during runtime, through the MRC:SEC[1:0] setting. See [30.10.1 Message Buffer Reconfiguration](#). However, if the header section of the message buffer is reconfigured during runtime, a received message in this communication cycle may be lost.

If multiple buffers match the filter criteria, the reception buffer with the smallest message buffer number is updated with the received message.

30.8.2 Frame Reception

The following procedure is required to prepare a message buffer for reception.

- Configure the reception buffer in the message RAM through WRHS1, WRHS2, and WRHS3.
- Transfer settings from the input buffer to the message RAM by writing the target message buffer number to the IBCR register.

When these steps are performed, the message buffer functions as an active reception buffer and participates in the acceptance filtering process, which is executed each time a message is received. The reception buffer that matches the filter criteria first is updated by the received message.

If a valid payload segment is stored in the data section of the message buffer, the respective ND flag of the NDAT1/2/3/4 register is set to '1'. Also, if the MBI bit of the header section of the message buffer is set to '1', the SIR:RXI flag is set to '1'. An interrupt is generated if the interrupt is enabled.

If the ND bit is already set to '1' when the message buffer is updated, MBS:MLST in the reception message buffer is set, and unprocessed message data is lost. If a slot has received no frame, a null frame, or a corrupted frame, the data section of the message buffer configured for this slot is not updated. In this case, only the respective message buffer status flags are updated.

When the status flag in the header section of the message buffer is updated, the respective MBS flag in the MBSC1/2/3/4 register is set to '1'. If the MBI bit of the header section of the message buffer is set to '1', the SIR:MBSI flag is set to '1'. An interrupt is generated if the interrupt is enabled.

If the payload length PLR[6:0] of a received frame is longer than the set value in PLC[6:0] in the header section of the

message buffer, the data field stored in the message buffer is truncated to that length.

For data transfer between the output buffer and the message RAM, follow the procedure described in [30.10.2 Host Access to Message RAM](#).

Note: The ND and MBS flags are cleared to '0' when the payload data and header, respectively, of the received message are transferred to the output buffer.

30.8.3 Null Frame Reception

The reception buffer does not reflect the payload segment of a received null frame. If a null frame is received, the header section of the reception buffer is updated by the received null frame. All bits in header 2 and 3 of the matching message buffer remain unchanged. They are updated from received data frames only.

Each MBS flag of the MBSC1/2/3/4 register is set to '1' when the status flag in the header section of the message buffer is updated. If the MBI bit of the header section of the message buffer is set to '1', the SIR:MBSI flag is set to '1'. An interrupt is generated if the interrupt is enabled.

30.9 FIFO Function

30.9.1 Details

A group of message buffers can be configured as the FIFO buffer. Message buffers belonging to the FIFO message group are adjacent to one another in the register map, starting with the message buffer referenced by MRC:FFB[7:0] and ending with the message buffer referenced by MRC:LCB[7:0]. A maximum of 127 message buffers can be assigned to the FIFO.

The FIFO stores all valid received messages that do not match the filter criteria of dedicated reception buffers and match the set FIFO filter criteria. In such cases, the frame ID, payload length, receive cycle count, and status bit of the specified FIFO message buffer are overwritten by a received frame. Bit SIR:RFNE shows that the FIFO is not empty, SIR:RFCL is set when the receive FIFO fill level FSR:RFLL[7:0] is equal or greater than the critical level as configured by FCL:CL[7:0], and EIR:RFO shows that a FIFO overrun has been detected. An interrupt is generated if the interrupt is enabled.

If null frames are not rejected by the FIFO rejection filter, the null frames will be treated like data frames when they are stored into the FIFO.

If the FIFO is full and a new message is written before the oldest message in the FIFO is read, the new message overwrites the oldest message in the FIFO. As a result, the EIR:RFO flag is set to '1'.

The FIFO rejection filter (FRF) defines a filter pattern for rejecting messages. The filter consists of the channel filter, frame ID, and cycle counter filter. If the FRF:RSS bit is set to '1', the FIFO filter rejects all the messages received in the static segment. If the FRF:RNF bit is set to '1', the FIFO does not store received null frames.

The FIFO rejection filter mask (FRFM) specifies which bits of the frame ID filter in the FIFO rejection filter register are not used for rejection filtering.

30.9.2 FIFO Settings

Re-configuration of message buffers belonging to the FIFO is possible only when the CC is in DEFAULT_CONFIG or CONFIG state. While the CC is in DEFAULT_CONFIG or CONFIG state, the FIFO function is not available.

Set the payload length PLC[6:0] to the same value on all message buffers belonging to the FIFO, through WRHS2. Also, set the data pointer to the first 32-bit word of the data section of each message buffer in the message RAM through WRHS3.

All the information required for the acceptance filter is set by the FIFO rejection filter and FIFO rejection filter mask. Therefore, there is no need to set the filter criteria in the header section of each message buffer belonging to the FIFO.

Note: Program the MBI bits of the message buffers belonging to the FIFO to '0' via WRHS1.MBI to avoid generation of RX interrupts

If the payload length of a received frame is longer than the set value of PLC[6:0] in the header section of the respective message buffer, the data field stored in the FIFO message buffer is truncated to the length of PLC[6:0].

30.9.3 Access to FIFO

A transfer from the message RAM to the output buffer is triggered by writing the first FIFO message buffer number (referenced by MRC:FFB[7:0]) to the OBCR register. That operation will transfer the message buffer pointed to by the FIFO internal read index and afterwards increment the read index.

30.10 Message Handling

The message handler controls data transfer between the input/output buffer and message RAM, and between the message RAM and the two transient buffer RAM units. Any access to the internal RAM is in units of 32+1 bits. The additional bit is used for a parity check.

The message buffers stored in the message RAM are accessed under the control of the message handler state machine. This prevents access collisions between the host

and the two FlexRay channel protocol controllers to the message RAM.

The frame IDs of the message buffers assigned to the static segment must be in the range of 1 to GTUC7:NSS[9:0] (which has the number of static slots). The frame IDs of the message buffers assigned to the dynamic segment must be in the range of GTUC7:NSS[9:0]+1 to 2047.

Any received message that does not match the filter criteria of the dedicated reception buffer (static segment or dynamic segment) but matches the filter criteria of the FIFO rejection filter is stored in the reception FIFO (if configured).

30.10.1 Message Buffer Reconfiguration

For an application that requires more than 128 message buffers, the static and dynamic message buffers may be reconfigured during FlexRay operation. This can be done by updating the header section of the respective message buffer via the WRHS1 to WRHS3 input buffer registers.

The MRC:SEC[1:0] control bits in the message RAM configuration register must enable that reconfiguration.

If the message buffer is not updated by a received frame or a transmission message is not transmitted from the message buffer before reconfiguration begins, the message is lost.

The time when a reconfigured message buffer is ready for transmission or reception according to the reconfigured frame ID depends on the actual state of the slot counter when the update of the header section has completed. Therefore, it may happen that a reconfigured message buffer is not transmitted or updated from a received frame in the cycle where it was reconfigured

The message RAM is scanned according to the following table.

Table 30-7. Scan of Message RAM

Start of Scan in Slot	Scan for Slots
1	2...15, 1 (next cycle)
8	16...23, 1 (next cycle)
16	24...31, 1 (next cycle)
24	32...39, 1 (next cycle)
...	...

A message RAM scan ends at the start of NIT even if it has not been completed. The message RAM scan for slots 2 to 15 starts at the beginning of slot 1 of the actual cycle. The message RAM scan for slot 1 is done in the previous cycle by checking in parallel to each scan of the message RAM whether there is a message buffer configured for slot 1 of the next cycle.

The first dynamic message buffer number is set by MRC:FDB[7:0]. If the message RAM scan starts while CC is

in the dynamic segment, the scan starts at the message buffer number set in MRC:FDB[7:0].

The following procedure is required to reconfigure a message buffer for use in slot 1 of the next cycle.

- If the message buffer to be reconfigured for slot 1 is part of the "Static Buffers", it will only be found if it is reconfigured before the last message RAM scan in the static segment of the actual cycle evaluates this message buffer.
- If the message buffer to be reconfigured for slot 1 is part of the "Static + Dynamic Buffers", it will be found if it is reconfigured before the last message RAM scan in the actual cycle evaluates this message buffer.
- The message RAM scan finishes at the start of NIT. If the message RAM scan has not yet evaluated the reconfigured message buffer by that time, the message buffer is not considered during the next cycle.

Note: Reconfiguration of a message buffer may lead to a loss of messages, so great care must be taken in the reconfiguration. If a message buffer is reconfigured in consecutive cycles, it may happen that a message buffer is never transmitted or updated from a received frame.

See [30.11.3 Message Buffer Assignment](#) for configuration of message buffer partition and the relevant registers

30.10.2 Host Access to Message RAM


Message transfer between the input buffer and message RAM, or between the message RAM and output buffer begins when the target or source message buffer number is written to the IBCR or OBCR register, respectively.

[Figure 30-10](#) shows the access image between the host and message RAM.

The IBCM and OBCM registers can be used to read/write the header and data sections of the selected message buffer separately. If bit IBCM.STXR is set to '1', the transmission request flag TXR of the selected message buffer is automatically set after the message buffer has been updated. If IBCM.STXR is reset to '0', TXR of the selected message buffer is reset. This can be used to stop transmission from message buffers operated in continuous mode.

The input buffer (IBF) and output buffer (OBF) compose a double buffer structure. The IBF and OBF host in this double-buffer configuration can be accessed from the host.

Figure 30-10. Host Access to Message RAM




30.10.2.1 Data Transfer from Input Buffer to Message RAM

To configure or update message buffers in the message RAM, data must be written to WRDSn, and the header must

be written to WRHS1 to WRHS3. A specific operation is selected with the IBCM setting.


The IBF host and IBF shadow are swapped by writing the target message buffer number in the message RAM to IBCR:IBRH[6:0]. See Figure 30-11 .

Figure 30-11. Double-buffer Structure of Input Buffer



Bits in the IBCM and IBCR registers are also swapped with each other to retain the association with the respective section in the IBF. See Figure 30-12.

Figure 30-12. IBCM Register and IBCR Register Bit Swapping



This write operation sets IBCR:IBSYS to '1'. The message handler starts transmitting the contents of the IBF shadow to the message buffer in the message RAM selected by IBCR:IBRS[6:0].

The next message can be written to the IBF host while data is transferred from the IBF shadow to the target message buffer in the message RAM. After the transfer between the IBF shadow and message RAM is completed and the IBCR:IBSYS bit is cleared to '0', the next transfer to the message RAM begins when the next target message buffer number is written to IBCR:IBRH[6:0].

If a write access to IBCR:IBRH[6:0] occurs while IBCR:IBSYS is '1', IBCR:IBSYH is set to '1'. After completion of the ongoing data transfer from IBF shadow to message RAM, the IBF host and IBF shadow are swapped, IBCR:IBSYH is reset to '0', IBCR:IBSYS remains set to '1', and the next transfer to the message RAM is started. In addition, the message buffer numbers stored under IBCR:IBRH[6:0] and IBCR:IBRS[6:0] and the command mask flags are also swapped.

Example of the input buffer setting procedure:

- Configure/update the first message buffer via the IBF.
 - Wait until IBCR:IBSYH is reset.

- Write the data section to WRDSn (n:1–3).
- Write the header section to WRHS1 to WRHS3.
- Command mask writing: Write to IBCM:LHSH, IBCM:LDSH, and IBCM:STXRH.
- Data transfer request to the target message buffer: Write to IBCR:IBRH[6:0].
- Configure/update the second message buffer via the IBF.
 - Wait until IBCR:IBSYH is reset.
 - Write the data section to WRDSn. (n:1–3).
 - Write the header section to WRHS1 to WRHS3.
 - Command mask writing: Write to IBCM:LHSH, IBCM:LDSH, and IBCM:STXRH.
 - Data transfer request to the target message buffer: Write to IBCR:IBRH[6:0] after IBCR:IBSYH is cleared to '0'.
- Configure/update the third message buffer via the IBF.
 - Repeat the procedure to configure/update the second message buffer.

Note: Any write access to IBF while IBCR:IBSYH is '1' will set error flag EIR.IIBA to '1'. In this case the write access has no effect.

Table 30-8. Assignment of Input Buffer Command Mask (IBCM) Register

Bit Field and Bit Name	Access	Function
STXRS	r	Set transmission request shadow ongoing
LDSS	r	Load data section shadow ongoing
LHSS	r	Load header section shadow ongoing
STXRH	r/w	Set transmission request host
LDSH	r/w	Load data section host
LHSH	r/w	Load header section host


Table 30-9. Assignment of Input Buffer Command Request (IBCR) Register

Bit Field and Bit Name	Access	Function
IBSYS	r	IBF busy shadow. Signals ongoing transfer from the IBF shadow to the message RAM
IBRS[6:0]	r	IBF request shadow. Message buffer number currently/latently updated
IBSYH	r	IBF busy host. Transfer request pending for the message buffer referenced by IBRH6:0
IBRH[6:0]	r/w	IBF request host, Message buffer number for the next update

30.10.2.2 Data Transmission from Message RAM to Output Buffer

To read a message buffer from the message RAM, it is required to write to the OBCR register to trigger data transfer as configured in OBCM. After the transfer is completed, the data can be read from RDDSn, RDHS1 to RDHS3, and MBS.

Figure 30-13. Double-buffer Structure of Output Buffer




OBF host and OBF shadow as well as the OBCM.RHSS, OBCM.RDSS, OBCM.RHSH, and bits OBCM.RDSH, and OBCR.OBRS[6:0] and OBCR.OBRH[6:0] bits are swapped under control of the OBCR.VIEW and OBCR.REQ bits.

Writing bit OBCR.REQ to '1' copies OBCM.RHSS and OBCM.RDSS, and OBCR.OBRS[6:0] bits to an internal storage (see Figure 30-14).

After setting OBCR.REQ to '1', OBCR.OBSYS is set to '1'; the message buffer selected by OBCR.OBRS[6:0] from the message RAM to OBF shadow is transferred. After the transfer is complete, the OBCR.OBSYS bit is set back to '0'. The OBCR.REQ and OBCR.VIEW bits can only be set to '1' while OBCR.OBSYS is '0'.

Figure 30-14. OBCM Register and OBCR Register Bit Swapping



If '1' is set in OBCR:VIEW while the OBCR:OBSYS bit is '0', the OBF host and OBF shadow are swapped with each other. (See Figure 30-13). In addition OBCR.OBRH[6:0] and OBCM.RHSH, OBCM.RDSH bits are swapped with the registers' internal storage thus assuring that the message buffer number stored in OBCR.OBRH[6:0] and the mask configuration stored in OBCM.RHSH, OBCM.RDSH matches the transferred data stored in OBF host. (See Figure 30-14)

After the swap, the transferred message buffer can be read from the OBF host, and the message handler can transfer the next message from the message RAM to the OBF shadow. If both REQ and VIEW are set to '1' simultaneously while OBSYS is '0', OBSYS is set to '1'. Subsequently, the OBF host and OBF shadow are swapped with each other. Also, the OBCM.RDSH and OBCM.RHSH mask bits are swapped with the internal registers to keep them attached to the respective output buffer transfer. Afterwards OBRH[6:0] is copied to the register internal storage, mask bits OBCM.RDSS and OBCM.RHSS are copied to register OBCM internal storage, and the transfer of the selected message buffer from the Message RAM to OBF shadow is started. While the transfer is in progress, the CPU can read the message buffer previously transferred from the OBF host. The OBSYS bit is cleared to '0' when the transfer between the message RAM and OBF shadow is completed.

Example of the output buffer setting procedure:

- Request transfer of the first message buffer to the OBF shadow.
 - Wait until OBCR:OBSYS is cleared to '0'.
 - Command mask writing for first message buffer: Write to OBCM:RHSS and OBCM:RDSS.
 - Request transfer of the first message buffer to OBF shadow: Write to OBCR:OBRH[6:0] and OBCR:REQ.
- Toggle OBF shadow and OBF host to read out first transferred message buffer and request transfer of second message buffer
 - Wait until OBCR:OBSYS is cleared to '0'.
 - Command mask writing for second message buffer: Write to OBCM:RHSS and OBCM:RDSS.
 - Swap OBF host and shadow and request transfer of second message buffer to OBF shadow simultane-

ously: Write to OBCR:VIEW, OBCR:REQ, and OBCR:OBRH[6:0].

- Read the first message.
- ... (Repeat the same steps.)
- Read the nth message buffer from the OBF host (assuming no more requests for message buffer transfer).
 - Wait until OBCR:OBSYS is cleared to '0'.
 - Request access to last transferred message buffer: Write to OBCR:VIEW.
 - Read the nth message.

Table 30-10. Assignment of Output Buffer Command Mask (OBCM) Register

Bit Field and Bit Name	Access	Function
RDSH	r	Data section available for host access.
RHSH	r	Header section available for host access.
RDSS	r/w	Read data section shadow.
RHSS	r/w	Read header section shadow.

Table 30-11. Assignment of Output Buffer Command Request (OBCR) Register

Bit Field and Bit Name	Access	Function
OBRH[6:0]	r	OBF request host Number of message buffer available for host access
OBSYS	r	OBF busy shadow Signals ongoing transfer from message RAM to OBF shadow
REQ	r/w	Request transfer from message RAM to OBF shadow
VIEW	r/w	View OBF shadow (swap OBF shadow and OBF host)
OBRH[6:0]	r/w	OBF request shadow Number of message buffer for next request

30.10.3 FlexRay Protocol Controller Access to Message RAM


Two transient buffer RAM units (TBF A and TBF B) are used to buffer data for the transfer between the two FlexRay channel protocol controllers and the message RAM. [Figure 30-15](#) shows the transient buffer RAM units.

The two transient buffer RAM units compose a double buffer, and they can store two complete FlexRay messages. One buffer is always accessible by the message handler, while the other one is assigned to the corresponding channel protocol controller.

For example, if the message handler writes a transmission message to the transient buffer Tx, the FlexRay channel protocol controller can store the received message in the transient buffer RX. While the message stored in TX is being transmitted, the message handler transfers the latest received message (if it passes the acceptance filter) stored in RX to the message RAM and updates the message buffer.

The data transfer between the transient buffer RAM and the shift register of the FlexRay channel protocol controller is in units of 32-bit words. This enables the use of a 32-bit shift register independent of the length of the FlexRay message.

Figure 30-15. Access to Transient Buffer RAM



30.11 Message RAM


Direct access to the message buffer in the message RAM is disabled to prevent collisions between host access to the message RAM and transmission and reception of FlexRay messages. Direct access to the message RAM and all other RAM blocks of the FlexRay controller is possible only in the test mode; see [30.14.3 RAM Test Mode](#).

Access is processed via the input buffer and output buffer. The message RAM can store up to 128 message buffers (depending on the configured payload length).

The message RAM consists of 2048 bytes × 33 bits, which equals 67,584 bits, and a parity bit protects each piece of 32-bit data. The structure of the message RAM can have a variable (0 to 254) number of data bytes per FlexRay frame, as shown in [Figure 30-16](#).

The data partition is allowed to start at message RAM word number: $(MRC:LCB+1) \times 4$

Figure 30-16. Example of Message Buffer Configuration in Message RAM



Header Partition

The header partition stores the header sections of the configured message buffers:

- A maximum of 128 message buffers are supported.
- Each message buffer has a four-word (1 word = 32 + 1 bits) header section.
- Header 3 of each message buffer has an 11-bit data pointer to the respective data section in the data partition.

- If each data section is 254 bytes, the maximum is 30 message buffers.
- If each data section is 128 bytes, the maximum is 56 message buffers.
- If each data section is 48 bytes, the maximum is 128 message buffers.

Note: Header partition + data partition cannot occupy more than 2048 words.

Data Partition

The data partition allows flexible storage of data sections with different lengths. Some example for the maximum number of message buffers for various data lengths are as follows.

30.11.1 Header Partition

The header partition of the message RAM stores the message buffer status and message buffer setting elements as shown in Figure 30-17. The header section of the message buffer is configured via the IBF (WRHS1 to WRHS3), and the header section is read via the OBF (RDHS1 to RDHS3 + MBS). The data pointer in the header section that defines the start position of the data section for the respective message buffer needs to be correctly configured and should not be changed during runtime. Reconfiguration of the message buffers belonging to the FIFO message group is only possible in the DEFAULT_CONFIG or CONFIG state.

The header section of each message buffer occupies four words (1 word = 32 + 1 bits) in the header partition of the message RAM. The header section of message buffer 0 starts at the first word of the message RAM.

The header CRC for the transmission buffers must be calculated by the host.

Payload Length Received PLR[6:0], Receive Cycle Count RCC[5:0], Received on Channel Indicator RCI, Startup Frame Indicator SFI, Sync Frame Indicator SYN, Null Frame Indicator NFI, Payload Preamble Indicator PPI, and Reserved Bit RES are updated from received valid data frames only

Header word 3 of each configured message buffer contains the corresponding message buffer status (MBS).

Figure 30-17. Header Section of Message Buffer in Message RAM

Bit Word	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	P			M B I	T X M	P P I T	C F G	C H B	C H A		Cycle Code										Frame ID												
1	P		Payload Length Received								Payload Length Configured										TX Buffer : Header CRC Configured RX Buffer : Header CRC Received												
2	P			R E S	P P I	N F I	S Y N	S F I	R C I		Receive Cycle Count										Data Pointer												
3	P			R E S S	P P I S	N F I S	S Y N S	S F I S	R C I S		Cycle Count Status						F T B	F T A		M L S T	E S B	E S A	T C I B	T C I A	S V O B	S V O A	C E O B	C E O A	S E O B	S E O A	V F R B	V F R A	
...	P	...																															
...	P	...																															

Frame Configuration	(PPIT, CFG, Frame ID, Payload Length Configured)
Filter Configuration	(CHB, CHA, Cycle Code)
Message Buffer Control	(MBI, TXM)
Message RAM Configuration	(Data Pointer)
Updated from received Frame	(Payload Length Received, RES, PPI, NFI, SYN, SFI, RCI and Receive Cycle Count)
Message Buffer Status MBS	(MLST, ESB, ESA, TCIB, TCIA, SVOB, SVOA, CEOB, CEOA, SEOB, SEOA, VFRB, VFRA, RESS, PPIS, NFIS, SYNS, SFIS, RCIS, Cycle Count Status, FTB, FTA)
P	Parity Bit
	Unused

Header 1 (word 0)

The following parameters are written via WRHS1 and read via RDHS1.

- Frame ID: Slot counter filtering setting
- Cycle code: Cycle counter filtering setting
- CHA, CHB: Channel filtering settings

- CFG: Message buffer direction setting: reception/Transmission
- PPIT: Payload preamble indicator transmission
- TXM: Transmission mode setting: Single-shot/Continuous
- MBI: Enable flag for message buffer transmission/reception interrupts

Header 2 (word 1)

The following parameters are written via WRHS2 and read via RDHS2.

- Header CRC
 - Transmission buffer: Configured by host (calculated from frame header)
 - Reception buffer: Updated by received frame
- Payload length configured: Data section length (2-byte units) as configured by host
- Payload length received: Length of payload segment (2-byte units) stored from received frame

Header 3 (word 2)

The following parameters are written via WRHS3 and read via RDHS3.

- Data pointer: Pointer to the start position of the corresponding data section in the data partition

The following parameters are read via RDHS3. They are valid only for reception buffers and updated by received frames

- Receive cycle count: Stored cycle count value from the received frame
- RCI: Reception channel indicator
- SFI: Startup frame indicator
- SYN: Sync frame indicator
- NFI: Null frame indicator
- PPI: Payload preamble indicator
- RES: Reserved bit

Header 4 (word 3)

Header 4 is read via MBS. It is updated at the end of the configured slot.

- VFRA: Valid frame reception channel A
- VFRB: Valid frame reception channel B
- SEOA: Syntax error observed channel A
- SEOB: Syntax error observed channel B
- CEOA: Content error observed channel A
- CEOB: Content error observed channel B
- SVOA: Slot boundary violation observed channel A
- SVOB: Slot boundary violation observed channel B
- TCIA: Transmission conflict indication channel A
- TCIB: Transmission conflict indication channel B
- ESA: Empty slot channel A
- ESB: Empty slot channel B
- MLST: Message lost
- FTA: Frame transmitted channel A
- FTB: Frame transmitted channel B

- Cycle Count Status: Actual cycle count when status was updated
- RCIS: Received on channel indicator status
- SFIS: Startup frame indicator status
- SYNS: Sync frame indicator status
- NFIS: Null frame indicator status
- PPIS: Payload preamble indicator status
- RESS: Reserved bit status

30.11.2 Data Partition

The data partition in the message RAM stores the data sections of the message buffers configured for reception or transmission as defined in the header partition. The number of data bytes that can be set for each message buffer ranges from 0 to 254 bytes. The bit width of the message RAM is set in 32 bits + 1 parity bit to optimize the data transfer between the host interface and message RAM and the data transfer between the shift register of two FlexRay channel protocol controllers and the message RAM.

The data partition starts after the last word of the header partition. The data pointers must point to an address within the data partition when configuring the message buffer in the message RAM. [Figure 30-18](#) shows an example of how the data section of the configured message buffers can be stored in the data partition in the message RAM.

The start and end of the data section of a message buffer is determined by the data pointer and payload length configured in the header section of the message buffer. This enables flexible use of the RAM space for message buffers with different data length. If the data section size is an odd number of 2-byte units, the remaining 16 bits in the last 32-bit word are not used. See [Figure 30-18](#).

Figure 30-18. Example of Data Section Structure in Message RAM

Bit Word	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
...	P																																
2046	P																																
2047	P																																

30.11.3 Message Buffer Assignment

The assignment of message buffers is done according to the scheme shown in Figure 30-19. The number N of available message buffers depends on the payload length of the configured message buffers. The maximum number of message buffers is 128. The maximum payload length supported is 254 bytes.

The message buffers are separated into three consecutive groups:

- Static Buffers – Transmit/receive buffers assigned to static segment
- Static + Dynamic Buffers – Transmit/receive buffers assigned to static or dynamic segment
- FIFO - Receive FIFO

The message buffer separation configuration can be changed only in DEFAULT_CONFIG or CONFIG state by programming register MRC.


The first group starts with message buffer 0 and consists of static message buffers only. Message buffer 0 is dedicated to hold the startup/sync frame or the single slot frame, if the node transmits one, as configured by SUCC1.TXST, SUCC1.TXSY, and SUCC1.TSM. In addition, message buffer 1 may be used for sync frame transmission when sync frames or single-slot frames should have different payloads on the two channels. In this case, the MRC.SPLM bit must be programmed to '1' and message buffers 0 and 1 must be configured with the key slot ID and can be reconfigured in DEFAULT_CONFIG or CONFIG state only.

The second group consists of message buffers assigned to the static or to the dynamic segment. Message buffers

belonging to this group may be reconfigured during runtime from dynamic to static or vice versa depending on the state of MRC.SEC[1:0].

The message buffers belonging to the third group are concatenated to a single receive FIFO.

Figure 30-19. Assignment of Message Buffers



30.11.4 Parity Check

The FlexRay controller is equipped with a parity check mechanism to guarantee the integrity of data stored in the seven RAM blocks. These RAM blocks include the parity generator/checker connected as shown in Figure 30-20, and the parity generator generates the parity bit when data is written to a RAM block. The FlexRay controller uses even parity. (The parity bit is generated as '0' when there is an even number of 1s in a 32-bit word.)


The parity bit is stored together with the corresponding data word. The parity is always checked when data is read from a RAM block. The bit width of the internal data bus of the FlexRay controller is 32 bits.

The respective error flag is set to '1' when a parity error is detected. The message handler status register contains the parity error flags (MHDS:PIBF, MHDS:POBF, MHDS:PMR, MHDS:PTBF1, and MHDS:PTBF2) and the faulty message buffer indicators (MHDS:FMBD, MHDS:MFMB, and

MHDS:FMB[6:0]). These error flags control the EIR:PERR error interrupt flag.

Figure 30-20 shows the data paths between the RAM blocks and the parity generators/checkers.

Figure 30-20. Parity Generation and Check



Note: The parity generator and parity checker are not part of the RAM block but are part of the RAM access logic.

The following operations are executed when a parity error is detected.

All cases:

- The corresponding parity error flag in the MHDS register is set.
- The parity error flag, EIR:PERR, is set. An interrupt is generated if the interrupt is enabled.

Additionally in special cases:

1. Parity error during data transfer from input buffer RAM 1 or 2 to the message RAM
 - a. Transfer of the header and/or data:
 - The MHDS:PIBF bit is set.
 - The MHDS:FMBD bit is set to indicate that MHDS:FMB[6:0] points to a faulty message buffer.
 - MHDS:FMB[6:0] indicates the message buffer number with the error.
 - No transmission request bit is set for a transmission buffer.
 - b. Transfer of data only

Parity error when the header of the respective message is read from the message RAM

- The MHDS:PMR bit is set.
- The MHDS:FMBD bit is set to indicate that MHDS:FMB[6:0] points to a faulty message buffer.
- MHDS:FMB[6:0] indicates the message buffer number with the error.
- Data section of the message buffer is not updated.
- No transmission request bit is set for a transmission buffer.

2. Parity error during host reading from input buffer RAM 1 or 2
 - The MHDS:PIBF bit is set.
3. Parity error when scanning the header section of the message RAM
 - The MHDS:PMR bit is set.
 - The MHDS:FMBD bit is set to indicate that MHDS:FMB[6:0] points to a faulty message buffer.
 - MHDS:FMB[6:0] indicates the message buffer number with the error.
 - The message buffer with the parity error is ignored.

4. Parity error during data transfer from the message RAM to transient buffer RAM 1 or 2
 - ❑ MHDS:PMR is set.
 - ❑ MHDS:FMBD is set to indicate that MHDS:FMB[6:0] points to a faulty message buffer.
 - ❑ MHDS:FMB[6:0] indicates the message buffer number with the error.
 - ❑ Frame is not transmitted, or frame is invalidated by setting frame CRC to zero in case frame was already in transmission.
5. Parity error during data transfer from transient buffer RAM 1 or 2 to channel protocol controller 1 or 2
 - ❑ MHDS:PTBF1 or MHDS:PTBF2 are set.
 - ❑ Frames in transmission are invalidated by setting frame CRC to zero.
6. Parity error during data transfer from transient buffer RAM 1 or 2 to the message RAM
 - a. When reading header section:
 - ❑ MHDS:PMR is set.
 - ❑ MHDS:FMBD is set to indicate that MHDS:FMB[6:0] points to a faulty message buffer.
 - ❑ MHDS:FMB[6:0] indicates the message buffer number with the error.
 - ❑ The data section of the message buffer is not updated.
 - b. When reading transient buffer RAM 1 or 2:
 - ❑ MHDS:PTBF1 or MHDS:PTBF2 is set.
 - ❑ MHDS:FMBD is set to indicate that MHDS:FMB[6:0] points to a faulty message buffer.
 - ❑ MHDS:FMB[6:0] indicates the message buffer number with the error.
7. Parity error during data transfer from the message RAM to the output buffer RAM
 - ❑ MHDS:PMR is set.
 - ❑ MHDS:FMBD is set to indicate that MHDS:FMB[6:0] points to a faulty message buffer.
 - ❑ MHDS:FMB[6:0] indicates the message buffer number with the error.
8. Parity error during host reading from the output buffer RAM
 - ❑ MHDS:POBF is set.
9. Parity error while data is read from transient buffer RAM 1 or 2

If a parity error occurs when the message handler reads a frame with network management information (PPI is 1) from transient buffer RAM 1 or 2, the NMV1 to NMV3 network management vectors are not updated from that frame.

30.11.5 Parity Error Handling

Parity errors caused by transient bit flips can be fixed by the following methods.

30.11.5.1 Self-healing

Parity errors in the following locations are overwritten with the next write access to the disturbed bit(s) caused by host access or by FlexRay communication.

- Input buffer RAM 1 and 2
- Output buffer RAM 1 and 2
- Data section in message RAM
- Transient buffer RAM A
- Transient buffer RAM B

30.11.5.2 CLEAR_RAMs Command

When executed in the DEFAULT_CONFIG or CONFIG state, the CLEAR_RAMs command initializes all RAMs to zero.

30.11.5.3 Temporary Unlocking of Header Section

Fixing a parity error in the header section of a locked message buffer is possible with a transfer from the input buffer to the header section of the locked buffer.

For this transfer, write access to IBCR (specifying the message buffer number) must be immediately preceded by the unlock sequence normally used to leave CONFIG state.

For this transfer, the respective message buffer header is unlocked and data is updated regardless of whether the buffer belongs to the FIFO or whether its lock is controlled by MRC:SEC[1:0].

30.12 Interrupts

30.12.1 Error and Status Interrupts

Interrupts provide a close link to the protocol timing as they are triggered almost immediately when an error or status change is detected by the CC, a frame is received or transmitted, a configured timer interrupt is activated, or a stop watch event occurred.

This enables quick responses to specific error conditions, status changes, or timer events. However, if too many interrupts are generated, the host may miss deadlines required for the application. Therefore, the FlexRay controller supports enabling/disabling each interrupt source individually.

An interrupt may be generated in the following cases.

- An error is detected.
- A status flag is set.
- A timer reaches a preconfigured value.
- A message transferred from the input buffer to the message RAM or from the message RAM to the output buffer has completed.
- A stop watch event occurs.

Tracking status and generating interrupts when a status change or an error occurs are two independent tasks. Regardless of whether an interrupt is enabled, the corresponding status is tracked and indicated by the CC. The host has access to the actual status and error information by reading registers EIR and SIR.

All status and error interrupts can be flexibly assigned to one of the INT0 and INT1 interrupt lines that go to the host. These lines are controlled by the enabled interrupts. Also, these two interrupt lines, INT0 and INT1, can be enabled or disabled separately with the ILE:EINT0 and ILE:EINT1 settings, respectively.

30.13 Timers and Stop Watch

30.13.1 Timer 0

Timer 0 is an absolute timer that can be used to generate an interrupt or a trigger signal at an arbitrary point in time based on the cycle count and macrotick value. Status interrupt SIR.TI0 is set to '1' when the configured time is reached. Additionally, a trigger signal can be generated (see [30.15.2 Timer 0 Trigger Output](#) for more information).

Timer 0 can be activated as long as the POC is either in NORMAL_ACTIVE state or in NORMAL_PASSIVE state. Timer 0 is deactivated when leaving NORMAL_ACTIVE state or NORMAL_PASSIVE state except for transitions between the two states.

Before reconfiguration of the timer, the timer must be halted first by writing bit T0RC to '0'.

30.13.2 Timer 1

Timer 1 is a relative timer that can be used to generate an interrupt when the specified number of macroticks since the start of the timer has elapsed. Status interrupt SIR.TI1 is set to '1' when the configured time has been reached.

Timer 1 can be activated as long as the POC is either in NORMAL_ACTIVE state or in NORMAL_PASSIVE state. Timer 1 is deactivated when leaving NORMAL_ACTIVE state or NORMAL_PASSIVE state except for transitions between the two states.

Before reconfiguration of the timer, the timer must be halted first by writing bit T1RC to '0'.

30.13.3 Stop Watch

The Stop Watch function stores the current FlexRay time (cycle number, macrotick value, slot counter channel A and B) at the next macrotick counter increment after a trigger event has occurred. The following trigger sources can be configured:

- Software Trigger
- Interrupt Line 0 event
- Interrupt Line 1 event
- External Trigger (see [30.15.3 Stop Watch Event Trigger Input](#) for more information)

Status interrupt SIR.SWE is set to '1' when the Stop Watch is triggered.

30.14 Test Modes

The FlexRay controller supports following test modes:

- Asynchronous Transmit Mode
- Internal and External Loop Back Mode
- RAM Test Mode
- I/O Test Mode
- Additional status information

The test features can be configured with TEST1 register and TEST2 register. Writing to the test registers is only possible after TEST1.WRTEN is set to 1. The write access to set this bit must be preceded by an unlock sequence (LCK.TMK).

30.14.1 Asynchronous Transmit Mode (ATM)

The asynchronous transmit mode is entered by writing SUCC1.CMD[3:0] = "1110" while the CC is in CONFIG state and bit TEST1.WRTEN is set to '1'. This write operation must be directly preceded by two consecutive write accesses to the Configuration Lock Key (unlock sequence).

When called in any other state or when bit TEST1.WRTEN is not set, SUCC1.CMD[3:0] will be reset to “0000” = command_not_accepted. Reading CCSV.POCS[5:0] will return “00 1110” while the FlexRay controller is in ATM mode. To leave ATM mode, write SUCC1.CMD[3:0] = 0001 (CHI command: CONFIG). In ATM mode, transmission of a FlexRay frame is triggered by writing the number of the respective message buffer to IBCR.IBRH[6:0] while IBCR.STXR is set to ‘1’. In this mode wakeup, startup, and clock synchronization are bypassed. The CHI command SEND_MTS results in the immediate transmission of an MTS symbol. The cycle counter value of frames sent in ATM mode can be programmed using MTCCV.CCV[5:0] (writable in ATM and loop back mode only).

30.14.2 Loop Back Mode

The loop back mode is entered by writing SUCC1.CMD[3:0] = 1111 while the CC is in CONFIG state and bit TEST1.WRTEN is set to ‘1’. This write operation must be preceded by two consecutive write accesses to the Configuration Lock Key (unlock sequence). When called in any other state or when TEST1.WRTEN is not set, SUCC1.CMD[3:0] will be reset to “0000” = command_not_accepted. Reading CCSV.POCS[5:0] will return “00 1001” while the FlexRay controller is in loop back mode.

Loop back mode can be left by writing SUCC1.CMD[3:0] = 0001 (CHI command: CONFIG). The loop back mode is intended to check the module's internal data paths. Normal, time-triggered operation is not possible in loop back mode. There are two possibilities to perform a loop back test. External loop back via physical layer (TEST1.ELBE = 1) or internal loop back for in-system self-test (TEST1.ELBE = 0). In an internal loop back, the txena_n and txenb_n pins are in their inactive state; txda and txdb pins are set to HIGH; and rxda and rxdb pins are not evaluated.

A loop back test is started by the host configuring the FlexRay controller, then writing a message to the input buffer, and requesting the transmission by writing to the IBCR register. The Message Handler will transfer the message into the message RAM and then into the transient buffer of the selected channel. The Channel Protocol Controller (PRT) will read (in 32-bit words) the message from the transmit part of the transient buffer and load it into its RX / TX shift register. The serial transmission is looped back into the shift register; its content is written into the receive part of the channel's transient buffer before the next word is loaded.

The PRT and Message Handler will then treat this transmitted message as a received message, perform an acceptance filtering on frame ID and receive channel, and store the message into the message RAM if it passed acceptance filtering. The loop back test ends with the host requesting this received message from the message RAM and checking the contents of the output buffer.

Each FlexRay channel is tested separately. The FlexRay controller cannot receive messages from the FlexRay bus while it is in the loop back mode. The cycle counter value of frames used in loop back mode can be programmed using MTCCV.CCV[5:0] (writable in ATM and loop back mode only).

Note: In an odd payload the last two bytes of the looped-back payload will be shifted by 16 bits to the right inside the last 32-bit data word.

30.14.3 RAM Test Mode

In RAM test mode (TEST1.TMC[1:0] = 01), one of the seven RAM blocks can be selected for direct read/write access by programming TEST2.RS[2:0]. For external access, the selected RAM block is mapped to address space 400h to 7FF (1024 byte addresses or 256 word addresses).

Because the length of the message RAM exceeds the available address space, the message RAM is divided into segments of 1024 bytes. The segments can be selected by programming TEST2.SSEL[2:0].

Furthermore, the parity bit state of the last read access is shown in TEST2.RDPB and the parity bit state for the next write access can be configured with TEST2.WRPB. This can also be used to inject parity errors in the RAM blocks for test purposes.

30.14.4 I/O Test Mode

In I/O test mode (TEST1.TMC[1:0] = 10) output pins txda, txdb, txena_n, and _txenb_n, are driven to the values defined by bits TXA, TXB, TXENA, and TXENB in the TEST1 register. The values applied to input pins rxda and rxdb can be read from the RXA and RXB register bits in the TEST1 register.

30.14.5 Additional Status Information

The TEST1 register provides additional status information on bits AOA, AOB (Activity on Channel A/B) and CERA, CERB (Coding Error Report Channel A/B).

30.15 Traveo II Specific Functions

The FlexRay controller in Traveo II offers extensions with regard to the E-Ray block that will be described in this section.

30.15.1 Enable/Disable FlexRay Controller

Before the FlexRay controller can be used, it needs to be globally enabled in the CTL register. While the controller is disabled, only the Traveo II-specific registers, CTL and DMA_CTL, can be accessed. All other registers will cause a bus error.

30.15.2 Timer 0 Trigger Output

A pulse will be generated on the trigger output signal tr_tint0_out when the timer 0 interrupt is asserted.

30.15.3 Stop Watch Event Trigger Input

The trigger input signal tr_stpwt_in can be used to activate the stop watch that will capture the current cycle counter, macrotick, and slot counter channel A/B values.

30.15.4 DMA Trigger Interface for Input/Output Buffer Access

Trigger input and output signals are provided to transfer data between the system memories and the FlexRay controller message RAM via the input/output buffer using the DMA controller.

This feature can offload the CPU from the tedious task of writing/reading the message RAM via the input/output buffer.

The available trigger signals are shown in [Table 30-12](#) and they are configured in the DMA_CTL register.

Table 30-12. DMA Trigger Signals

Signal	Direction	Description.
tr_ibf_in	Input	Trigger input that indicates that the DMA transfer to write IBCR has completed.
tr_ibf_out	Output	Trigger output for triggering the DMA transfer from system memory to IBF.
tr_obf_in	Input	Trigger input that indicates that the DMA transfer from OBF to the system memory has completed.
tr_obf_out	Output	Trigger output for triggering the DMA transfer to write OBCR.

30.15.4.1 Input Buffer Access by DMA

The tr_ibf_* triggers are used to control two DMA channels. [Figure 30-21](#) and [Figure 30-22](#) show how the DMA channels and FlexRay trigger signals need to be configured and connected:

- Channel 0 will be triggered by tr_ibf_out and is responsible for writing the header (optional) and data sections of the input buffer. When the transfer is complete, it needs to trigger channel 1 using the Traveo II trigger multiplexer.
- Channel 1 is responsible for writing the IBCM (optional) and the IBCR. When it is done, it needs to assert the tr_ibf_in trigger signal using the Traveo II trigger multiplexer. Note that the DMA configuration may differ depending on whether IBCM will also be written. The two cases are depicted in [Figure 30-21](#) and [Figure 30-22](#). After channel 1 completes the write, it may raise an interrupt to notify the application that all message buffers have been processed.

Figure 30-21. DMA Channel Configuration for Input Buffer Access (IBCR only)





Figure 30-22. DMA Channel Configuration for Input Buffer Access (IBCM and IBCR)



The functional behavior of the FlexRay controller trigger signals related to the IBF is shown in [Figure 30-23](#).


The output trigger tr_ibf_out requests the DMA to transfer the message header/data from the system memory to the IBF. This trigger is generated when the following conditions are met:

- A trigger is detected on tr_ibf_in (an indication that DMA has configured IBCR for the previous message that was written to IBF).

- IBF is not busy (eray_busy = 0) – that is, the previous message has been transferred from IBF to message RAM.

For the first message, software needs to set tr_ibf_in manually using the Traveo II trigger multiplexer block to start processing (at this time the IBF is considered to be not busy).

Figure 30-23. Trigger Interface Operation for IBF



30.15.4.2 Output Buffer Access by DMA

The `tr_obf_*` triggers are used to control two DMA channels. Figure 30-24 and Figure 30-25 show how the DMA channels and the FlexRay trigger signals need to be configured and connected

- Channel 0 will be triggered by `tr_obf_out` and is responsible for writing the OBCM (optional) and the OBCR. Note that the DMA configuration differs depending on whether IBCM is also written. The two cases are depicted in Figure 30-24 and Figure 30-25. When the transfer is complete, it needs to trigger channel 1 using the Traveo II trigger multiplexer.

- Channel 1 is responsible for transferring the Header (optional) and Data sections from the OBF to the system memory. When it is done, it needs to assert the `tr_obf_in` trigger signal using the Traveo II trigger multiplexer. After channel 1 completes the transfer, it may raise an interrupt to notify the application that all message buffers have been processed.

Figure 30-24. DMA Channel Configuration for Output Buffer Access (OBCR only)





Figure 30-25. DMA Channel Configuration for Output Buffer Access (OBCM and OBCR)



Notes:

- For the first message, software needs to start DMA channel 0 manually.
- The first transfer of DMA channel 1 will transfer undefined dummy data because requesting a message buffer and reading the corresponding message data happens in a pipelined manner.

The tr_obf_out trigger signal is generated when the following conditions are met:

- An indication that the DMA read of the previous message is complete (tr_obf_in pulse causing internal logic node dma_transfer = 0).
- Previous OBCR command is complete and previous message has been moved from message RAM to OBF (eray_obusy = 0).
- If enabled, the receive FIFO is not empty and there is more data to read ((DMA_CTL.ODMAFFE and !FSR.RFNE) = 0).

The functional behavior of the FlexRay controller trigger signals related to the OBF for various scenarios is shown in [Figure 30-26](#), [Figure 30-27](#), [Figure 30-28](#), and [Figure 30-29](#).

Figure 30-26. Case 1: DMA Transfer Takes Longer than Message Buffer RAM to OBF Transfer




Figure 30-27. Case 2: Message Buffer RAM to OBF Transfer Takes Longer than DMA Transfer




Figure 30-28. Case 3: Receive FIFO is Not Empty after OBCR





Figure 30-29. Case 4: Receive FIFO is Empty after OBCR



30.15.4.3 Mirror Registers for DMA

In the register map of the original E-Ray block, the header registers in the input or output buffer are located after the corresponding data registers. This means that even if the actual messages have a small payload size, the DMA must be configured for the full payload size (64 words), if data and header sections are accessed by the DMA. This causes an unnecessary increase of DMA transfer times and inefficient utilization of system memory.

The Traveo II FlexRay controller has additional mirror registers to overcome this problem and provide a more efficient implementation of DMA transfers.

In the mirror address space, the header registers will be located before the corresponding data registers as shown in [Table 30-13](#).

Table 30-13. Mirror Registers

Address Offset	Register	Remarks
0xBF0	WRHS1_MIR2 (second mirror)	Header and data sections of input buffer can be written by input buffer DMA channel 0
0xBF4	WRHS2_MIR2 (second mirror)	
0xBF8	WRHS3_MIR2 (second mirror)	
0xBFC	Reserved	
0xC00 – 0xCFC	WRDSn_MIR (mirror)	These registers are kept for compatibility with original input buffer address space
0xD00	WRHS1_MIR (mirror)	
0xD04	WRHS2_MIR (mirror)	
0xD08	WRHS3_MIR (mirror)	
0xD0C	Reserved	Input Buffer Command Mask and Request register can be written by input buffer DMA channel 1
0xD10	IBCM_MIR	
0xD14	IBCR_MIR	
0xD18 – 0xDEC	Reserved	
0xDF0	RDHS1_MIR2 (second mirror)	Header and data sections of output buffer can be read by output buffer DMA channel 1
0xDF4	RDHS2_MIR2 (second mirror)	
0xDF8	RDHS3_MIR2 (second mirror)	
0xDFC	MBS_MIR2 (second mirror)	
0xE00 – 0xEFC	RDDSn_MIR (mirror)	These registers are kept for compatibility with original output buffer address space
0xF00	RDHS1_MIR (mirror)	
0xF04	RDHS2_MIR (mirror)	
0xF08	RDHS3_MIR (mirror)	
0xF0C	MBS_MIR (mirror)	Output Buffer Command Mask and Request register can be written by Output Buffer DMA channel 0
0xF10	OBCM_MIR (mirror)	
0xF14	OBCR_MIR (mirror)	

30.16 FlexRay Registers

The FlexRay controller has a set of 32-bit registers, listed in [Table 30-14](#). For more information on these registers, see the *Traveo II Body Controller High Registers TRM*.

Notes:

- 32-bit accesses must be used to read/write any FlexRay controller register.
- The mirror registers of the input/output buffer are not shown in [Table 30-14](#).

Table 30-14. List of FlexRay Controller Registers

Category	Register Name	Description
Traveo II Specific Registers	CTL	Control Register
	DMA_CTL	DMA Control Register
Special Register	TEST1	Test Register 1
	TEST2	Test Register 2
	LCK	Lock Register
Interrupt-related Registers	EIR	Error Interrupt Register
	SIR	Status Interrupt Register
	EILS	Error Interrupt Line Selection Register
	SILS	Status Interrupt Line Selection Register
	EIES	Error Interrupt Enable Register (set)
	EIER	Error Interrupt Enable Register (reset)
	SIES	Status Interrupt Enable Register (set)
	SIER	Status Interrupt Enable Register (reset)
	ILE	Interrupt Line Enable Register
	T0C	Timer 0 Configuration Register
	T1C	Timer 1 Configuration Register
	STPW1	Stop Watch Register 1
	STPW2	Stop Watch Register 2
Communication Controller (CC) Control Registers	SUCC1	SUC Configuration Register 1
	SUCC2	SUC Configuration Register 2
	SUCC3	SUC Configuration Register 3
	NEMC	NEM Configuration Register
	PRTC1	PRT Configuration Register 1
	PRTC2	PRT Configuration Register 2
	MHDC	MHD Configuration Register
	GTUC1	GTU Configuration Register 1
	GTUC2	GTU Configuration Register 2
	GTUC3	GTU Configuration Register 3
	GTUC4	GTU Configuration Register 4
	GTUC5	GTU Configuration Register 5
	GTUC6	GTU Configuration Register 6
	GTUC7	GTU Configuration Register 7
	GTUC8	GTU Configuration Register 8
	GTUC9	GTU Configuration Register 9
	GTUC10	GTU Configuration Register 10
	GTUC11	GTU Configuration Register 11

Table 30-14. List of FlexRay Controller Registers

Category	Register Name	Description
Communication Controller (CC) Status Registers	CCSV	CC Status Vector Register
	CCEV	CC Error Vector Register
	SCV	Slot Counter Value Register
	MTCCV	Macrotick and Cycle Counter Value Register
	RCV	Rate Correction Value Register
	OCV	Offset Correction Value Register
	SFS	Sync Frame Status Register
	SWNIT	Symbol Window and NIT Status Register
	ACS	Aggregated Channel Status Register
	ESID1	Even Cycle Sync Frame ID Register 1
	ESID2	Even Cycle Sync Frame ID Register 2
	ESID3	Even Cycle Sync Frame ID Register 3
Communication Controller (CC) Status Registers	ESID4	Even Cycle Sync Frame ID Register 4
	ESID5	Even Cycle Sync Frame ID Register 5
	ESID6	Even Cycle Sync Frame ID Register 6
	ESID7	Even Cycle Sync Frame ID Register 7
	ESID8	Even Cycle Sync Frame ID Register 8
	ESID9	Even Cycle Sync Frame ID Register 9
	ESID10	Even Cycle Sync Frame ID Register 10
	ESID11	Even Cycle Sync Frame ID Register 11
	ESID12	Even Cycle Sync Frame ID Register 12
	ESID13	Even Cycle Sync Frame ID Register 13
	ESID14	Even Cycle Sync Frame ID Register 14
	ESID15	Even Cycle Sync Frame ID Register 15
	OSID1	Odd Cycle Sync Frame ID Register 1
	OSID2	Odd Cycle Sync Frame ID Register 2
	OSID3	Odd Cycle Sync Frame ID Register 3
	OSID4	Odd Cycle Sync Frame ID Register 4
	OSID5	Odd Cycle Sync Frame ID Register 5
	OSID6	Odd Cycle Sync Frame ID Register 6
	OSID7	Odd Cycle Sync Frame ID Register 7
	OSID8	Odd Cycle Sync Frame ID Register 8
	OSID9	Odd Cycle Sync Frame ID Register 9
	OSID10	Odd Cycle Sync Frame ID Register 10
	OSID11	Odd Cycle Sync Frame ID Register 11
	OSID12	Odd Cycle Sync Frame ID Register 12
	OSID13	Odd Cycle Sync Frame ID Register 13
	OSID14	Odd Cycle Sync Frame ID Register 14
	OSID15	Odd Cycle Sync Frame ID Register 15
	NMV1	Network Management Register 1
	NMV2	Network Management Register 2
	NMV3	Network Management Register 3

Table 30-14. List of FlexRay Controller Registers

Category	Register Name	Description
Message Buffer Control Registers	MRC	Message RAM Configuration Register
	FRF	FIFO Rejection Filter Register
	FRFM	FIFO Rejection Filter Mask Register
	FCL	FIFO Critical Level Register
Message Buffer Status Registers	MHDS	Message Handler Status Register
	LDTS	Last Dynamic Transmission Slot Register
	FSR	FIFO Status Register
	MHDF	Message Handler Constraints Flags
	TXRQ1	Transmission Request Register 1
	TXRQ2	Transmission Request Register 2
	TXRQ3	Transmission Request Register 3
	TXRQ4	Transmission Request Register 4
	NDAT1	New Data Register 1
	NDAT2	New Data Register 2
	NDAT3	New Data Register 3
	NDAT4	New Data Register 4
	MBSC1	Message Buffer Status Changed Register 1
	MBSC2	Message Buffer Status Changed Register 2
	MBSC3	Message Buffer Status Changed Register 3
	MBSC4	Message Buffer Status Changed Register 4
Identification Registers	CREL	Core Release Register
	ENDN	Endian Register
Input Buffer	WRDSn	Write Data Section Register [1 to 64]
	WRHS1	Write Header Section Register 1
	WRHS2	Write Header Section Register 2
	WRHS3	Write Header Section Register 3
	IBCM	Input Buffer Command Mask Register
	IBCR	Input Buffer Command Request Register
Output Buffer	RDDSn	Read Data Section Register [1 to 64]
	RDHS1	Read Header Section Register 1
	RDHS2	Read Header Section Register 2
	RDHS3	Read Header Section Register 3
	MBS	Message Buffer Status Register
	OBCM	Output Buffer Command Mask Register
	OBCR	Output Buffer Command Request Register

31. Ethernet MAC



31.1 Overview

The Ethernet Media Access Controller (MAC) module in the device implements a 10/100/1000 Mbps Ethernet MAC compatible with the IEEE 802.3 standard, supporting MII, RMII, GMII, and RGMII PHY interfaces to support several automotive applications.¹


31.1.1 Supported Features

- Both Full Store and Forward mode and Partial Store and Forward mode for full-duplex operation
- 10 Mbit/s, 100 Mbit/s, or 1 Gbit/s operation
- MII, RMII, GMII, and RGMII PHY interface modes
- OPEN Alliance specified RGMII V2.2
- 1536 bytes of maximum frame length
- Three transmit and receive priority queues
- IEEE Std 802.1Qav – Forwarding and Queuing Enhancements for Time-Sensitive Streams
- IEEE Std 802.1AS – Timing and Synchronization for Time-Sensitive Application in Bridged LANs
- IEEE Std 1588 – Precision Time Protocol
- IEEE Std 802.1Qbb – Priority Based Flow Control
- 16 of each Screening Type 1 and Type 2 registers
- IEEE 802.3 Pause frame and MAC PFC priority based pause frame support
- Receive and transmit IP, TCP, and UDP checksum offload
- Automatic pad and CRC generation on transmitted frames
- MDIO interface for PHY management
- Strict priority, DWRR, or Enhanced Transmission Selection (ETS – 802.1Qaz) on transmit queues
- Support for 802.3az EEE

1. Please check device specific datasheet to confirm which of these interfaces are supported in the device.

31.2 Block Diagram

Figure 31-1. Ethernet MAC Block Diagram



31.3 Ethernet MAC Operation

31.3.1 DMA Interface

Ethernet MAC accesses data from other available system memory through DMA interface and stores fetched data in local dedicated TX/RX packet buffer. DMA is attached to the Ethernet MAC external FIFO interface to provide a scatter gather type capability for packet data storage. Configured for packet buffering mode, DMA uses dual port memory to store fetched data. This configuration makes application to use either of these mentioned operation modes to store and forward data.

- Full Store and Forward Mode
- Partial Store and Forward Mode

In full store and forward mode, a packet will automatically be replayed directly from the packet buffer memory rather than having to re-fetch from system memory through AXI. If transmission fails, received erroneous packets are automatically dropped before they are sent to system memory, thus reducing AXI activity.

In partial store and forward mode, the transmitter will only forward the packet to the MAC when there is enough frame data stored in the packet buffer. Similarly, in case of receive operation, the receiver will only begin to forward the packet to the external AHB or AXI slave when enough frame data are stored in the local packet buffer.

The following features have also become available due to this approach:

- Transmit TCP/IP checksum offload
- Priority queuing
- RX packet flush when there is lack of resource
- Burst padding at end of packet and end of buffer to maximize AXI efficiency
- TX/RX timestamp capture to buffer descriptor entry

31.3.1.1 AXI Interface

AXI master interface attached to the Ethernet MAC provides separate data channels and common address channels for read and write operations. With stated channels, interface supports two outstanding transactions on both the Read and Write channels.

EMAC requires to store configuration parameters for each transmit and receive frame through descriptors.

TX and RX descriptor reads are issued up-front and stored in a local buffer to feed the underlying DMA when required. This optimizes performance and avoids the need for the underlying DMA to pause while new descriptor fetches are sent to the system bus.

TX and RX descriptor writes issued by the underlying DMA are buffered locally to avoid holding up the underlying DMA when the system delays the completion of descriptor writes. Note that a descriptor write transaction is not considered complete until the write response (BRESP) associated with that transaction has arrived.

The maximum burst lengths the DMA can use are programmable. Single accesses and bursts with up to four

beats can be selected. With 64-bit data path and a burst length setting of 4, 32 bytes transfers can be made with a single request. The burst length is controlled via the *ETHx_dma_config* register.

31.3.1.2 Partial Store and Forward Using Packet Buffer DMA

This feature is enabled via the TX and RX partial store and forward programmable registers. When the transmit partial store and forward mode is activated, the transmitter will only begin to forward the packet to the MAC when there is enough packet data stored in the packet buffer. Similarly, when the receive partial store and forward mode is activated, the receiver will only begin to forward the packet to the external AHB or AXI slave when enough packet data is stored in the packet buffer.

The amount of packet data required to activate the forwarding process is programmable via watermark registers, which are located at the same address as the partial store and forward enable bits. Note that the minimum operational value for the TX partial store and forward watermark is 20. There is no operational limit for the RX partial store and forward watermark. To reduce the bandwidth requirements of the receive buffer manager the receive buffer size can be increased above its default value of 128 bytes by writing to the DMA configuration register.

Enabling partial store and forward is a useful means to reduce latency and increase Ethernet throughput.

31.3.1.3 Full Store and Forward Mode Using Packet Buffer DMA

In full store and forward mode, MAC only starts transmission when the complete transmit frame is written into the local TX buffer. It will be flushed from the local buffer only after MAC completes the transmission and TX BD is updated with the status fields.

In receive process, DMA starts forwarding data to configured memory address only after the entire frame is received and does not contain any error. Received frame will be flushed from local packet buffer only after frame is copied and RX BD is updated with the status.

When the EMAC DMA is configured in the full store and forward mode, a receive over run condition occurs when the receive packet buffer memory is full, or because an AXI error occurred. In partial store and forward mode, a receive overrun condition occurs when either the AXI bus was not granted quickly enough, or because an AXI error occurred, or because a new frame is detected by the receive block when the status update or write back for the previous frame is not yet finished. For a receive over run condition, the receive over run interrupt is asserted and the buffer being written is recovered. The next frame that is received whose address is recognized reuses the buffer.

The benefits of full store and forward mode are:

- Discard packets that are received with errors before they are partially written out of DMA thus saving AXI bandwidth and driver processing overhead
- Retry failed transmit frames from the packet buffer itself, thus saving AXI bus bandwidth
- Implement transmit IP/TCP/UDP checksum offload
- Allows multi-buffer frames

31.3.1.4 DMA Transaction

EMAC DMA uses separate transmit and receive lists of buffer descriptors, with each descriptor describing a buffer area in system memory. This allows Ethernet packets to be broken up and scattered around the system memory.

The DMA controller performs six types of operation on the AMBA bus. In order of priority these are:

- Receive buffer manager write/read
- Transmit buffer manager write/read
- Receive data DMA write
- Transmit data DMA read

All read operations are routed to the AXI read channel and all write operations to the AXI write channel. Both read and write channels may operate simultaneously. Arbitration logic is used when multiple requests are active on the same channel. For example: when transmit DMA requests a transmit data read at the same time receive DMA requests a receive descriptor read; in these case, the receive DMA is granted the bus before the transmit DMA. However, majority of requests are either receive data writes or transmit data reads both of which can operate in parallel and can execute simultaneously.

Transfer size is set to 64-bit words by default in the *ETHx_network_config* register and burst length can be programmed in the range from single access up to 16 accesses per burst using the *ETHx_dma_config* register. It is recommended to set burst length maximum to 4 to have quicker arbitration for all masters accessing the bus.

31.3.1.5 Receive Buffers

Received frames, optionally including FCS, are written to receive buffers located in system memory. The receive buffer depth (*rx_buf_size*[7:0]) is programmable in the range of 64 bytes to 16 Kbytes in the DMA configuration register, with the default being 1536 bytes. If received frames are being routed to different priority queues via screening registers, it is possible to program different receive buffer depths for each queue. For queue 0, the receive buffer depth is programmed through the DMA configuration register (offset 0x10). For the other queues, they are programmed through specific queue configuration registers (starting from offset 0x4a0). Default is 128 bytes.

Start address for each receive buffer is stored in system memory in a list of receive buffer descriptors at an address location described by the receive buffer queue pointer. The base address of the receive buffer queue pointer (also referred as list of buffer descriptors) must be configured by software using the receive buffer queue base address registers (*ETHx_receive_q_ptr*, *ETHx_receive_q1_ptr*, *ETHx_receive_q2_ptr*).

Each buffer descriptor can be either of two or four words, depending on configured buffer descriptor (BD) mode (*ETHx_dma_config[28]*), whereas word is defined as 32 bits. The first two words (Word 0 and Word 1) are used in both BD modes.

In Extended Buffer Descriptor mode (*ETHx_dma_config[28]* = 1), two BD words (Word 2 and Word 3) are added for timestamp capture if Timestamp Capture mode is enabled (*ETHx_rx_bd_control[5:4]* > 0_h). Therefore, BD modes will be either of two or four words size and each BD will have the same size. To summarize,

- Each BD must be of 64 bits when Descriptor Time Capture mode is disabled
- Each BD must be of 128 bits when Descriptor Time Capture mode is enabled

Following description details about Word 0 and Word 1 of each BD. Word 0 contains the start location of the receive buffer and Word 1 contains the receive status. If the length of a received frame exceeds the DMA buffer length, the status word (Word 1) in relevant BD is written with zeros except for the start of frame bit, which is always set for the first BD in a frame. Bit zero of the address field is set to '1' to show the buffer has been used. The receive buffer manager

then reads the location of the next receive buffer and fills that with the next part of the received frame data. Receive buffers are filled until the frame is complete, and the final buffer descriptor status word contains the complete frame status. See [Table 31-1](#) for details of the receive buffer descriptor list

When using receive descriptor timestamp capture (*ETHx_dma_config[28]* = 1), bit 2 of Word 0 is used to indicate a valid timestamp is captured in the BD. The use of bit 2 for this purpose also necessitates the data buffer being located on a 64-bit address boundary (EMAC only supports 32-bit address).

Each receive buffer start location is a word address. The start of the first buffer in a frame can be offset by up to three bytes depending on the value written to bits 15 and 14 of the network configuration register (*ETHx_receive_buffer_offset[1:0]*) and bit 2 of Word 0.

Table 31-1. Receive Buffer Byte Offset Configuration

Receive Buffer Offset Configuration Bit 2 of Word 0	ETHx_receive_buffer_offset[1]	ETHx_receive_buffer_offset[0]	Number of Bytes Offset
0	0	0	0
0	0	1	1
0	1	0	2
0	1	1	3

If the start location of the buffer is offset the available length of the first buffer is reduced by the corresponding number of bytes.

Table 31-2. Word 0 and Word 1 Description of Each BD

Bit	Function
Word 0	
31:3	Address [31:3] of beginning of buffer
2	Address [2] of beginning of buffer or In Extended Buffer Descriptor mode, indicates a valid timestamp in the BD entry.
1	Wrap - marks last descriptor in receive buffer descriptor list.
0	Ownership - needs to be '0' for the EMAC to write data to the receive buffer. EMAC sets this to '1' after it has successfully written a frame to memory. Software must clear this bit before the buffer can be used again.
Word 1	
31	Global all ones broadcast address detected.
30	Multicast hash match.
29	Unicast hash match.
28	External address match.
27	Unused.

Table 31-2. Word 0 and Word 1 Description of Each BD

Bit	Function
26:25	Specific Address register match. Encoded as follows: 00-Specific Address 1 register match (lowest priority) 01-Specific Address 2 register match 10-Specific Address 3 register match 11-Specific Address 4 register match (highest priority) If more than one specific address is matched only one of them is indicated with priority 4 down to 1.
24	This bit has a different meaning depending on whether RX checksum offloading is enabled (ETHx_network_config[24] = 1). With RX checksum offloading disabled: Type ID register match found, bit 22 and bit 23 indicate which Type ID register causes the match. With RX checksum offloading enabled: 0-The frame was not SNAP encoded and/or had a VLAN tag with the CFI bit set. 1-The frame was SNAP encoded and had either no VLAN tag or a VLAN tag with the CFI bit not set.
23:22	This bit has a different meaning depending on whether RX checksum offloading is enabled. With RX checksum offloading disabled: Type ID register match. Encoded as follows: 00-Type ID Match 1 register 01-Type ID Match 2 register 10-Type ID Match 3 register 11-Type ID Match 4 register If more than one Type ID is matched only one of them is indicated with priority 4 down to 1. With RX checksum offloading enabled: 00-Neither the IP header checksum nor the TCP/UDP checksum was checked. 01-The IP header checksum was checked and was correct. Neither the TCP nor UCP checksum was checked. 10-Both the IP header and TCP checksum were checked and were correct. 11-Both the IP header and UDP checksum were checked and were correct.
21	VLAN tag detected - Type ID of 8100 _h . For packets incorporating the stacked VLAN processing feature, this bit will be set if the second VLAN tag has a Type ID of 8100 _h .
20	Priority tag detected - Type ID of 8100 _h and null VLAN identifier. For packets incorporating the stacked VLAN processing feature, this bit will be set if the second VLAN tag has a Type ID of 8100 _h and a null VLAN identifier.
19:17	VLAN priority - only valid if bit 21 is set. 000-Priority 0 (lowest) BK Background 001-Priority 1 BE Best Effort 010-Priority 2 EE Excellent Effort 011-Priority 3 CA Critical Applications 100-Priority 4 VI Video, <100 ms latency and jitter 101-Priority 5 VO Voice, <10 ms latency and jitter 110-Priority 6 IC Internetwork Control 111-Priority 7 (highest) NC Network Control
16	Canonical Format Indicator (CFI) bit - only valid if bit 21 is set.

Table 31-2. Word 0 and Word 1 Description of Each BD

Bit	Function
15	End of Frame - when set the buffer contains the end of a frame. If End of Frame is not set, then the only valid status bit is Start of Frame (bit 14).
14	Start of Frame - when set the buffer contains the start of a frame. If both bits 15 and 14 are set, the buffer contains a whole frame.
13	This bit has a different meaning depending on whether jumbo frames and ignore FCS mode are enabled (network_configuration[3], network_configuration[26]). If neither mode is enabled this bit will be 0. With jumbo frame mode enabled: Additional bit for length of frame (bit 13), that is concatenated with bits [12:0] With ignore FCS mode enabled and jumbo frames disabled: This indicates per frame FCS status as follows: 0 -Frame had good FCS 1 -Frame had bad FCS, but was copied to memory as ignore FCS is enabled.
12:0	These bits represent the length of the received frame, which may or may not include FCS depending on whether FCS discard mode is enabled (network_configuration[17] = 1). With FCS discard mode disabled: Least significant 12 bits for length of frame including FCS. If jumbo frames are enabled, these 12 bits are concatenated with bit 13 of the descriptor. With FCS discard mode enabled: Least significant 12 bits for length of frame excluding FCS. If jumbo frames are enabled, these 12 bits are concatenated with bit 13 of the descriptor.

When Descriptor Timestamp Capture mode is enabled, the following table identifies the added descriptor words.

Table 31-3. Word 3 and Word 4 Description for Receive BD

Bit	Function
Word 2	
31:30	Timestamp seconds [1:0] ^a
29:0	Timestamp nanoseconds [29:0] ^a
Word 3	
31:10	Unused
9:0	Timestamp seconds [11:2] ^a

a. The Timestamp mode is controlled using the RX BD control register (ETHx_rx_bd_control). The timestamp bits are written back to the last buffer descriptor of a frame only.

To receive frames, the receive buffer descriptors must be initialized by writing an appropriate address to bits [31:2] (or [31:3] for timestamp capture mode) in the Word 0 of each BD. Bit 0 must be written as '0'. Bit 1 is the wrap bit and indicates the last entry in the buffer descriptor list.

The start location of the receive buffer descriptor list must be written with the receive buffer queue base address before reception is enabled (ETHx_network_control[2]=1). When reception is enabled, any writes to the receive buffer queue base address register are ignored.

Note: Writing receive buffer queue base address register may require three AXI clock cycles to take effect. Therefore, reception cannot be enabled until three AXI clock cycles after receive buffer queue base address register is updated. This restriction need to be taken care by firmware.

The receive buffer queue pointer increments by two or four words after each buffer is used. It re-initializes to the receive buffer queue base address if any descriptor has its wrap bit set. When receive buffer queue base address register is read, it returns the current pointer position in the descriptor list, though this is only valid and stable when receive is disabled.

If the filter block indicates that a frame should be copied to memory, the receive data DMA operation starts writing data into the receive buffer. As receive buffers are used, the receive buffer manager sets bit 0 of the first word of the descriptor to '1' indicating the buffer has been used.

Software should search through the "Used" bits in the buffer descriptors to find out how many frames are received, checking the Start of Frame and End of Frame bits.

If full store and forward mode is configured, only good received frames are written to the receive buffers, so no fragments will exist in the local packet buffers due to MAC receiver errors. However, there is still possibility of fragments due to EMAC DMA errors; for example, used bit read on the second buffer of a multi-buffer frame.

If bit 0 in Word 0 of the receive BD is already set when the receive buffer manager reads the BD, then the buffer has already been used and cannot be used again until software has processed the frame and clear bit 0. In this case, the “buffer not available” bit in the receive status register is set and an interrupt triggered. The Receive Resource Errors statistics register is also incremented.

When EMAC DMA is configured in the full store and forward mode, it can be selected to indicate whether received frames should be automatically discarded when no buffer resource is available; that is, “Used” bit is set for all receive BDs. This feature is selected via bit 24 of the *ETHx_dma_config* register. By default, the received frames are not automatically discarded. If this feature is off, then received packets will remain stored in the RX packet buffer memory until system memory resource becomes available again. This may lead to packet buffer overflow if packets continue to be received and still bit 0 (“Used” bit) of the receive BD remains set. Note that after a “Used” bit is read, the receive buffer manager will re-read the receive BD every time a new packet is received.

When the Ethernet MAC DMA is configured for packet buffer mode, the upper bits of the data buffer address stored in bits [31:2] in the first word of each list entry can be dynamically altered in real-time without physically changing the system memory holding the list entry. This feature is useful if the destination must be selected based on CPU usage or other flow control hardware. It is achieved using a mux structure whereby it can be defined whether the upper four bits of the 32-bit data-buffer AXI address should come from the descriptor list entry or from a programmable register. See the Receive DMA Data Buffer Address Mask register for further details. Note that any changes to this register will be ignored while the Ethernet MAC DMA is processing a receive packet. It will only affect the next full packet to be written to system memory.

31.3.1.6 Transmit Buffers

Frames to be transmitted can be stored in one or more transmit buffers. Transmit frames can be between 1 and 1536 bytes long. Note that zero length buffers are allowed and the maximum number of buffers permitted for each transmit frame is 128.

The start addresses of each transmit buffer is stored in system memory in a list of transmit buffer descriptors located at the transmit buffer queue pointer. The base addresses of transmit BD list must be configured by software using the transmit buffer queue base address registers (*ETHx_transmit_q_ptr*, *ETHx_transmit_q1_ptr*, *ETHx_transmit_q2_ptr*).

Each buffer descriptor can be either of two or four words, depending on the configured BD mode, whereas word is defined as 32 bits. The first two words (Word 0 and Word 1) are used in both BD modes.

In Extended Buffer Descriptor mode (*ETHx_dma_config*[29] = 1), two BD words (Word 2 and Word 3) are added for timestamp capture if timestamp capture mode is enabled (*ETHx_tx_bd_control*[5:4] > 0_h). Therefore, Transmit BDs will be either of two or four words size and each BD will have the same size. To summarize,

- Each transmit BD must be of 64 bits when descriptor time capture mode is disabled
- Each transmit BD must be of 128 bits when descriptor time capture mode is enabled

The following description details Word 0 and Word 1 of TX BD. Word 0 of the each transmit BD is the start address of the transmit buffer and the Word 1 consist of transmit control and status bits. For the packet buffer DMA, the start location for each transmit buffer is a byte address, the bottom bits of the address being used to offset the start of the data from the data-word boundary (that is, bits 2, 1, and 0 are used to offset the address for 64-bit data paths).

Frames can be transmitted with or without automatic CRC generation. If it is configured to generate CRC automatically, pad bytes will also be automatically generated to take frames to a minimum length of 64 bytes. If it is not configured to generate CRC automatically (as defined in Word 1 of the transmit buffer descriptor), the frame is assumed to be at least 64 bytes long and pad bytes are not generated.

Table 31-4. Word 0 and Word 1 Description of Transmit Buffer Descriptors

Bit	Function
Word 0	
31:0	Byte address of buffer
Word 1	
31	Used – must be 0 for the EMAC to read data to the transmit buffer. The EMAC sets this to ‘1’ for the first buffer of a frame after it is successfully transmitted. Software must clear this bit before the buffer can be used again.
30	Wrap – marks last descriptor in transmit buffer descriptor list. This can be set for any buffer within the frame.

Table 31-4. Word 0 and Word 1 Description of Transmit Buffer Descriptors

Bit	Function
29	Retry limit exceeded, transmit error detected
28	Unused
27	Transmit frame corruption due to AXI error – set if an error occurs whilst midway through reading through reading transmit frame from the AXI, including RRESP/BRESP errors and buffers exhausted mid frame (if the buffers run out during transmission of a frame then transmission stops, FCS shall be bad and TX_ER asserted). Also set if single frame is too large for the transmit packet buffer memory size.
26	Transmit error detected.
25:24	Reserved
23	For Extended Buffer Descriptor Mode, this bit indicates a timestamp is captured in the BD. Otherwise it is unused.
22:20	Transmit IP/TCP/UDP checksum generation offload errors: 000 – No error 001 – The packet was identified as a VLAN type, but the header was not fully complete, or had an error in it. 010 – The packet was identified as a SNAP type, but the header was not fully complete, or had an error in it. 011 – The packet was not of an IP type, or the IP packet was invalidly short, or the IP was not of type IPv4/IPv6 100 – The packet was not identified as VLAN, SNAP, or IP. 101 – Unsupported packet fragmentation occurred. For IPv4 packets, the IP checksum was generated and inserted. 110 – Packet type detected was not TCP or UDP, TCP/UDP checksum was therefore not generated. For IPv4 packets, the IP checksum was generated and inserted. 111 – A premature end of packet was detected and the TCP/UDP checksum could not be generated.
19:17	Reserved. must be set to 3'b000
16	No CRC to be appended by MAC. When set this implies that the data in the buffers already contains a valid CRC and hence no CRC or padding is to be appended to the current frame by the MAC. This control bit must be set for the first buffer in a frame and will be ignored for the subsequent buffers of a frame. Note that this bit must be “0” when using the transmit IP/TCP/UDP checksum generation offload, otherwise checksum generation and substitution will not occur. Note this bit must also be “0” when TX Partial Store and Forward mode is active.
15	Last buffer; when “1”, this bit will indicate the last buffer in the current frame is reached.
14	Reserved
13:0	Length of buffer.

When Descriptor Timestamp Capture mode is enabled, the following table identifies the added descriptor words.

Table 31-5. Word 3 and Word 4 of TX BD When Timestamp Capture Mode Enabled

Bit	Function
Word 2	
31:30	Timestamp seconds [1:0] ^a
29:0	Timestamp nanoseconds [29:0] ^a
Word 3	
31:10	Unused
9:0	Timestamp seconds [11:2] ^a

a. The Timestamp mode is controlled using the TX BD register (*ETHx_tx_bd_control*). After transmission the timestamp bits are written back only to the first buffer descriptor.

To transmit frames, the buffer descriptors must be initialized by writing the start address of the buffers to bits [31:0] in the first word (Word 0) of each descriptor.

Word 1 of the transmit buffer descriptor must be initialized with control information that indicates the length of the

frame, whether the MAC is required to append CRC and whether the buffer is the last buffer in the frame.

After transmission, the status bits of Word 1 of the first BD for a frame are updated by EMAC along with the “Used” bit. “Used” bit is written to ‘1’, after the frame is transmitted. Bits

[29:20] indicate various transmit error conditions. Bit 23 indicates a valid timestamp is captured in the BD. Bit 30 is the "Wrap" bit, which can be set for any buffer within a frame. If no wrap bit is encountered the queue pointer continues to increment to fetch next BD.

The transmit buffer queue base address register can only be updated whilst transmission is disabled or halted; otherwise any attempted write will be ignored.

- When transmission is halted the transmit buffer queue pointer will maintain its value. Therefore, when transmission is restarted the next descriptor read from the queue will be from the frame immediately after the last successfully transmitted frame.
- When transmit is disabled (*ETHx_network_control*[3] = 0), the transmit buffer queue pointer resets to point to the address indicated by the transmit buffer queue base address register.

Note that disabling receive does not have the same effect on the receive buffer queue pointer.

After the transmit queue is initialized, transmit is activated by writing to the transmit start bit (*ETHx_network_control*[9]). Transmit is halted when a buffer descriptor with its "Used" bit set is read or a transmit error occurs, or by writing to the transmit halt bit of the network control register (*ETHx_network_control*[10]). Transmission is suspended if a pause frame is received while the pause enable bit is set in the network control register (*ETHx_network_control*[13]). Rewriting the start bit while transmission is active is allowed. *transmit_go* bit (*ETHx_transmit_status*[3]) reset under following conditions:

- Transmit is disabled.
- A buffer descriptor with its ownership bit set is read.
- Bit 10, *tx_halt_clk*, of the network control register is written.
- There is a transmit error such as too many retries or a transmit under run.

When *transmit_go* is cleared, DMA will stop to fetch new packet from system memory and EMAC will not complete transmission until packet buffer is empty.

To set *transmit_go* write to bit 9, *tx_start_clk*, of the *ETHx_network_control*. Transmit halt does not take effect until any ongoing transmit finishes.

If the transmit BD list is incorrectly set up, for example a "used" bit set is read mid-way through a multi buffer frame, transmission will stop. If cut-through is in operation and the MAC has actually started transmitting the frame that has its used bit set, the MAC treats it as a transmit error, and asserts *tx_er* truncates the frame and corrupts the FCS.

DMA Burst

When performing data transfers, the burst length used can be programmed using bits [4:0] of the DMA configuration

register. Either single accesses (burst length = 1) or incrementing bursts of up to 4 can be used.

When there is sufficient space and enough data to be transferred, the burst of programmed length will be used. If there is not enough data or space available, for example when at the end of a packet or buffer, burst lengths of less than the programmed burst length value will be issued. Single accesses will be used when a 4 Kbyte boundary will be crossed by the burst to not violate the AXI specification.

EMAC DMA can also be configured to pad the remaining bursts at the end of a buffer to the programmed burst length value available via bits 26 and 25 of the *ETHx_dma_config* register. Bit 26 will control the transactions for TX and bit 25 for RX. For RX, the data to burst is padded with "0"s up to the burst boundary defined by burst length. For TX, the extra data that is read is ignored by the DMA. This feature is included for performance reasons when AXI slaves that are being accessed by EMAC perform better when accessed using fixed length bursts.


EMAC DMA will not terminate fixed length bursts early if receive/transmit operation is disabled by writing to *ETHx_network_control* register bit 2/3.

31.3.1.7 DMA Packet Buffer

The packet buffer DMA mode allows multiple packets to be buffered in both transmit and receive directions and allows the DMA to withstand variable levels of access latencies on the AXI fabric. Using packet buffers, AXI bandwidth has been used most efficiently in the device.

Figure 31-2 illustrates the structure of the EMAC data paths.

Figure 31-2. EMAC Data Path Structure



In the transmit direction, the DMA will continue to fetch packet data up to a limit of 256 packets, or until the TX Packet Buffer Memory is full. In the receive direction, if the RX Packet Buffer Memory becomes full, then an overflow will occur. An overflow will also occur if the limit of 256 packets is reached.

Transmit Packet Buffer

The transmit packet buffer (TX packet buffer) will continue attempting to fetch frame data from the system memory until the TX packet buffer memory is full or up to a limit of 256 packets, at which point it will attempt to maintain its full level.

To accommodate the status and statistics associated with each frame, two words per packet are reserved at the end of the packet data. If the packet was bad and requires to be dropped, the status and statistics are the only information held on that packet. Storing the status in the packet buffer is required to decouple the TX DMA interface of the buffer from the MAC transmitter interface, to update the MAC status/statistics and to generate interrupts in the order in which the packets that they represent were fetched from the system memory.

If any errors occur on the AXI whilst reading the transmit frame, then fetching of packet data from memory is halted. The MAC transmitter will continue to fetch packet data,

thereby emptying the TX packet buffer memory and allowing any good non-erroneous frames to be transmitted successfully. When these are fully transmitted, the status/statistics for the erroneous frame will be updated and software will be informed via an interrupt that an AXI error occurred. This way, the error is reported in the correct packet order.

The TX packet buffer will only attempt to read more frame data from the system memory when space is available in the TX packet buffer memory. If space is not available, it must wait until a packet fetched by the MAC transmitter completes transmission and is subsequently removed from the TX packet buffer memory. Note that if full store and forward mode is active and if a single frame is fetched that is too large for the TX packet buffer memory, the frame is flushed, and the TX DMA halted with an error status. This is because a complete frame must be written into the TX packet buffer memory before transmission can begin (If frame is split into multiple buffers, DMA will measure the buffer length against the available TX packet buffer room to decide if it will keep fetching).

In full store and forward mode, after the complete transmit frame is written into the TX packet buffer memory, a trigger is sent across to the MAC transmitter, which will then begin reading the frame from the TX packet buffer memory.

Because the whole frame is present and stable in the TX packet buffer memory, an underflow of the MAC transmitter is not possible. The frame is kept in the TX packet buffer memory until notification is received from the MAC transmitter that the frame data has either been successfully transmitted or can no longer be re-transmitted. When this notification is received the frame is flushed from TX packet buffer memory to make room for a new frame to be fetched from system memory. The frame is removed from the packet buffer on the fly after the frame is successfully transmitted.

In partial store and forward mode, a trigger is sent across to the MAC transmitter as soon as sufficient packet data is available in the internal buffer, which will then begin fetching the frame from the packet buffer memory. Later, if MAC transmitter is fetching data from the packet buffer faster than the DMA can fill it, an underflow of the transmitter will occur. In this case, the transmission is terminated early, and the packet buffer is flushed. Transmission can only be restarted by writing to the transmit start bit.

Receive Packet Buffer

The Receive Packet Buffer (RX Packet Buffer) stores frames from the MAC receiver along with their status and statistics. Frames with errors are flushed from the RX packet buffer memory, while good frames are pushed onto the AXI master interface.

When programmed in full store and forward mode, if the frame has an error the frame data is immediately flushed from the RX packet buffer memory allowing subsequent frames to use the freed-up space. The status and statistics for bad frames are still used to update the EMAC registers.

To accommodate the status and statistics associated with each frame, up to two words (one being for descriptor timestamp capture when enabled) per packet are reserved at the end of the packet data. If the packet was bad and requires to be dropped, the status and statistics are the only information held on that packet.

The RX packet buffer will also indicate a full condition such that an overflow condition can be detected. If this occurs, subsequent packets will be dropped, and an RX overflow interrupt is raised.

For full store and forward, the RX DMA will only begin packet fetches when the status and statistics for a frame are available. If the frame has a bad status due to a frame error, the status and statistics are passed onto the EMAC registers. If the frame has a good status, the information is used to read the frame from the RX packet buffer memory and burst onto the AXI master interface using the DMA buffer management protocol.

If partial store and forward mode is active, the DMA will begin fetching the packet data before the status is available. As soon as the status becomes available, the DMA will fetch this information before continuing to fetch the remainder of the frame. When the last frame data has been transferred to

the buffer in system memory, the status and statistics are updated to the EMAC registers.

31.3.1.8 Priority Queuing in EMAC DMA

Ethernet MAC supports three transmit and receive priority queues. Each queue has an independent list of buffer descriptors pointing to separate data streams. By default, each queue is active. Queues can be disabled by setting bit 0 of the Transmit or Receive Buffer Queue Base Address register. Note that at least one queue must always remain enabled and only the top indexed queues may ever be disabled. For example, if only two queues are being used, the user would disable Queue 2 by setting bit 0 of the Transmit or Receive Buffer Queue Base Address register.

In the transmit direction, higher priority queues are always serviced before lower priority queues. This strict priority scheme requires the user to ensure that high-priority traffic is constrained such that lower priority traffic will have the required bandwidth. The DMA will determine the next queue to service by initiating a sequence of buffer descriptor reads interrogating the ownership bits of each.

The buffer descriptor corresponding to the highest priority queue is read first. If the ownership bit of this descriptor is set, then the DMA will progress to reading the second highest priority queue's descriptor. If ownership bit of the second highest priority queue is also set, then the DMA will read the third highest priority queue's descriptor. If all the descriptors return an ownership bit set, then a resource error occurs, an interrupt is generated, and transmission is automatically halted.

Transmission can only be restarted by setting the START bit in the ETHx_network_control register. The DMA will identify the highest available queue to transmit from when the START bit in the ETHx_network_control register is written to and the TX is in a halted state, or when the last word of any packet has been fetched from external memory. The transmit DMA will maximize the effectiveness of priority queuing by ensuring that high priority traffic be transmitted as early as possible after being fetched from external memory.

For each queue, there is an associated Transmit Buffer Queue Base Address register. For the lowest priority queue (or the only queue when a single queue is selected), the Transmit Buffer Queue Base Address is located at offset address of 0x1C. For all other queues, the Transmit Buffer Queue Base Address registers are located at sequential addresses starting at offset 0x440.

In the receive direction each data packet is written to the internal packet buffer in the order that it is received. For each queue, there is an independent set of receive buffers. Therefore, a separate Receive Buffer Queue Base Address register for each queue is allocated. For the lowest priority queue, the Receive Buffer Queue Base Address register is located at address offset of 0x18. For all other queues,

these registers are located at sequential addresses starting at the address offset of 0x480. Every received packet will pass through a programmable screening algorithm, which will allocate that frame to a specific queue. The user interface to the screener is via two banks of programmable registers – ETHx_screener_type_1 and ETHx_screener_type_2.

Screener type 1 registers allow routing the received frames based on IP and UDP fields extracted from the received frames. Specifically, these fields are DS (differentiated services field of IPv4 frames), TC (traffic class field of IPv6 frames), and/or the UDP destination port. These fields are compared against the values stored in each of the screener type 1 match registers. If the result of this comparison is positive, then the received frame is routed to the priority queue specified in that screener type 1 register. TVII-B-H supports 16 screener type 1 registers.

Screener type 2 match registers operate independent of screener type 1 registers and offer additional match capabilities, extending the capabilities into vendor specific protocols. The type 2 screening allows configuring a screen, which is a combination of all or any of the following comparisons:

- An Enabled VLAN Priority. A VLAN priority match will be performed if the VLAN priority enable is set. The extracted priority field in the VLAN header is compared against three bits within the screener type 2 register.
- An Enabled EtherType. The ethertype field inside the screener type 2 register maps to one of eight ethertype match registers. The extracted ethertype is compared against the ethertype register designated by this ethertype field.
- An enabled Field Compare A.
- An enabled Field Compare B.
- An enabled Field Compare C.

Compare A, B, and C fields of the screener type 2 match register are pointers to a pool of up to 32 compare registers. When enabled, the compare is true if the data at the OFFSET into the frame ANDed with the MASK value (if the mask is enabled) is equal to the COMPARE value. Either a 16-bit comparison or a 32-bit comparison is done. This selection is made via a control bit in the associated compare word1. If a 16-bit comparison is selected, then a 16-bit mask is also available to the user to select which bits should be compared. If the 32-bit compare option is selected, then no mask is available. The byte at the OFFSET number of bytes from the index start is compared through bits 7:0 of the configured VALUE. The byte at the OFFSET number of bytes + 1 from the index start is compared through bits 15:8 of the configured VALUE, and so on. The OFFSET can be configured to be from 0 to 127 bytes from either the start of the frame, the byte following the ethertype field, the byte following the end of the IP header (IPv4 or IPv6), or the byte following the end of the TCP/UDP header. Note that the logic to decode the IP header or the TCP/UDP header is

reused from the TCP/UDP/IP checksum offload logic and has the same restrictions on use (the main limitation is that IP fragmentation is not supported). See [Checksum Offload for IP, TCP, and UDP on page 558](#) for further details. The Compare Register field points to a single pool of 32 compare registers. Compare A, B, and C use a common set of compare registers.

Note compare A, B, and C together allow matching an arbitrary 48 bits of data; therefore, they can be used to match a MAC address.

All enabled comparisons are ANDed together to form the overall type 2 screener match. TVII-B-H supports 16 screener type 2 registers.

Each screener register is programmable. Although it is not recommended, it is possible that more than one screener register can be programmed to match against a single frame. If this happens, consider the following:

- If a received frame matches multiple screeners of the same type, then the frame will route to the queue mapped by the screener located at the lowest numeric APB address. For example, if screener type 2 #0 and screener type 2 #1 both match, then the frame will route to the queue identified in bits [3:0] of the screener type 2 #0 register.
- If a received frame matches a type 2 screener and a type 1 screener, then the type 1 screener will take precedence.

When a screener is matched, the received frame will be routed to a queue defined inside bits 3:0 of the screener register. Unmatched frames are routed to queue 0.

The interrupt outputs from the Ethernet MAC match the number of supported priority queues. Only Ethernet MAC DMA-related events are reported using the individual interrupt outputs, because the Ethernet MAC can relate these events to specific queues. All other events generated within the Ethernet MAC are reported in the interrupt associated with the lowest priority queue (Queue 0). For the lowest priority queue, the Interrupt Status register is located at offset address 0x024. For all other priority queues, this register is located at sequential offset addresses starting at 0x400.

31.3.2 Transmit Scheduling Algorithm

When multiple priority queues are selected, the transmit scheduler is automatically included in the design and is responsible for selecting the next queue to be serviced from the attached DMA. There are four scheduling algorithms available to the user; with some exceptions detailed further below, the user can select one of the four modes for each queue.

31.3.2.1 802.1Qav Support - Credit Based Shaping

A credit-based shaping algorithm is available on the two highest priority active queues and is defined in 802.1Qav: Forwarding and Queuing Enhancements for Time-Sensitive Streams. Traffic shaping is enabled via the ETHx_CBS_control register (0x4bc) or the ETHx_tx_sched_ctrl register (offset 0x580). These two registers are aliased, so updating one register will automatically update the other. Note that it is permitted to enable CBS only on the second-highest-priority queue and not on the highest, in which case the highest-priority queue would always take precedence.

Enabling CBS on a queue will enable a counter, which stores the amount of transmit 'credit', measured in bytes, that a queue has. A queue may only transmit if it has non-negative credit.

If a queue has data to send but is held off from doing as another queue is transmitting, then credit will accumulate in the credit counter at the rate defined in the IdleSlope register for that queue. IdleSlope is the rate of change of credit when waiting to transmit and must be less than the value of the portTransmitRate. When this queue is transmitting, the credit counter is decremented at the rate of sendSlope, which is defined as the portTransmitRate - IdleSlope. A queue can accumulate negative credit when transmitting, which will hold off any other transfers from that queue until credit returns to a non-negative value. No transfers are halted when a queue's credit becomes negative; it will accumulate negative credit until the transfer completes.

To ensure that the CBS scheduling is completely accurate, a single transmit buffer should be used per Ethernet frame (rather than multi-buffer transmit frames).

31.3.2.2 Fixed Priority

Any of the active queues can be selected as fixed priority; this is the default mode of operation for all queues. The queue index is used as the priority, where a higher index will have a higher priority than a lower index. The scheduler will always attempt to transmit from fixed-priority queues with the highest priority. This means that a fixed-priority queue with a high queue index will always take precedence over a priority queue with a lower index.

31.3.2.3 Deficit Weighted Round Robin (DWRR)

Any of the active queues can be selected as DWRR. If DWRR is required, then at least two of the active queues should be selected as DWRR. It should not be used in conjunction with ETS, as both algorithms operating together would make little practical sense. A DWRR-enabled queue has lower priority than a CBS-enabled queue or a fixed-priority queue with a higher index.

The DWRR algorithm works by scanning all non-empty queues in sequence. Each queue is allocated a 'deficit counter' and an 8-bit weighting (or quantum) value. The value of the deficit counter is the maximum number of bytes that can be sent at the current time. If the deficit counter of the scanned queue is greater than the length of the packet waiting for transmission, then the packet will be transmitted and the value of the deficit counter is decremented by the packet size. If it is not greater, the scheduler will skip to the next DWRR-enabled queue. If there is insufficient credit to transmit, the queue is simply skipped. If the queue is empty, the value of the deficit counter is reset to '0'. If all queues have insufficient credit, at each tx_clk cycle, every queue's deficit counter is incremented by its quantum value until a queue's deficit counter obtains sufficient credit to transmit its first queued frame. The higher the quantum value chosen, the quicker the deficit counter will reach the required value. If all DWRR queues have the same weighting, then all queues will be granted the same overall bandwidth. The weighting value is stored in four programmable registers starting at offset 0x590. See the register descriptions for further details.

Note that if fixed-priority queues are to be used in conjunction with DWRR, the fixed-priority queues must be at a higher index value than the DWRR queues. A consequence of this is that the enabled DWRR queues will form a contiguous set of queues starting from queue 0.

If CBS is also used in conjunction with DWRR, the DWRR queues will share the remaining bandwidth after the CBS allocation is deducted.

Note: Transmit cut-thru (Partial Store and Forward Mode) should not be enabled if the transmit scheduler is used.

31.3.2.4 Enhanced Transmission Selection (ETS)

The ETS algorithm is defined in 802.1Qaz: Enhanced Transmission Selection for Bandwidth Sharing between Traffic, and allows traffic on specific queues to be bandwidth limited. Any of the active queues can be selected as ETS. If ETS is required, then at least two of the active queues should be selected as ETS. It should not be used in conjunction with DWRR as both algorithms operating together would make little practical sense. An ETS-enabled queue has a lower priority than a CBS-enabled queue or a fixed-priority queue with a higher index.

For each ETS-enabled queue, the user should program the bandwidth requirement for each queue as a percentage of the total bandwidth (an 8-bit register is used and the sum of values programmed should not exceed decimal 100). This will be the maximum bandwidth to be granted to that queue. The actual scheduling algorithm operates in a round-robin style from lowest indexed queues up to the highest indexed queue in sequence. The bandwidth allocation percentage is stored in four programmable registers starting at offset

0x590 – these are the same registers used for DWRR. See the register descriptions for further details.

If CBS is also used in conjunction with ETS, the sum of the ETS queue percentages should equal the remaining bandwidth after the CBS allocation is deducted.

Note: Transmit cut-thru (Partial Store and Forward Mode) should not be enabled if the transmit scheduler is used.

31.3.3 MAC Transmitter

The MAC transmitter operates in full duplex and transmits frames in accordance with the Ethernet IEEE 802.3 standard.

A small input buffer receives data through the DMA, which, depending on the `data_bus_width` control bits in the `ETHx_network_config` register, will extract data in 128-bit form. All subsequent processing before the final output is performed in bytes.

Transmit data can be output using one of four PHY interface – MII, RMII, GMII, or RGMII.

Frame assembly starts by adding preamble and the start frame delimiter (SFD). Data is taken from the TX packet buffer a word at a time. If necessary, padding is added to take the frame length to 60 bytes. CRC is calculated using an order 32-bit polynomial. This is inverted and appended to the end of the frame taking the frame length to a minimum of 64 bytes. If the “No CRC” bit (bit 16) is set in the second word (Word 1) of the last buffer descriptor of a transmit frame neither pad nor CRC are appended.

In full duplex mode, frames are transmitted immediately. Back-to-back frames are transmitted 96-bit times apart to guarantee the Interpacket Gap.

In all modes of operation, if the TX DMA under runs, a bad CRC is automatically appended using the same mechanism as jam insertion and the TX_ER signal is asserted. For a properly configured system this should never happen, as the complete frame is buffered in TX Packet Buffer Memory.

31.3.4 MAC Receiver

The MAC Receiver block checks for valid preamble, FCS, alignment, and length, presents received frames to the DMA and, stores the frame destination address for use by the address checking block.

Ethernet frames are normally stored in the receive buffer in the AXI memory complete with the FCS. Setting the `fcs_remove` bit in the network configuration register (`ETHx_network_config[17]`) causes frames to be stored without their corresponding FCS. The reported frame length field is reduced by four bytes to reflect this operation.

The MAC Receive block signals to the register block to increment the alignment, CRC (FCS), short frame, long

frame, jabber, or receive symbol errors when any of these exception conditions occur.

If bit 26 is set in the `ETHx_network_config` register, CRC errors will be ignored, and CRC erroneous frames will not be discarded, though the `ETHx_fcs_errors` statistics register will still be incremented. Additionally, if configured to use the DMA and not enabled for jumbo frames mode, then bit 13 of the receive buffer descriptor word 1 will be updated to indicate the FCS validity for the particular frame. This is useful for applications such as EtherCAT whereby individual frames with FCS errors must be identified.

Received frames can be checked for length field error by setting the `length_field_error_frame_discard` bit of the `ETHx_network_config` register (bit 16). When this bit is “1”, the receiver compares a frame’s measured length with the length field (bytes 13 and 14) extracted from the frame. The frame is discarded if the measured length is shorter. This checking procedure is for received frames between 64 bytes and 1518 bytes in length.

Each discarded frame is counted in the 10-bit length field `ETHx_fcs_errors` statistics register. Frames where the length field value is greater than or equal to 0600_h (1536) will not be checked.

31.3.5 Checksum Offload for IP, TCP, and UDP

Ethernet IP can be programmed to perform IP, TCP, and UDP checksum offloading in both receive and transmit directions, which is enabled by setting bit 24 in the `ETHx_network_config` register for receive and bit 11 in the `ETHx_dma_config` register for transmit.

IPv4 packets contain a 16-bit checksum field, which is the 16-bit 1’s complement of the 1’s complement sum of all 16-bit words in the header. TCP and UDP packets contain a 16-bit checksum field, which is the 16-bit 1’s complement of the 1’s complement sum of all 16-bit words in the header, the data, and a conceptual IP pseudo header.

To calculate these checksums in software requires processing each byte of the packet. For TCP and UDP this can use a large amount of processing power. Offloading the checksum calculation to hardware can result in significant performance improvements.

For IP, TCP, and UDP checksum offload to be useful, the operating system containing the protocol stack must be aware that this offload is available so that it can make use of the fact that the hardware can either generate or verify the checksum.

31.3.5.1 Receive Checksum Offload

When receive checksum offloading is enabled in the Ethernet IP, the IPv4 header checksum is checked as per RFC791, where the packet meets the following criteria:

- If present, the VLAN header must be four octets long and the CFI bit must not be '1'. (for receive, one stacked VLAN is supported.)
- Encapsulation must be RFC 894 Ethernet Type Encoding or RFC 1042 SNAP Encoding or PPPoE Encoding.
- IPv4 packet.
- IP header is of valid length.
- IP options are supported.

The Ethernet IP also checks the TCP checksum as per RFC 793, or UDP checksum as per RFC 768, if the following criteria are met:

- IPv4 or IPv6 packet.
- IP options and all IPv6 extension headers (hop-by-hop, routing, and destination) are supported (except fragmentation headers).
- Good IP header checksum (if IPv4).
- IP fragmentation is not supported. (If a packet is fragmented, then the checksum will not be checked.)
- TCP or UDP packet.

When an IP, TCP, or UDP frame is received, the receive buffer descriptor gives an indication if the Ethernet IP was able to verify the checksums. There is also an indication if the frame had SNAP encapsulation. These indication bits will replace the Type ID match indication bits when the receive checksum offload is enabled. For details of these indication bits, refer to the [Table 31-1](#).

If any of the checksums are incorrectly verified by the Ethernet IP, the packet is discarded, and the appropriate statistics counter is incremented.

31.3.5.2 Transmit Checksum Offload

The transmitter checksum offload is only available if Ethernet IP is configured to use the DMA in packet buffer mode, and full store and forward mode is enabled. This is because the complete frame to be transmitted must be read into the TX Packet Buffer Memory before the checksum can be calculated and written back into the headers at the beginning of the frame.

Transmitter checksum offload is enabled by setting bit 11 in the `ETHx_dma_config` register. When enabled, it will monitor the frame as it is written into the TX Packet Buffer Memory to automatically detect the protocol of the frame. Protocol support is identical to the receiver checksum offload.

For transmit checksum generation and substitution to occur, the protocol of the frame must be recognized, and the frame must be provided without the FCS field, by making sure that bit 16 of the transmit descriptor Word 1 is clear (VLAN tagged frames will be recognized but stacked VLAN tagged frames will not be recognized). If the frame data already had

the FCS field, this would be corrupted by the substitution of the new checksum fields.

If these conditions are met, the transmit checksum offload engine will calculate the IP, TCP, and UDP checksums as appropriate. When the full packet is completely written into TX Packet Buffer Memory, the checksums will be valid and the relevant memory locations will be updated for the new checksum fields as per standard IP/TCP and UDP packet structures.

If the transmitter checksum engine is prevented from generating the relevant checksums, bits [22:20] of the Transmit Buffer Descriptor Entry will be updated to identify the reason for the error. Note that the frame will still be transmitted but without the checksum substitution; this is because, the reason the substitution did not occur is typically because the protocol was not recognized.

31.3.6 Jumbo Frame Support

The jumbo frames enable bit in the `ETHx_network_config` register (bit 3) allows EMAC to receive jumbo frames up to a software configurable number of bytes in size. This operation is not part of IEEE Std 802.3 specification and is by default disabled. When jumbo frames are enabled, frames received with a frame size greater than the configured value are discarded.

Oversize frame received register (`ETHx_excessive_rx_length`) will count the number of very long frames received.

The jumbo frames maximum length can be controlled using the Jumbo-Frame Maximum Length register (`ETHx_jumbo_max_length`). In EMAC, the maximum length of jumbo frame is 1536 bytes.

The jumbo frames maximum length can be controlled using the Jumbo-Frame Maximum Length register (`ETHx_jumbo_max_length`). Its default value is 1536 bytes.

- If jumbo frame is enabled (`ETHx_network_config[3]`):
 - The `ETHx_jumbo_max_length` register has default value 1536; the user does not need to set it for a 1536 bytes transfer.
 - The user can modify the `ETHx_jumbo_max_length` register, and the maximum length of the frame received is determined by this register.
 - The value of `ETHx_receive_1536_byte_frames` does not matter.
- If jumbo_frame is disabled:
 - If `ETHx_network_config.receive_1536_byte_frames` is set, maximum length will be 1518.
 - If `ETHx_network_config.receive_1536_byte_frames` is not set, maximum length will be 1500.
 - The value of `ETHx_jumbo_max_length` does not matter.

In EMAC, the maximum length of jumbo frame is 1536 bytes. The TX and RX buffer size does not allow longer frames to be transmitted or received in full store and forward mode.

31.3.7 MAC Filtering Block

The filter block determines which frames should be written to the EMAC DMA.

Whether a frame is passed depends on what is enabled in the *ETHx_network_config* register, contents of the Specific Address (*ETHx_spec_addi_top*, *ETHx_spec_addi_bottom*), Type ID Match (*ETHx_spec_typei*), and Hash (*ETHx_hash_top*, *ETHx_hash_bottom*) registers and the frame's destination address and type fields.

Ethernet frames are transmitted one byte at a time, least significant bit first. The first six bytes of an Ethernet frame make up the destination address. The first bit of the destination address, which is the LSB of the first byte of the frame, is the group or individual bit. This is '1' for multicast addresses and '0' for unicast ones. The "all ones" address is the broadcast address and a special case of multicast.

The EMAC supports recognition of specific source or destination addresses. The number of specific source or destination address filters is four. Each specific address filter consists of two registers, Specific Address Bottom i register and Specific Address Top i register. Specific Address Bottom i register stores the first four bytes of the compare source or destination address. Specific Address Top i register contains the last two bytes of this address, a control bit to select between source or destination address filtering, and a 6-bit byte mask field to allow masking certain bytes during the comparison. The first filter (Filter 1) is slightly different to all other filters in that there is no byte mask. Instead address comparison against individual bits of Specific Address 1 register can be masked using the unique Specific Address Mask 1 register. The addresses stored in all filters can be specific (unicast), group (multicast), local, or universal.

The destination or source address of received frames is compared against the data stored in the Specific Address registers after they are activated. The addresses are deactivated at reset or when their corresponding Specific Address Bottom i register is written. They are activated when the corresponding Specific Address Top i register is written. If a receive frame address matches an active address, the frame is written to the DMA memory if used.

Frames may be filtered using the Type ID field for matching. Four Type ID Match registers exist and each can be enabled for matching by writing a "1" to the MSB (bit 31) of the respective register. When a frame is received, the matching is implemented as an OR function of the various types of match.

The content of each Type ID register (when enabled) is compared against the length/Type ID of the frame being received (for example, bytes 13 and 14 in non-VLAN and

non-SNAP encapsulated frames) and copied to system memory if a match is found. The encoded Type ID match bits (Word 0, bit 22, and bit 23) in the receive buffer descriptor status are set indicating which Type ID match register generated the match, if the receive checksum offload is disabled.

The reset state of the Type ID match registers is "0", hence each is initially disabled.

The following example illustrates the use of the address and Type ID match registers for a MAC address of 21:43:65:87:A9:CB:

Preamble	55h
SFD	D5h
DA (Octet 0 - LSB)	21h
DA (Octet 1)	43h
DA (Octet 2)	65h
DA (Octet 3)	87h
DA (Octet 4)	A9h
DA (Octet 5 - MSB)	CBh
SA (Octet 0 - LSB)	00*
SA (Octet 1)	00*
SA (Octet 2)	00*
SA (Octet 3)	00*
SA (Octet 4)	00*
SA (Octet 5 - MSB)	00*
Type ID (MSB)	43h
Type ID (LSB)	21h

* Contains the address of the transmitting device.

The sequence above shows the beginning of an Ethernet frame. Byte order of transmission is from top to bottom as shown. For a successful match to Specific Address 1, the following address match registers must be set up:

ETHx_spec_add1_bottom (0x088) 87654321h

ETHx_spec_add1_top (0x08C) 0000CBA9h

And for a successful match to the Type ID, the Specific Address 1 register must be set up:

ETHx_spec_type1 (0x0A8) 80004321h

31.3.7.1 Broadcast Address

Frames with the broadcast address of FFFFFFFFh are stored to memory only if the no_broadcast bit in the *ETHx_network_config* register is set to '0'.

31.3.7.2 Hash Addressing

The 64-bit Hash register (*ETHx_hash_top/ETHx_hash_bottom*) takes up two locations in the memory map. The least significant bits are stored in the *ETHx_hash_bottom* register and the most significant bits in the *ETHx_hash_top* register. The unicast hash enable and the multicast hash enable bits in the *ETHx_network_config* register enable the reception of hash matched frames. The destination address is reduced to a 6-bit index into the 64-bit Hash register using the following hash function ($\{ETHx\_hash\_top, ETHx\_hash\_bottom\}[2hash\_index[5:0]]$). The hash function is an XOR of every sixth bit of the destination address.

$hash\_index[05] = da[05] \wedge da[11] \wedge da[17] \wedge da[23] \wedge da[29] \wedge da[35] \wedge da[41] \wedge da[47]$

$hash\_index[04] = da[04] \wedge da[10] \wedge da[16] \wedge da[22] \wedge da[28] \wedge da[34] \wedge da[40] \wedge da[46]$

$hash\_index[03] = da[03] \wedge da[09] \wedge da[15] \wedge da[21] \wedge da[27] \wedge da[33] \wedge da[39] \wedge da[45]$

$hash\_index[02] = da[02] \wedge da[08] \wedge da[14] \wedge da[20] \wedge da[26] \wedge da[32] \wedge da[38] \wedge da[44]$

$hash\_index[01] = da[01] \wedge da[07] \wedge da[13] \wedge da[19] \wedge da[25] \wedge da[31] \wedge da[37] \wedge da[43]$

$hash\_index[00] = da[00] \wedge da[06] \wedge da[12] \wedge da[18] \wedge da[24] \wedge da[30] \wedge da[36] \wedge da[42]$

- $da[00]$ represents the least significant bit of the first byte received, that is, the multicast/unicast indicator, and $da[47]$ represents the most significant bit of the last byte received.
- If the hash index points to a bit that is set in the Hash register, then the frame will be matched according to whether the frame is multicast or unicast.
- A multicast match will be signaled if the multicast hash enable bit is set, $da[00]$ is '1' and the hash index points to a bit set in the Hash register.
- A unicast match will be signaled if the unicast hash enable bit is set, $da[00]$ is '1' and the hash index points to a bit set in the Hash register.
- To receive all multicast frames, the Hash register must be set with all '1' and the multicast hash enable bit must be set to '1' in the *ETHx_network_config* register.

31.3.7.3 Copy All Frames

If the copy all frames bit is set in the *ETHx_network_config* register then all frames (except those that are too long, too short, have FCS errors, or have RX_ER asserted during reception) will be copied to memory. Frames with FCS errors will be copied if bit 26 is set in the *ETHx_network_config* register.

31.3.7.4 Disable Copy of Pause Frames

Pause frames can be prevented from being written to memory by setting the disable copying of pause frames control bit 23 in the *ETHx_network_config* register. When set, pause frames are not copied to system memory regardless of the copy all frames bit, whether a hash match is found, a type ID match is identified, or if a destination address match is found.

31.3.7.5 VLAN Support

An Ethernet encoded 802.1Q VLAN tag appears as follows:

Tag Protocol Identifier (TPID) 16 bits	Tag Control Information (TCI) 16 bits
8100h	First 3 bits priority, then CFI bit, last 12 bits VID

The VLAN tag is inserted at the thirteenth byte of the frame adding an extra four bytes to the frame. To support these extra four bytes, the EMAC can accept frame lengths up to 1536 bytes by setting bit 8 in the *ETHx_network_config* register. If the VID (VLAN identifier) is null (000h) this indicates a priority-tagged frame.

The following bits in the receive buffer descriptor status word give information about VLAN tagged frames:

- Bit 21 set if receive frame is VLAN tagged (Type ID of 8100h).
- Bit 20 set if receive frame is priority tagged (Type ID of 8100h and null VID). If bit 20 is set bit 21 will also be set.
- Bit 19, 18, and 17 set to priority if bit 21 is set.
- Bit 16 set to CFI if bit 21 is set.

EMAC can be configured to reject all frames except VLAN-tagged frames by setting the discard non-VLAN frames bit in the *ETHx_network_config* register.

31.3.8 IEEE 1588 and IEEE 802.1AS Support

IEEE Std 1588 is a standard for precision time synchronization in local area networks. It works with the exchange of special Precision Time Protocol (PTP) frames. The PTP messages can be transported over IEEE Std 802.3/Ethernet, over Internet Protocol Version 4 (IPv4) or over Internet Protocol Version 6 (IPv6) as described in the annex of IEEE Std 1588-2008.

Most IEEE Std 1588 functionality can be implemented in software but for greatest accuracy, hardware assist is required to detect when PTP event messages pass the MII interface (clock timestamp point).

Synchronization between master and slave clocks is a two-stage process.

First, the offset between the master and slave clocks is corrected by the master sending a Sync frame to the slave with a follow-up frame containing the exact time the Sync frame was sent. Hardware assist modules at the master and slave side detect exactly when the Sync frame was sent by the master and received by the slave. The slave then corrects its clock to match the master clock.

Second, the transmission delay between the master and slave is corrected. The slave sends a delay request frame to the master, which sends a delay response frame in reply. Hardware assist modules at the master and slave side detect exactly when the delay request frame was sent by the slave and received by the master. The slave will now have enough information to adjust its clock to account for delay. For example, if the slave assumes zero delay the actual delay will be half the difference between transmit and receive time of the delay request frame (if transmit and receive times are equal) because the slave clock will be lagging the master clock by the delay time already.

For hardware assist, it is necessary to timestamp when Sync and Delay_Req messages are sent and received. The timestamp is taken when the message timestamp point passes the clock timestamp point. For Ethernet, the message timestamp point is the SFD and the clock timestamp point is the MII interface. (The IEEE Std 1588 specification refers to Sync and Delay_Req messages as event messages because these require time stamping. Follow up, delay response, and management messages do not require time stamping and are referred to as general messages.)

IEEE Std 1588 version 2 (IEEE Std 1588-2008) defines two additional PTP event messages. These are the peer delay request (Pdelay_Req) and peer delay response (Pdelay_Resp) messages. These messages are used to calculate the delay on a link. Nodes at both ends of a link send both types of frames (regardless of whether they contain a master or slave clock). The Pdelay_Resp message contains the time at which a Pdelay_Req was

received and is itself an event message. The time at which a Pdelay_Resp message is received is returned in a Pdelay_Resp_Follow_Up message.

IEEE Std 1588 version 2 introduces two kinds of transparent clocks, peer-to-peer (P2P) and end-to-end (E2E). Transparent clocks measure the transit time of event messages through a bridge and amend a correction field within the message to allow for the transit time. P2P transparent clocks additionally correct the delay in the receive path of the link using the information gathered from the peer delay frames. With P2P transparent clocks Delay_Req messages are not used to measure link delay. This simplifies the protocol and makes larger systems more stable.

The Ethernet MAC recognizes ten different encapsulations for PTP event messages:

- IEEE Std 1588 version 1 (UDP/IPv4 multicast)
- IEEE Std 1588 version 1 (UDP/IPv4 multicast with VLAN)
- IEEE Std 1588 version 2 (UDP/IPv4 multicast)
- IEEE Std 1588 version 2 (UDP/IPv4 multicast with VLAN)
- IEEE Std 1588 version 2 (UDP/IPv4 unicast)
- IEEE Std 1588 version 2 (UDP/IPv4 unicast with VLAN)
- IEEE Std 1588 version 2 (UDP/IPv6 multicast)
- IEEE Std 1588 version 2 (UDP/IPv6 multicast with VLAN)
- IEEE Std 1588 version 2 (Ethernet multicast)
- IEEE Std 1588 version 2 (Ethernet multicast with VLAN)

Note: IEEE Std 1588 version 1 (IEEE Std 1588-2002)

Unicast PTP frame recognition is enabled via bit 20 of the *ETHx_network_config* register. The unicast addresses themselves are programmable via the PTP Unicast IP Destination Address (*ETHx_rx_ptp_unicast/ETHx_tx_ptp_unicast*) register. The first holds the RX unicast IP destination address and the other, the TX unicast destination address. The PTP Unicast IP Destination Address register should only be changed when the unicast PTP frame recognition is disabled.

Example of a Sync frame in the IEEE Std 1588 version 1 format:

Preamble/SFD	55555555555555D5h
DA (Octets 0 - 5)	
SA (Octets 6 - 11)	
Type (Octets 12 - 13)	0800h
IP stuff (Octets 14 - 22)	
UDP (Octet 23)	11h
IP stuff (Octets 24 - 29)	
IP DA (Octets 30 - 32)	E00001h
IP DA (Octet 33)	81h or 82h or 83h or 84h

source IP port (Octets 34 - 35)
 dest IP port (Octets 36 - 37) 013Fh
 other stuff (Octets 38 - 42)
 version PTP (Octet 43) 01h
 other stuff (Octets 44 - 73)
 control (Octet 74) 00h
 other stuff (Octets 75 - 168)

Example of a Delay_Req frame in the IEEE Std 1588 version 1 format:

Preamble/SFD 55555555555555D5h
 DA (Octets 0 - 5)
 SA (Octets 6 - 11)
 Type (Octets 12 - 13) 0800h
 IP stuff (Octets 14 - 22)
 UDP (Octet 23) 11h
 IP stuff (Octets 24 - 29)
 IP DA (Octets 30 - 32) E00001h
 IP DA (Octet 33) 81h or 82h or 83h or 84h
 source IP port (Octets 34 - 35)
 dest IP port (Octets 36 - 37) 013Fh
 other stuff (Octets 38 - 42)
 version PTP (Octet 43) 01h
 other stuff (Octets 44 - 73)
 control (Octet 74) 01h
 other stuff (Octets 75 - 168)

For IEEE Std 1588 version 1 messages Sync and Delay_Req frames are indicated by the EMAC if the frames type field indicates TCP/IP, UDP protocol is indicated, the destination IP address is 224.0.1.129/130/131/132, the destination UDP port is 319, and the control field is correct. The control field is 00h for Sync frames and 01h for Delay_Req frames.

For IEEE Std 1588 version 2 messages, the type of frame is determined by the message type field in the first byte of the PTP frame. Whether a frame is version 1 or version 2 can be determined by the version PTP field in the second byte of version 1 and version 2 PTP frames.

In version 2 messages, Sync frames have a message type value of 0h, Delay_Req frames have 1h, Pdelay_Req frames have 2h, and Pdelay_Resp frames have 3h.

Example of a Sync frame in the IEEE Std 1588 version 2 (UDP/IPv4) format:

Preamble/SFD 55555555555555D5h
 DA (Octets 0 - 5)
 SA (Octets 6 - 11)
 Type (Octets 12 - 13) 0800h
 IP stuff (Octets 14 - 22)

UDP (Octet 23) 11h
 IP stuff (Octets 24 - 29)
 IP DA (Octets 30 - 33) E0000181h
 source IP port (Octets 34 - 35)
 dest IP port (Octets 36 - 37) 013Fh
 other stuff (Octets 38 - 41)
 message type (Octet 42) 00h
 version PTP (Octet 43) 02h

Example of a Pdelay_Req frame in the IEEE Std 1588 version 2 (UDP/IPv4) format:

Preamble/SFD 55555555555555D5h
 DA (Octets 0 - 5)
 SA (Octets 6 - 11)
 Type (Octets 12 - 13) 0800h
 IP stuff (Octets 14 - 22)
 UDP (Octet 23) 11h
 IP stuff (Octets 24 - 29)
 IP DA (Octets 30 - 33) E000006Bh
 source IP port (Octets 34 - 35)
 dest IP port (Octets 36 - 37) 013Fh
 other stuff (Octets 38 - 41)
 message type (Octet 42) 02h
 version PTP (Octet 43) 02h

Example of a Sync frame in the IEEE Std 1588 version 2 (UDP/IPv6) format:

Preamble/SFD 55555555555555D5h
 DA (Octets 0 - 5)
 SA (Octets 6 - 11)
 Type (Octets 12 - 13) 86DDh
 IP stuff (Octets 14 - 19)
 UDP (Octet 20) 11h
 IP stuff (Octets 21 - 37)
 IP DA (Octets 38 - 53) FF0X000000000181h
 source IP port (Octets 54 - 55)
 dest IP port (Octets 56 - 57) 013Fh
 other stuff (Octets 58 - 61)
 message type (Octet 62) 00h
 other stuff (Octets 63 - 93)
 version PTP (Octet 94) 02h

Example of a Pdelay_Req frame in the IEEE Std 1588 version 2 (UDP/IPV6) format:

Preamble/SFD	55555555555555D5h
DA (Octets 0 - 5)	
SA (Octets 6 - 11)	
Type (Octets 12 - 13)	86DDh
IP stuff (Octets 14 - 19)	
UDP (Octet 20)	11h
IP stuff (Octets 21 - 37)	
IP DA (Octets 38 - 53)	FF0200000000006Bh
source IP port (Octets 54 - 55)	
dest IP port (Octets 56 - 57)	013Fh
other stuff (Octets 58 - 61)	
message type (Octet 62)	03h
other stuff (Octets 63 - 93)	
version PTP (Octet 94)	02h

Example of a Sync frame in the IEEE Std 1588 version 2 (Ethernet multicast) format. For the multicast address 011B19000000h Sync and Delay_Req frames are recognized depending on the message type field, 00h for Sync, and 01h for Delay_Req:

Preamble/SFD	55555555555555D5h
DA (Octets 0 - 5)	011B19000000h
SA (Octets 6 - 11)	
Type (Octets 12 - 13)	88F7h
message type (Octet 14)	00h
version PTP (Octet 15)	02h

Example of a Pdelay_Req frame in the IEEE Std 1588 version 2 (Ethernet multicast) format is given here. These need a special multicast address so they can get through ports blocked by the spanning tree protocol. For the multicast address 0180C200000Eh Sync, Pdelay_Req, and Pdelay_Resp frames are recognized depending on the message type field, 00h for Sync, 02h for Pdelay_Req, and 03h for Pdelay_Resp:

Preamble/SFD	55555555555555D5h
DA (Octets 0 - 5)	0180C200000Eh
SA (Octets 6 - 11)	
Type (Octets 12 - 13)	88F7h
message type (Octet 14)	02h
version PTP (Octet 15)	02h

The EMAC contains a timestamp unit (TSU), which consists of a timer and registers to capture the time at which PTP event frames cross the message timestamp point. The registers are accessible through the EMAC's AHB slave interface. An interrupt is issued when a capture register is updated.

31.3.8.1 Support for Time Stamping and Timestamp Accuracy

The MAC has the responsibility of sampling the TSU timer value when the TX or RX SOF event of the frame passes the MII boundary. This event is an existing signal synchronous to MAC TX/RX clock domains. The MAC uses the sampled timestamp to insert the timestamp into transmitted PTP Sync frames (if one step sync feature is enabled), to pass to the Ethernet MAC's register block to capture the timestamp (TS) in registers, or to pass to the Ethernet MAC DMA to insert into TX or RX buffer descriptors. For each of these, the SOF event, which is captured in the TX and RX clock domains respectively, is synchronized to the TSU clock domain and the resulting signal is used to sample the TSU count value. This value will be kept stable for an entire frame, or specifically at least 64 TX/RX clock cycles, because the minimum frame size in Ethernet is 64 bytes and worst case is a transfer rate of one byte per cycle. It is used as the source for all the components within the Ethernet MAC that require the timestamp value. The SOF event must pass a clock boundary and there may be a degree of inaccuracy in the captured timestamp. The level of inaccuracy depends on the frequency of the TSU clock (clk_tsu). There will be no more than one clock cycle of inaccuracy.

In the best case, the SOF event (which is in the TX/RX clock domain) just meets the setup time of the TSU clock domain at the input to the first synchronization flip-flop. The captured TS is N+2, but it really should be N+1. In the worst case, the captured TS is also N+2, but it really should be N.

31.3.8.2 Single Step Time Stamping

Support of one step clock for TX Sync frames can be enabled by setting bit 24 in the *ETHx_network_control* register. In this mode, the timestamp field within the IEEE Std 1588 version 2 Sync frame, is replaced by the TSU timestamp value at the time the Sync frame SOF passes the MII interface. To use single step time stamping, the sampled timestamp must be stable before the point at which EMAC requires to insert the timestamp. This can be guaranteed by enforcing a rule that TSU clock (clk_tsu) is greater than one-eighth the frequency of TX clock (TX_CLK) or RX clock (RX_CLK).

31.3.8.3 Timestamp Capture Registers

Four 80-bit timestamp status registers capture the time at which PTP event frames are transmitted and received.

- *ETHx_tsu_ptp_rx_msb_sec*, *ETHx_tsu_ptp_rx_sec*, *ETHx_tsu_ptp_rx_nsec*
- *ETHx_tsu_ptp_tx_msb_sec*, *ETHx_tsu_ptp_tx_sec*, *ETHx_tsu_ptp_tx_nsec*
- *ETHx_tsu_peer_rx_msb_sec*, *ETHx_tsu_peer_rx_sec*, *ETHx_tsu_peer_rx_nsec*

- *ETHx_tsu_peer_tx_msb_sec*, *ETHx_tsu_peer_tx_sec*, *ETHx_tsu_peer_tx_nsec*

An interrupt is issued when these registers are updated.

31.3.8.4 Timestamp Capture in DMA Descriptors

The TX and RX timestamp can optionally be captured in an extended buffer descriptor when configured using bits [29:28] in the *ETHx_dma_config* register. The timestamp can be captured for a number of frame types (PTP event, PTP general, all frames, or none as defined in the *ETHx_tx_bd_control*/*ETHx_rx_bd_control* registers) and a bit within Buffer Descriptor Word 0/1 is used to indicate that the timestamp is present.

31.3.8.5 Controlling Timestamp Unit

The timer is implemented as a 102-bit register with the upper 48 bits counting seconds, the next 30 bits counting nanoseconds, and the lowest 24 bits counting sub-nanoseconds. The lower 54 bits roll over when they have counted to one second. An interrupt is generated when the seconds increment. The timer value can be read, written, and adjusted through the AHB slave interface. The timer is clocked with TSU clock (*clk_tsu*).

The two operation modes that control the way the timer varies over time are:

- Increment mode – Increments the timer by a fixed value every TSU clock.
- Alternative increment mode – Increments the timer by a fixed value for a fixed number of clocks, followed by an alternative increment value for a single clock. This is the legacy mode and is not recommended for use.

The timer count value can be compared to a programmable comparison value. For the comparison, the 48 bits of the seconds value and the upper 22 bits of the nanoseconds value are used. A signal is output from the core to indicate when the TSU timer count value is equal to the comparison value stored in the TSU timer comparison value registers (0x0DC, 0x0E0, and 0x0E4). An interrupt can be issued when the timer count value and the comparison value in the IEEE 1588 Timer Comparison Value registers (*ETHx_tsu_msb_sec_cmp*, *ETHx_tsu_sec_cmp* and *ETHx_tsu_nsec_cmp*) are equal. The interrupt can be enabled with bit 29 in the interrupt enable (*ETHx_int_enable*) register.

EMAC is designed to recognize Sync frames with both IEEE Std 1588 and IEEE Std 802.1AS addresses, and can support their frame recognition simultaneously.

31.3.8.6 Increment Mode

The amount by which the timer increments each clock cycle is controlled by the timer increment (*ETHx_tsu_timer_incr*) register. Bits [7:0] are the default increment value in

nanoseconds and an additional 16 bits of sub-nanosecond resolution are available using the timer increment *ETHx_sub_nsec* register (*ETHx_tsu_timer_incr_sub_nsec*). If the rest of the timer increment register is written with '0' the timer increments by the value in bits [7:0] and the value in timer increment *ETHx_sub_nsec* register, each clock cycle.

31.3.8.7 Alternate Increment Mode

Bits [15:8] of the timer increment register are the alternative increment value in nanoseconds and bits [23:16] are the number of increments after which the alternative increment value is used. If bits [23:16] are 00h, then the alternative increment value will never be used.

31.3.9 MAC 802.3 Pause Frame Support

The EMAC supports both hardware-controlled pause of the transmitter upon reception of a pause frame and hardware-generated pause frame transmission.

31.3.9.1 IEEE Std 802.3 Pause Frame Reception

Bit 13 of the *ETHx_network_config* register is the pause enable control for reception. If this bit is set, transmission will pause when a non-zero pause quantum frame is received.

If a valid pause frame is received, then the Receive Pause Quantum (*ETHx_pause_time*) register is updated with the new frame's pause time regardless of whether a previous pause frame is active. An interrupt (either bit 12 or 13 of the *ETHx_int_status* register) is triggered when a pause frame is received, but only if the interrupt is enabled. Pause frames received with non-zero quanta are indicated through the interrupt bit 12 of the Interrupt Status (*ETHx_int_status*) register. Pause frames received with zero quanta are indicated on bit 13 of *ETHx_int_status*.

After the Receive Pause Quantum (*ETHx_pause_time*) is loaded and the frame currently being transmitted is sent, no new frames are transmitted until the pause time reaches zero. The loading of a new pause time, and hence pausing of transmission, occurs because the EMAC is operating in full duplex mode. A valid pause frame is defined as having a destination address that matches either the address stored in Specific Address 1 register or the reserved address of 0180C2000001h. It must also have the MAC control frame Type ID of 8808h and the pause opcode of 0001h.

Pause frames that have FCS or other errors will be treated as invalid and will be discarded. IEEE Std 802.3 pause frames that are received after priority-based flow control (PFC) is negotiated will also be discarded. Valid pause frames received will increment the Pause Frames Received statistics register.

The Receive Pause Quantum register decrements every 512-bit times when transmission stops. For test purposes, the retry test bit can be set (bit 12 in *ETHx_network_config*), which causes the Receive Pause Quantum register to decrement every TX_CLK cycle after transmission has stopped.

The interrupt (bit 13 in the *ETHx_int_status* register) is asserted whenever the Receive Pause Quantum register decrements to zero and it is enabled. This interrupt is also set when a zero quantum pause frame is received.

31.3.9.2 IEEE Std 802.3 Pause Frame Transmission

Automatic transmission of pause frames is supported through the transmit pause frame bits of the *ETHx_network_control* register. If either bit 11 or bit 12 of *ETHx_network_control* is set to '1', an IEEE Std 802.3 pause frame will be transmitted, provided the MAC transmitter is enabled (bit 3) in the *ETHx_network_control* register.

Pause frame transmission will happen immediately if transmit is inactive or if transmit is active between the current frame and the next frame due to be transmitted.

Transmitted pause frames comprise of the following:

- A destination address of 0180C2000001h
- A source address taken from Specific Address 1 register
- A Type ID of 8808h (MAC control frame)
- A pause opcode of 0001h
- A pause quantum register
- Fill of 00h to take the frame to minimum frame length
- Valid FCS

The pause quantum used in the generated frame will depend on the trigger source for the frame as follows:

- If bit 11 is written with '1', the pause quantum will be taken from the Transmit Pause Quantum (*ETHx_tx_pause_quantum*) register. This register resets to a value of 0xFFFF giving maximum pause quantum as the initial value.
- If bit 12 is written with '1', the pause quantum will be zero.

After transmission, a pause frame transmitted interrupt will be generated (bit 14 of *ETHx_int_status*) and the only statistics register that will be incremented will be the Pause Frames Transmitted (*ETHx_pause_frame_txed*) register.

Pause frames can also be transmitted by the MAC using normal frame transmission methods.

31.3.10 MAC PFC Priority-based Pause Frame Support

EMAC supports PFC priority-based pause transmission and reception. Before PFC pause frames can be received, bit 16 of the Network Control register must be set to '1'.

31.3.10.1 PFC Pause Frame Reception

The ability to receive and decode priority-based pause frames is enabled by setting bit 16 of the Network Control (*ETHx_network_control*) register. When this bit is set, EMAC will match either classic IEEE Std 802.3 pause frames or PFC priority-based pause frames. After a priority-based pause frame is received and matched, the EMAC will only match priority-based pause frames (this is IEEE Std 802.1Qbb requirement, known as PFC negotiation). When a priority-based pause is negotiated, any received IEEE Std 802.3x format pause frames will not be acted upon.

If a valid priority-based pause frame is received, then EMAC will decode the frame and determine which, if any, of the eight priorities require to be paused. Up to eight pause time registers are then updated with the eight pause times extracted from the frame regardless of whether a previous pause operation is active. An interrupt (*ETHx_int_status.pause_frame_with_non_zero_pause_quantum_received*) is triggered when a non-zero PFC pause frame is received, but only if the interrupt is enabled (bit 12 of the interrupt mask register). Pause frames received with non-zero quantum are indicated through the interrupt bit 12 of the interrupt status register. PFC pause frames received with zero quantum cannot trigger an interrupt; that is, bit 13 is never set for PFC pause frames. The loading of a new pause time occurs because the EMAC is operating in full duplex mode. A valid pause frame is defined as having a destination address that matches either the address stored in Specific Address 1 register or if it matches the reserved address of 0180C2000001h. It must also have the MAC control frame Type ID of 8808h and the pause opcode 0101h.

Pause frames that have FCS or other errors will be treated as invalid and will be discarded. Valid pause frames received will increment the Pause Frames Received statistic register.

The Receive Pause Quantum (*ETHx_pause_time*) register decrement every 512 bit times immediately, following the PFC frame reception. For test purposes, the retry test bit can be set (bit 12 in the *ETHx_network_config* register), which causes the Receive Pause Quantum (*ETHx_pause_time*) register to decrement every RX_CLK cycle when transmission has stopped.

31.3.10.2 PFC Pause Frame Transmission

Automatic transmission of pause frames is supported through the transmit priority-based pause frame bit of the Network Control (*ETHx_network_control*) register. If

ETHx_network_control.transmit_pfc_priority_based_pause_frame is set to '1', a PFC pause frame will be transmitted, provided the MAC transmitter is enabled (bit 3) in *ETHx_network_control*. When bit 17 of *ETHx_network_control* is set to '1', the fields of the priority base pause frame will be built using the values stored in the Transmit PFC Pause (*ETHx_tx_pfc_pause*) register.

Pause frame transmission will happen immediately if transmit is inactive or if transmit is active between the current frame and the next frame due to be transmitted.

Transmitted pause frames comprise of the following:

- A destination address of 0180C2000001h
- A source address taken from Specific Address 1 register
- A Type ID of 8808h (MAC control frame)
- A pause opcode of 0101h
- A priority enable vector taken from the Transmit PFC Pause (*ETHx_tx_pfc_pause*) register
- Eight pause quanta in four registers (*ETHx_tx_pause_quantum*, *ETHx_tx_pause_quantum1*, *ETHx_tx_pause_quantum2*, *ETHx_tx_pause_quantum3*)
- Fill of 00h to take the frame to minimum frame length
- Valid FCS

The pause quantum registers used in the generated frame will depend on the trigger source for the frame as follows:

- If bit 17 of the Network Control (*ETHx_network_control*) register is set to '1', the priority enable vector of the priority-based pause frame will be equal to the value stored in the Transmit PFC Pause (*ETHx_tx_pfc_pause*) register bits [7:0]. For each entry equal to '0' in *ETHx_tx_pfc_pause* [15:8], the pause quantum field of the pause frame associated with that entry will be taken from the Transmit Pause Quantum (*ETHx_tx_pause_quantum*) register. For each entry equal to '1' in *ETHx_tx_pfc_pause* [15:8], the pause quantum associated with that entry will be '0'. The *ETHx_tx_pause_quantum* register resets to a value of FFFFh giving maximum pause quantum as the initial value.
- The pause quantum registers are classified as static and these registers should be updated only when no PFC frame is transmitted.
- To use the eight priority pause quanta stored in the four *ETHx_tx_pause_quantum* registers, set bit 24 to '1' in the *ETHx_network_control* register.

After transmission, a pause frame transmit interrupt will be issued (bit 14 of the *ETHx_int_status* register) and the only statistics register that will be incremented will be the Pause Frames Transmitted (*ETHx_pause_frame_txed*) register.

PFC Pause frames can also be transmitted by the MAC using normal frame transmission methods.

31.3.11 Energy Efficient Ethernet Support

IEEE Std 802.3az adds support for energy efficiency to Ethernet (EEE). These are the key features of IEEE Std 802.3az:

- Allows a system's transmit path to enter a low-power mode if there is nothing to transmit.
- Allows a PHY to detect whether its link partner's transmit path is in low-power mode, therefore allowing the system's receive path to enter low-power mode.
- Link remains up during low-power mode and no frames are dropped.
- Asymmetric, one direction can be in low-power mode while the other is transmitting normally.
- Low-power idle (LPI) signaling is used to control entry and exit to and from low-power modes.
- LPI signaling can only take place if both sides have indicated support for it through auto-negotiation.

IEEE Std 802.3az operation:

- Low-power control is done at MII/GMII (reconciliation sublayer).
- As an architectural convenience in writing the 802.3az, it is assumed that transmission is deferred by asserting carrier sense; in practice it will not be done this way. This system will know when it has nothing to transmit and only enter low-power mode when it is not transmitting.
- LPI should not be requested unless the link has been up for at least one second.
- LPI is signaled on the GMII transmit path by asserting 0x01 on TXD with TX_EN low and TX_ER high.
- A PHY on seeing LPI requested on MII/GMII will send the sleep signal before going quiet. After going quiet, it will periodically emit refresh signals.
- The sleep, quiet, and refresh periods are defined in Table 78-2 of the IEEE Std 802.3az. For 1000BASE-X, the sleep period is 20 microseconds, the quiet period is 2.5 milliseconds and the refresh period is 20 microseconds. For 100BASE-TX, the sleep period (T_s) is 100 microseconds, the quiet period is (T_{qt})/(T_{qr}) are 20/24 milliseconds, and the refresh period (T_r) is 100 microseconds. 10BASE-T is not supported.
- 1000BASE-X/100BASE-TX is required to go quiet after sleep is signaled.
- LPI mode ends by transmitting normal idle for the wake time. This has a default time, which can be adjusted in software using the Link Layer Discovery Protocol (LLDP) described in Clause 79 of IEEE Std 802.3az.
- LPI is indicated at the receive side when sleep and refresh signaling is detected.

31.3.11.1 LPI Operation in EMAC

Auto-negotiation:

- IndicateEEE capability using auto-negotiation.

For the transmit path:

- If the link has been up for 1 second and nothing is being transmitted, write to the LPI bit (bit 19) in the Network Control register.
- Wake up by clearing the LPI bit in the Network Control register.

For the receive path:

- Wait for an interrupt to indicate that LPI is received.
- Wait for an interrupt to indicate that regular idle is received and then re-enable the receive path.

31.3.12 MDIO Interface

MDIO is a single bi-directional tristate signal between the EMAC and PHY.

PHY maintenance register (*ETHx_phy_management*) is implemented as a shift register. Writing to the register starts a shift operation, which is signaled as complete when bit two is set in the network status register (about 2000 pclk cycles later when bits [18:16] are set to 010 in the *ETHx_network_config* register). An interrupt is generated as this bit is set.

During this time, the MSb of the register is output on the *mdio_out* pin and the LSb updated from the *mdio_in* pin with each MDC cycle. This causes the transmission of a PHY management frame on MDIO.

Reading during the shift operation will return the current contents of the shift register. At the end of the management

operation the bits will have shifted back to their original locations. For a read operation, the data bits will be updated with data read from the PHY. It is important to write the correct values to the register to ensure a valid PHY management frame is produced.

MDC should not toggle faster than 2.5 MHz (minimum period of 400 ns), as defined by the IEEE 802.3 standard. MDC is generated by dividing *clk_sys*. Three bits in the network configuration register determine by how much pclk should be divided to produce MDC.

31.3.13 Interrupts

Several interrupt conditions can be enabled in EMAC. The interrupt outputs from the Ethernet MAC match the number of supported priority queues. Only EMAC DMA related events are reported using the individual interrupt outputs, as the Ethernet MAC can relate these events to specific queues. All other events generated within the Ethernet MAC are reported in the interrupt associated with the lowest priority queue (Queue 0). For the lowest priority queue, the interrupt status (*ETHx_int_status*) register is located at offset address 0x024. For all other priority queues, this register is located at sequential offset addresses starting at 0x400.

At reset all interrupts are disabled. To enable an interrupt, write to interrupt enable register with the pertinent interrupt bit set to 1. To disable an interrupt, write to interrupt disable register with the pertinent interrupt bit set to 1. To check whether an interrupt is enabled or disabled, read interrupt mask register: if the bit is set to 1, the interrupt is disabled.

31.3.14 Media Independent Interfaces

31.3.14.1 MII/GMII

EMAC connects to PHY using the MII/GMII interface (*ETHx_TXD*, *ETHx_TX_ER*, *ETHx_TX_CLK*, *ETHx_TX_CTL*, *ETHx_RXD*, *ETHx_RX_ER*, *ETHx_RX_CLK*, and *ETHx_RX_CTL*).

Configuration and status information is exchanged between the EMAC and PHY using the management interface (*ETHx_MDC* and *ETHx_MDIO*).

Table 31-6. Signals Used in MII/GMII

Signal Name	Description	Direction
<i>ETHx_TXD</i>	4 Transmit data lines for MII 8 Transmit data lines for GMII	MAC to PHY
<i>ETHx_TX_ER</i>	Transmit error	MAC to PHY
<i>ETHx_TX_CTL</i>	Transmit enable	MAC to PHY
<i>ETHx_TX_CLK</i>	Transmit clock	PHY to MAC for MII [10/100 Mbps] MAC to PHY for GMII [1000 Mbps]
<i>ETHx_RXD</i>	4 Receive data lines for MII 8 Receive data lines for GMII	PHY to MAC

Table 31-6. Signals Used in MII/GMII

Signal Name	Description	Direction
ETHx_RX_CTL	Receive data valid	PHY to MAC
ETHx_RX_ER	Receive error	PHY to MAC
ETHx_RX_CLK	Receive clock	PHY to MAC
ETHx_REF_CLK	Reference clock for GMII	Input to MAC

31.3.14.2 RGMII

In RGMII mode, the core receives a data nibble on ETHx_RXD[3:0] on both the edges of the ETHx_RX_CLK along with a single-bit control signal. The lower nibble of the data byte is received on the rising edge of ETHx_RX_CLK, and the upper nibble is received on the falling edge. The RGMII module receives a control bit on ETHx_RX_CTL on both edges of ETHx_RX_CLK. This control bit contains the data valid signal ETHx_RX_DV and encoded receiver error information of ETHx_RX_ER. The RGMII also uses the data and control signals to extract the in-band PHY status for carrier sense, collision, speed, link, and duplex.

Similarly, RGMII transmits the lower nibble of the data byte on ETHx_TXD[3:0] on the rising edge of ETHx_TX_CLK and the upper nibble is transmitted on the following falling edge. Control signals ETHx_TX_EN and ETHx_TX_ER are multiplexed on ETHx_TX_CTL.

The RGMII standard specifies a source synchronous clock. It relies on the clock having a longer path delay than the data so that the data is resampled using the same edge of the clock on which it was generated. Therefore, delay on clock signals are essential, which can be achieved either through trace lines or through delaying clocks from the source. EMAC supports RGMII specifications from OPEN Alliance. During the transmission process MAC operates in Delay-on-Destination (DoD) mode; hence required delay must be accomplished by the PHY or through trace lines. For receive operation, MAC operates in Delay-on-Source (DoS) mode, in which delay must be achieved either by the PHY or through trace lines.

Configuration and status information is exchanged between the EMAC and PHY using the management interface (ETHx_MDC and ETHx_MDIO).

Table 31-7. Signals Used in RGMII

Signal Name	Description	Direction
ETHx_TXD	4 Transmit data lines	MAC to PHY
ETHx_TX_CTL	Transmit enable	MAC to PHY
ETHx_TX_CLK	Transmit clock	MAC to PHY
ETHx_RXD	4 Receive data lines	PHY to MAC
ETHx_RX_CTL	Receive data valid and error signal	PHY to MAC
ETHx_RX_CLK	Receive clock	PHY to MAC
ETHx_REF_CLK	High-speed reference clock	Input to MAC

31.3.14.3 RMII

As the name suggests, Reduced Media Independent Interface (RMII) communicates data to PHY through lower number of pins compared to MII. Transmit and received data lines are reduced from four to two, and clock is doubled. The clock source for RMII can be selected either from internal PLL source² or external high-precision clock through clock line ETHx_REF_CLK; required clock for transmit and receive operations will be generated internally from this reference clock.

2. Please check device specific datasheet to know availability of internal reference clock for RMII.

Configuration and status information is exchanged between the EMAC and PHY using the management interface (ETHx_MDC and ETHx_MDIO).

Table 31-8. Signals Used in RMII

Signal Name	Description	Direction
ETHx_TXD	2 Transmit data lines	MAC to PHY
ETHx_TX_CTL	Transmit enable	MAC to PHY
ETHx_RXD	4 Receive data lines	PHY to MAC
ETHx_RX_ER	Receive Error	PHY to MAC
ETHx_RX_CTL	Receive data valid	PHY to MAC
ETHx_REF_CLK	External 50 MHz reference clock	Input to MAC
ETHx_TX_CLK	Reference clock to PHY when internal reference clock is selected	MAC to PHY

As described, EMAC supports four different PHY interfaces – MII, RMII, GMII, and RGMII. Table 31-9 shows the required settings to select any of the interfaces.

Table 31-9. PHY Interface Selection

ETHx_CTL.ETH_MODE	ETHx_network_config[0] (speed)	ETHx_network_config[10] (gigabit_mode_enable)	PHY Mode
2'd0	0	0	MII - 10 Mbps
2'd0	1	0	MII - 100 Mbps
2'd1	0	1	GMII - 1000 Mbps
2'd2	0	0	RGMII - 10 Mbps (4 bits/Cycle)
2'd2	1	0	RGMII - 100 Mbps (4 bits/Cycle)
2'd2	0	1	RGMII - 1000 Mbps (8 bits/Cycle)
2'd3	0	0	RMII - 10 Mbps
2'd3	1	0	RMII - 100 Mbps

Notes:

- ETH_MODE must be configured before configuring network_config register.
- In the RGMII interface:
 - MAC is in Delay-on-Destination (DoD) mode for transmission; the delay needs to be accomplished by PHY.
 - MAC is in Delay-on-Source (DoS) mode for receive operation; the delay needs to be accomplished by PHY.

31.3.14.4 Clock Sources for PHY Interface

Clock requirements and configurations are different for each interface. Following are the required clocks and source of clocks.

■ MII

- Both TX and RX clocks are supplied from external PHY.

■ RMII

- Both TX and RX clocks can be supplied from either internal reference clock or from external clock source.³
- ETHx_CTL.REFCLK_SRC_SEL must be used to select reference clock source from on-chip system resource or from HSIO.
- a. Clock out will be enabled internally when the internal clock source is selected; TX reference clock to PHY can be provided through ETHx_TX_CLK.
- ETHx_CTL.REFCLK_DIV must be used to divide reference clock and generate required frequency of 50 MHz.

■ GMII

- RX clock is supplied from PHY.
- TX clock source can be selected either from internal clock source or from HSIO.
- ETHx_CTL.REFCLK_SRC_SEL must be used to select clock source for TX functionality.
- ETHx_CTL.REFCLK_DIV is used to divide reference clock and to generate required clock of 125 MHz.
- Clock out will be enabled internally when the internal clock source is selected; TX reference clock to PHY can be provided through ETHx_TX_CLK.

■ RGMII


- RX clock is supplied from PHY.
- TX clock source can be selected either from internal clock source or from HSIO.
- ETHx_CTL.REFCLK_SRC_SEL must be used to select clock source for TX functionality.
- ETHx_CTL.REFCLK_DIV is used to divide reference clock and to generate required clock of 125 MHz.
- Clock out will be enabled internally when the internal clock source is selected; TX reference clock to PHY can be provided through ETHx_TX_CLK.

Note: Use a more precise external clock source than the internal PLL for RGMII and GMII transmit operations.

3. Please check device specific datasheet to know availability of internal reference clock for RMII.

31.3.14.5 PHY Interface

Figure 31-3. PHY Interface



31.3.15 Clocks to EMAC

Other than clocks described in [31.3.14.4 Clock Sources for PHY Interface](#), EMAC requires clocks to perform internal operation such as buffer data transfers or TSU operations. To perform those operations the following clocks are used. For more information about configuring mentioned clocks, see the [Clocking System chapter on page 197](#).

Table 31-10. Clocks to EMAC

Clock	Function
clk_sys	To generate MDC clock and for AHB operation, clk_sys is derived from CLK_PERI
clk_tsu	TSU clock for TSU timer
clk_mem	Fast clock for AXI operation
ref_int_clk	Internal reference clock supplied from PLL

Note: Refer to the device datasheet for assigned clocks to different EMAC instances.

31.3.16 Power Modes

[Table 31-11](#) shows EMAC availability in different device power modes. See the [Device Power Modes chapter on page 186](#) for more details.

Table 31-11. EMAC Status during Different Device Power Modes

Device Power Mode	IP Status
Active	EMAC is fully operational in Active power mode with power on and clocks running
LPActive	EMAC is fully operational in LPActive power mode with power on and clocks running; clocks can be limited to save some power during LPActive mode
Sleep	EMAC is fully operational in Sleep power mode
LPSleep	EMAC is fully operational in LPSleep power mode with power on and clocks running; clocks can be limited to save some power during LPSleep mode
DeepSleep	No clock is provided during DeepSleep power mode, hence the logic is not functional. All retention registers will hold the values in DeepSleep mode
Hibernate	Entire EMAC including retention register are not functional during Hibernate power mode

31.4 Register List

Table 31-12. EMAC Registers

Register	Name	Description
ETHx_CTL	Control Register	Control register to select type PHY mode and reference clock, and to enable the IP
ETHx_STATUS	Status Register	Status register shows status about PFC frames
ETHx_network_control	Network Control Register	The network control register contains general MAC control functions for both receiver and transmitter
ETHx_network_config	Network Configuration Register	The network configuration register contains functions to set the mode of operation for the Gigabit Ethernet MAC
ETHx_network_status	Network Status Register	The network status register shows status regarding PHY management interface
ETHx_dma_config	DMA Configuration Register	Register to configure DMA transfers for transmit and receive operation
ETHx_transmit_status	Transmit Status Register	Register provides the status of a transmit. After it is read, individual bits can be cleared by writing 1 to them. It is not possible to set a bit to 1 by writing to the register
ETHx_receive_q_ptr	Receive Queue 0 Pointer	This register holds the start address of the receive buffer queue (receive buffers descriptor list).
ETHx_transmit_q_ptr	Transmit Queue 0 Pointer	This register holds the start address of the transmit buffer queue (transmit buffers descriptor list).
ETHx_receive_status	Receive Status Register	Register provides the status of receive operation. After it is read, individual bits may be cleared by writing 1 to them. It is not possible to set a bit to 1 by writing to the register
ETHx_int_status	Interrupt Status Register	If not configured for priority queuing, the EMAC generates a single interrupt. This register indicates the source of this interrupt
ETHx_int_enable	Interrupt Enable	At reset all interrupts are disabled. Writing a one to the relevant bit location enables the required interrupt. This register is write only and when read will return zero.
ETHx_int_disable	Interrupt Disable	Register to disable interrupts
ETHx_int_mask	Interrupt Mask Register	The interrupt mask register is a read only register indicating which interrupts are masked. All bits are set at reset and can be reset individually by writing to the interrupt enable register or set individually by writing to the interrupt disable register
ETHx_phy_management	PHY Management Register	This register is implemented as a shift register. Writing to the register starts a shift operation which is signaled as complete when bit-2 is set in the network status register.
ETHx_pause_time	Pause Quantum Register	Received pause quantum register
ETHx_tx_pause_quantum	Quantum Register	Register to hold quantum to be transmitted
ETHx_pbuf_txcutthru	Packet Buffer Cut Through Configuration	Register is used to configure partial store and forward mode for transmit operation
ETHx_pbuf_rxcutthru	Packet Buffer Cut Through Configuration	Register to configure partial store and forward mode for receive operation
ETHx_jumbo_max_length	Jumbo Length Configuration	Jumbo frame size configuration
ETHx_axi_max_pipeline	AXI Pipeline Configuration	Used to set the maximum amount of outstanding transactions on the AXI bus between AR/R channels and AW/W channels.

Table 31-12. EMAC Registers

Register	Name	Description
ETHx_int_moderation	Interrupt Moderation Register	Used to moderate the number of transmit and receive complete interrupts issued. With interrupt moderation enabled receive and transmit interrupts are not generated immediately a frame is transmitted or received. Instead when a receive or transmit event occurs a timer is started, and the interrupt is asserted after it times out.
ETHx_sys_wake_time		Used to pause transmission after low-power idle is deasserted
ETHx_hash_bottom	Hash Configuration Register	The unicast hash enable and the multicast hash enable bits in the network configuration register enable the reception of hash matched frames. Hash Register Bottom (31 to 0 bits)
ETHx_hash_top	Hash Top Configuration Register	Top 32 bits of Hash configuration
ETHx_spec_add1_bottom to ETHx_spec_add4_bottom	Specific Address (Bottom) Configuration	The addresses stored in the specific address registers are deactivated at reset or when their corresponding specific address register bottom is written. They are activated when specific address register top is written
ETHx_spec_add1_top to ETHx_spec_add4_top	Specific Address (Top) Configuration	Specific top address
ETHx_spec_type1 to ETHx_spec_type4	Type ID Match Register	Type ID match register
ETHx_stretch_ratio	IPG Stretch Configuration	Inter packet gap stretch register
ETHx_stacked_vlan	Stacked VLAN	Stacked VLAN register
ETHx_tx_pfc_pause	Transmit Pause Register	Transmit PFC pause register
ETHx_mask_add1_bottom	Address Mask Register	Specific address mask 1 bottom (31 to 0 bits)
ETHx_mask_add1_top	Mask Register	Specific address mask 1 top (47 to 32 bits)
ETHx_dma_addr_or_mask	Receive DMA Data Buffer Address Mask	Receive DMA data buffer address mask
ETHx_rx_ptp_unicast	PTP RX Unicast IP Destination Address	Unicast IP destination address to be detected during receiving PTP frame process.
ETHx_tx_ptp_unicast	PTP TX Unicast IP Destination Address	Unicast IP destination address to be detected during transmitting PTP frame process
ETHx_tsu_nsec_cmp	TSU Timer Comparison Value Nanoseconds	TSU timer comparison value nanoseconds
ETHx_tsu_sec_cmp	TSU Timer Comparison Value Seconds	TSU timer comparison value seconds (31 to 0 bits)
ETHx_tsu_msb_sec_cmp	TSU Timer Comparison Value Seconds	TSU timer comparison value seconds (47 to 37 bits)
ETHx_tsu_ptp_tx_msb_sec	PTP Event Frame Transmitted (Seconds) Register	PTP event frame TX seconds. The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP transmit primary event crosses the MII interface.
ETHx_tsu_ptp_rx_msb_sec	PTP Event Frame Received (Seconds) Register	PTP event frame RX seconds. The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP receive primary event crosses the MII interface
ETHx_tsu_peer_tx_msb_sec	PTP Peer Event Frame Transmitted Seconds Register (47 to 32 bits)	PTP Peer event frame TX seconds. The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP transmit peer event crosses the MII interface.

Table 31-12. EMAC Registers

Register	Name	Description
ETHx_tsu_peer_rx_msb_sec	PTP Peer Event Frame Received Seconds Register (47 to 32 bits)	PTP Peer event frame RX seconds. The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP transmit peer event crosses the MII interface.
ETHx_dpram_fill_dbg	DMA Packet Buffer Fill Level	Register provides fill level of TX and RX packet buffer
ETHx_revision_req	Module Identification Number	Contains module identification number and revision number
ETHx_octets_tx_bottom	Octets Transmitted	Register shows total valid octets transmitted (lower 32 bits)
ETHx_octets_tx_top	Octets Transmitted	Register shows total valid octets transmitted (47 to 32 bits)
ETHx_frames_txed_ok	Transmitted Frames	Register shows number of frames transmitted without error
ETHx_broadcast_txed	Broadcast Frames Transmitted	Number of broadcast frames transmitted without error
ETHx_multicast_txed	Multicast Frames Transmitted	Number of multicast frames transmitted without error
ETHx_pause_frames_txed	Pause Frames Transmitted	Number of pause frames transmitted without error
ETHx_frames_txed_64	64 Byte Frames Transmitted	Number of 64 byte frames transmitted without error
ETHx_frames_txed_65	65 to 127 Byte Frames Transmitted	Number of 65 to 127 byte frames transmitted without error
ETHx_frames_txed_128	128 to 255 Byte Frames Transmitted	Number of 128 to 255 byte frames transmitted without error
ETHx_frames_txed_256	256 to 511 Byte Frames Transmitted	Number of 256 to 511 byte frames transmitted without error
ETHx_frames_txed_512	512 to 1023 Byte Frames Transmitted	Number of 512 to 1023 byte frames transmitted without error
ETHx_frames_txed_1024	1024 to 1518 Byte Frames Transmitted	Number of 1024 to 1518 byte frames transmitted without error
ETHx_frames_txed_1519	Greater than 1518 Byte Frames Transmitted	Greater than 1518 byte frames transmitted without error
ETHx_tx_underruns	Transmit Underrun	10-bit register counting the number of frames not transmitted due to a transmit under run
ETHx_octets_rxed_bottom	Octets Received (31 to 0 bits)	Number of octets received
ETHx_octets_rxed_top	Octets Received (47 to 32 bits)	Number of octets received
ETHx_frames_rxed_ok	Frames Received	Number of frames received without error
ETHx_broadcast_rxed	Broadcast Frames Received	Number of broadcast frames received without error
ETHx_multicast_rxed	Multicast Frames Received	Number of multicast frames received without errors
ETHx_pause_frames_rxed	Pause Frames Received	Number of pause frames received without errors
ETHx_frames_rxed_64	64 Byte Frames Received	Number of frames with 64 bytes received without error
ETHx_frames_rxed_65	65 to 127 Byte Frames Received	Number of frames with 65 to 127 bytes received without error
ETHx_frames_rxed_128	128 to 255 Byte Frames Received	Number of frames with 128 to 255 bytes received without error
ETHx_frames_rxed_256	256 to 511 Byte Frames Received	Number of frames with 256 to 511 bytes received without error
ETHx_frames_rxed_512	512 to 1023 Byte Frames Received	Number of frames with 512 to 1023 bytes received without error
ETHx_frames_rxed_1024	1024 to 1518 Byte Frames Received	Number of frames with 1024 to 1518 bytes received without error

Table 31-12. EMAC Registers

Register	Name	Description
ETHx_frames_rxed_1519	1519 to Maximum Byte Frames Received	Number of frames with 1519 or higher bytes received without error
ETHx_undersize_frames	Undersized Frames Received	10-bit register counting the number of frames received less than 64 bytes in length
ETHx_excessive_rx_length	Oversize Frames Received	10-bit register counting the number of frames received exceeding 1518 bytes
ETHx_rx_jabbers	Jabbers Received	Jabbers frames received - a 10-bit register counting the number of frames received exceeding 1518 bytes
ETHx_fcs_errors	Frame Check Sequence Errors	Frame check sequence errors - a 10-bit register counting frames that are an integral number of bytes, have bad CRC and are between 64 and 1518 bytes in length
ETHx_rx_length_errors	Length Field Frames Errors	Length field frame errors - this 10-bit register counts the number of frames received that have a measured length shorter than that extracted from the length field (bytes 13 and 14).
ETHx_rx_symbol_errors	Receive Symbol Errors	Receive symbol errors - a 10-bit register counting the number of frames that had rx_er asserted during reception
ETHx_alignment_errors	Alignment Errors	10-bit register counts alignment errors in received frames. Symbol error also increments this register
ETHx_rx_resource_errors	Receive Resource Errors	18-bit register counting the number of frames that were successfully received by the MAC but could not be copied to memory because no receive buffer was available.
ETHx_rx_overruns	Receive Overrun Error Counter	10-bit register counts number of frames which were received correctly but were dropped due to receive overrun condition
ETHx_rx_ip_ck_errs	IP Header Checksum Errors	8-bit register counting the number of frames discarded due to an incorrect header checksum
ETHx_rx_tcp_ck_errors	TCP Checksum Errors	8-bit register counting the number of frames discarded due to an incorrect TCP checksum
ETHx_rx_udp_ck_errors	UDP Checksum Errors	8-bit register counting the number of frames discarded due to an incorrect UDP checksum
ETHx_auto_flushed_pkts	Receive DMA Flushed Packets	16-bit register counting the number of frames that are flushed from the receive SRAM based packet buffer under some specific conditions
ETHx_tsu_timer_incr_sub_nsec	Timer Increment Register	1588 timer increment register sub nsec
ETHx_tsu_timer_msb_sec	Timer Seconds Register	1588 timer seconds register (47 to 32 bits)
ETHx_tsu_strobe_msb_sec	Timer Sync Strobe Seconds Register	1588 timer sync strobe seconds register (47 to 32 bits)
ETHx_tsu_strobe_sec	Timer Sync Strobe Seconds Register	1588 timer sync strobe seconds register (31 to 0 bits)
ETHx_tsu_strobe_nsec	Timer Sync Strobe Seconds Register	1588 timer sync strobe nanoseconds register
ETHx_tsu_timer_sec	Timer Seconds Register	1588 timer seconds register (31 to 0 bits)
ETHx_tsu_timer_nsec	Timer Nanoseconds Register	1588 timer nanoseconds register
ETHx_tsu_timer_adjust	Timer Adjust Register	This register is used to adjust the value of the timer in the TSU. It allows an integral number of nanoseconds to be added or subtracted from the timer in a one-off operation. This register returns all zeros when read.
ETHx_tsu_timer_incr	1588 Timer Increment Register	Timer increment register is used to set a count of nanoseconds by which the 1588 timer nanoseconds register will be incremented each clock cycle

Table 31-12. EMAC Registers

Register	Name	Description
ETHx_tsu_ptp_tx_sec	PTP Event Frame Transmitted Seconds	The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP transmit primary event crosses the MII interface
ETHx_tsu_ptp_tx_nsec	PTP Event Frame Transmitted Nanoseconds	The register is updated with the value that the 1588 timer nanoseconds register held when the SFD of a PTP transmit primary event crosses the MII interface
ETHx_tsu_ptp_rx_sec	PTP Event Frame Received Seconds	The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP receive primary event crosses the MII interface
ETHx_tsu_ptp_rx_nsec	PTP Event Frame Received Nanoseconds	The register is updated with the value that the 1588 timer nanoseconds register held when the SFD of a PTP receive primary event crosses the MII interface
ETHx_tsu_peer_tx_Sec	PTP Peer Event Frame Transmitted Seconds Register	The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP transmit peer event crosses the MII interface
ETHx_tsu_peer_tx_nsec	PTP Peer Event Frame Transmitted Nanoseconds Register	The register is updated with the value that the 1588 timer nanoseconds register held when the SFD of a PTP transmit peer event crosses the MII interface
ETHx_tsu_peer_rx_sec	PTP Peer Event Frame Received Seconds Register	The register is updated with the value that the 1588 timer seconds register held when the SFD of a PTP receive peer event crosses the MII interface
ETHx_tsu_peer_rx_nsec	PTP Peer Event Frame Received Nanoseconds Register	The register is updated with the value that the 1588 timer nanoseconds register held when the SFD of a PTP receive peer event crosses the MII interface
ETHx_tx_pause_quantum1	Transmit Pause Quantum Register 1	Transmit pause quantum - written with the pause quantum value for pause frame transmission of priority 2 and 3
ETHx_tx_pause_quantum2	Transmit Pause Quantum Register 2	Transmit pause quantum - written with the pause quantum value for pause frame transmission of priority 4 and 6
ETHx_tx_pause_quantum3	Transmit Pause Quantum Register 3	Transmit pause quantum - written with the pause quantum value for pause frame transmission of priority 7 and 8
ETHx_rx_lpi	Received LPI Transitions	Count of RX LPI transitions. A count of the number of times there is a transition from receiving normal idle to receiving low-power idle
ETHx_rx_lpi_time	Received LPI Time	Time in LPI. This register increments once every 16 pclk cycles when the LPI indication bit 20 is set in the receive configuration register
ETHx_tx_lpi	Transmit LPI Transitions	Count of LPI transmissions. A count of the number of times the enable LPI transmission bit 20 goes from low to high in the transmit control register.
ETHx_tx_lpi_time	Transmit LPI Time	Time in LPI. This register increments once every 16 pclk cycles when the LPI indication bit 20 is set in the transmit configuration register
ETHx_int_q1_status	Priority Queue Interrupt Status Register	Priority queue 1 interrupt status register
ETHx_int_q2_status	Priority Queue Interrupt Status Register	Priority queue 2 interrupt status register
ETHx_transmit_q1_ptr	Transmit Q1 Pointer	This register holds the start address of the transmit buffers descriptor list for queue 1
ETHx_transmit_q2_ptr	Transmit Q2 Pointer	This register holds the start address of the transmit buffers descriptor list for queue 2

Table 31-12. EMAC Registers

Register	Name	Description
ETHx_receive_q1_ptr	Receive Q1 Pointer	This register holds the start address of the receive buffers descriptor list for queue 1
ETHx_receive_q2_ptr	Receive Q2 Pointer	This register holds the start address of the receive buffers descriptor list for queue 2
ETHx_dma_rxbuf_size_q1	Receive Buffer Queue 1 Size	DMA receive buffer size in system memory. The value defined by these bits determines the size of buffer to use in main system memory when writing received data
ETHx_dma_rxbuf_size_q2	Receive Buffer Queue 2 Size	DMA receive buffer size in system memory. The value defined by these bits determines the size of buffer to use in main system memory when writing received data
ETHx_cbs_control	Credit Based Shaping Control Register	Register is used to enable credit based shaping on high-priority queues
ETHx_cbs_idleslope_q_a	Queue A Idle Slope Configuration Register	Register is used to configure idle slope for highest priority queue, which is queue 2
ETHx_cbs_idleslope_q_b	Queue B Idle Slope Configuration Register	Register is used to configure idle slope for second highest priority queue, which is queue 1
ETHx_tx_bd_control	TX BD Control	Transmit buffer descriptor mode configuration register
ETHx_rx_bd_control	RX BD Control	Receive buffer descriptor mode configuration register
ETHx_screening_type_1_register_0 to ETHx_screening_type_1_register_15	Screening Type 1 Register	Screening type 1 registers are used to allocate up to 16 priority queues to received frames based on certain block or UDP fields of incoming frames
ETHx_screening_type_2_register_0 to ETHx_screening_type_2_register_15	Screening Type 2 Register	Screening type 2 registers operate independently of screening type 1 registers and offer additional match capabilities, extending the capabilities into vendor specific protocols
ETHx_tx_sched_ctrl	Transmit Scheduling Control Register	This register controls the transmit scheduling algorithm the user can select for each active transmit queue. By default, all queues are initialized to fixed priority, with the top indexed queue having overall priority
ETHx_bw_rate_limit_q0toq3	Bandwidth Allocation Register	This register holds the DWRR weighting value or the ETS bandwidth percentage value used by the transmit scheduler for queues 0 to 2
ETHx_int_q1_enable	Interrupt Enable for Q1	At reset all interrupts are disabled. Writing a one to the relevant bit location enables the required interrupt. This register is write-only and when read will return zero
ETHx_int_q2_enable	Interrupt Enable for Queue 2	At reset all interrupts are disabled. Writing a one to the relevant bit location enables the required interrupt. This register is write-only and when read will return zero
ETHx_int_q1_disable	Interrupt Disable for Queue 1	Writing a one to the relevant bit location disables that interrupt. This register is write-only and when read will return zero
ETHx_int_q2_disable	Interrupt Disable for Queue 2	Writing a one to the relevant bit location disables that interrupt. This register is write-only and when read will return zero
ETHx_int_q1_mask	Interrupt Mask Register for Queue 1	The interrupt mask register is a read-only register indicating which interrupts are masked.

Table 31-12. EMAC Registers

Register	Name	Description
ETHx_int_q2_mask	Interrupt Mask Register for Queue 2	The interrupt mask register is a read-only register indicating which interrupts are masked.
ETHx_screening_type_2_ethertype_reg_0 to ETHx_screening_type_2_ethertype_reg_7	Screener Type 2 Ethertype Compare Register	Screening register to compare Ethertype
ETHx_type2_compare_0_word_0 to ETHx_type2_compare_31_word_0	Screener Type 2 Compare Register	Type 2 compare register word 0; see the <i>Traveo II Body Controller High Registers TRM</i> for full description
ETHx_type2_compare_0_word_1 to ETHx_type2_compare_31_word_1	Screener Type 2 Compare Register	Type 2 compare register word 1; see the <i>Traveo II Body Controller High Registers TRM</i> for full description

32. Serial Memory Interface



The Serial Memory InterFace (SMIF) block is a master that provides a low pin count connection to off-chip (single/dual/quad/octal) SPI devices, such as EEPROM, FRAM, MRAM, or NAND in SDR or DDR mode, and HyperBus devices such as HyperFlash (NOR flash) and HyperRAM (PSRAM and pseudo static RAM).

This chapter explains the features, implementation, and operational modes of SMIF.

32.1 Features


- SPI or HyperBus master functionality only
- HyperBus protocol
- SPI protocol
 - SPI mode 0 only, with configurable MISO sampling timing
 - Supports single/dual/quad/octal SPI
 - Supports dual-quad SPI mode
 - Supports single data rate (SDR) and dual data rate (DDR) transfers
- Memory device
 - Supports overall device capacity in the range of 64 KB to 4 GB in power of two multiples
 - Supports configurable external device capacities
 - Supports two external memory devices
- Memory mapped I/O (MMIO) operation mode
- XIP mode
 - eXecute-In-Place (XIP) operation mode for both read and write accesses
 - XIP mode supports on-the-fly encryption and decryption
 - XIP operation mode via AHB interface for CM0 and AXI interface for CM7 core
 - Supports up to four outstanding transactions
- Memory interface logic
 - Supports stalling of SPI and HyperBus transfers to address back pressure on FIFOs
 - Supports an asynchronous SPI/HyperBus transmit and receive interface clock
 - Supports read-write-data-strobe (RWDS)
 - Supports multiple interface receive clocks
 - Supports flexible external SPI memory devices data signal connections
 - Independent SPI interface transmitter clock from PLL/FLL
 - SPI interface logic supports flexible external memory devices data signal connections

32.2 System Diagram

The SMIF allows the Traveo II device to connect to off-chip NOR/NAND flash, SRAM, FRAM, nonvolatile SRAM, or DRAM (using single/dual/quad SPI or HyperBus interface) memories.

Figure 32-1 shows how to connect to external serial devices; up to four SPI/HyperBus devices can be connected.

Figure 32-1. QSPI Device Example System Diagram




32.3 Block Diagram

Figure 32-2 gives a high-level overview of the SMIF. The bottom part of the figure shows the SPI signal connections to the I/O subsystem (IOSS). The top part of the figure shows the interface to the CM7 core (one fast AXI interface and one slow AHB interface) and the AHB-Lite slave interface to the peripheral group. CM7 uses the fast AXI interface, CM4 uses the fast AHB-Lite interface and the slow interface is used by the CM0+, crypto, and P-DMA components.

The memory interface logic supports an asynchronous interface clock (clk_if) from which the interface transmit clock (clk_if_tx) and the interface receive clock (clk_if_rx) are derived.

Note that each XIP AHB-Lite interface has a dedicated cache. Cache coherency is not supported by the hardware. For example, an XIP interface 0 write to an address in the XIP interface 0 cache invalidates the associated cache subsector in the XIP interface 0 cache, but not in the XIP interface 1 cache.

Figure 32-2. High-level Block Diagram of mxsmif (SMIF)



Note: The inverted clock signal 'sphb_clk_inv' is generated by a special I/O cell and not inside the SMIF.

32.3.1 Clocks

The SMIF uses four different clock domains:

- **CLK_MEM:** It clocks the CPUSS fast infrastructure and is a divided version of CLK_HF. CLK_MEM is used for the XIP AXI interface in the SMIF configuration CM7 core. It is also used for the shared XIP-related blocks (port arbiter, cryptography, XIP, and memory interface TX and RX FIFOs).
- **CLK_SLOW:** This is the source clock for the Cortex-M0+. It is a divided version of CLK_MEM. CLK_SLOW is used for the XIP AHB-Lite interface.
- **CLK_SYS:** This clock is used as the interface clock of the MMIO interface. However, the MMIO registers are clocked by CLK_MEM (same clock domain as the blocks where MMIO registers are used or depend on). CLK_SYS is a gated clock, which can be synchronous or asynchronous to CLK_MEM.
- **CLK_IF:** Memory interface clock signals CLK_IF and sphb_clk_in are asynchronous to the signals in the other groups. These clocks are used for the SPI/HyperBus interface logic. To generate the sphb_clk_out memory clock, a divided by 2 CLK_IF is used.

32.4 Functional Description

This section describes the basic functions of SMIF.

32.4.1 Operating Modes

The SMIF has the following interfaces:

- An AHB-Lite interface to access the MMIO registers.
- An AHB-Lite and an AXI interface to support XIP.

All interfaces provide access to external memory devices. At any time, either the MMIO AHB-Lite interface or the two XIP interfaces have access to the SPI interface logic and external memory devices. The operation mode is specified by XIP_MODE in the CTL register. The operation mode should not be modified when the SMIF is busy, indicated by the BUSY bit in the STATUS register.

In the MMIO AHB-Lite interface, access is supported through software writes to transmit (TX) FIFOs and software reads from a receive (RX) FIFO. The FIFOs are mapped on MMIO registers. This interface provides the flexibility to implement any SPI device transfer. For example, the SPI device transfers to set up, program, or erase the external memory devices.

In a XIP AHB-Lite interface, access is supported through XIP: AHB-Lite read and write transfers are automatically translated (by the hardware) in SPI device read and write transfers. This interface provides efficient implementation of SPI device read and write transfers, but does not support other types of SPI device transfers. To improve XIP performance, the XIP AHB-Lite interface has a 4-KB cache.

As mentioned, MMIO mode and XIP mode are mutually exclusive. The operation modes share TX and RX FIFOs and SPI interface logic. In MMIO mode, the TX and RX FIFOs are mapped on MMIO registers and under software control. In XIP mode, the TX and RX FIFOs are under hardware control. The SPI interface logic is controlled through the TX and RX FIFOs and is agnostic of the operation mode.

32.4.1.1 MMIO_MODE

The MMIO mode can be activated by writing '0' to the XIP_MODE bit in the CTL register. The software generates SPI or HyperBus (HB) transfers by accessing the TX FIFOs and RX FIFO. The TX FIFOs are write-accessible and read accesses are done from the RX FIFO. The TX command FIFO has formatted commands (TX, TX_COUNT, RX_COUNT, DUMMY_COUNT) that are described in the memory map.

The software should ensure that it generates correct memory transfers and access the FIFOs correctly. For example, if a memory transfer is generated to read four bytes from a memory device, software should read the four bytes from the RX data FIFO. Similarly, if a memory transfer is generated to write four bytes to a memory device, the

software should write the four bytes to the TX command FIFO or TX data FIFO.

Incorrect software behavior can lock up the memory interface. For example, a memory transfer to read 32 bytes from a memory device, without the software reading the RX data FIFO will lock up the memory transfer as the memory interface cannot provide more than eight bytes to the RX data FIFO (the RX data FIFO has eight entries). This would prevent any successive memory transfers from taking place. Note that a locked memory transfer due to TX or RX FIFO states is still compliant to the memory bus protocol (but undesirable): the SPI or HB protocol allows shutting down the interface clock SPIHB_CLK in the middle of a memory transfer.

32.4.1.2 XIP_MODE

If the XIP_MODE bit is set to '1' in the CTL register, the SMIF is in XIP mode. The hardware automatically generates (without software intervention) memory transfers by accessing the TX FIFOs and RX FIFO. The hardware only supports memory read and write transfers.

- Hardware generates memory read transfers for AHB-Lite or AXI read transfers (that is, only for AHB-Lite read transfers that miss in the cache or AXI read transfers).
- Hardware generates memory write transfers for AHB-Lite or AXI write transfers.

This is done in the XIP block, which:

- translates read or write transfer requests from the AHB-Lite or AXI interfaces to commands in the TX command FIFO
- sends/receives data to/from the TX/RX data FIFOs

As different memory devices support different types of memory read and write transfer, it is necessary to provide the hardware with device specifics, such that it can perform the automatic translations. To this end, each memory device has a set of MMIO registers that specify its memory read and write transfers. This specification includes:

- Presence and value of the SPI or HB command byte
- Number of address bytes
- Presence and value of the mode byte
- Number of dummy cycles

In addition, the data transfer widths and data transfer mode (SDR or DDR) are specified.

The XIP interface logic produces an AHB-Lite/AXI bus error under the following conditions:

- The SMIF is disabled (ENABLED bit is set to '0' in CTL register).
- The SMIF is not in XIP_MODE (XIP_MODE bit is set to '0' in CTL register).
- The transfer request is not in a memory region.

- The transfer is a write and the identified memory region does not support writes (WR_EN bit is set to '0' in CTL register).
- In XIP mode (XIP_MODE bit is set to '1' in CTL register) and dual-quad SPI mode (DIV2 bit is set to '1' in the ADDR_CTL register), the transfer address of a write access is not a multiple of 2.
- In XIP mode (XIP_MODE bit is set to '1' in CTL register) and dual-quad SPI mode (DIV2 bit is set to '1' in ADDR_CTL register), the transfer size of a write access is not a multiple of 2.
- In XIP mode (XIP_MODE bit is set to '1' in CTL register) and octal SPI DDR mode or HyperBus mode, the transfer address of a write access is not a multiple of 2 and memory write byte masking is not supported (RWDS_EN bit is set to '0' in WR_DUMMY_CTL register).
- In XIP mode (XIP_MODE is set to '1' in CTL register) and octal SPI DDR mode or HyperBus mode, the transfer size of a write access is not a multiple of 2 and memory write byte masking is not supported (RWDS_EN bit is set to '0' in WR_DUMMY_CTL register).

32.4.1.3 Continuous Transfer Merging

To improve performance of multiple linear continuous transfers (subsequent transfer starts at address following the final address of the previous transfer), the transfers can be merged to a single transfer at the memory interface. This

is especially useful to improve XIP performance over the AXI interface, which splits longer transactions into multiple transfer requests of 16 byte each and allows it to merge these split transfers back to a single transfer at the memory interface. This avoids the overhead of multiple command, address, mode, and especially dummy (latency) cycles. However, not only split AXI transactions, but also sequential AXI or AHB transactions can be merged.

Continuous transfer merging can be done in MMIO and XIP modes.

■ MMIO mode


In MMIO mode, this is under full software control. For each TX, TX_COUNT, or RX_COUNT command in the TX command FIFO it can be specified whether it is the last command. That is, if the memory is deselected after the end of that command processing. This way TX, TX_COUNT, or RX_COUNT commands can be executed in a sequence without deselecting the memory.

■ XIP mode

In XIP mode, the continuous transfer merging can be enabled or disabled per memory device.

If disabled, the XIP block sets the last bit in a TX_COUNT or RX_COUNT causing the memory interface logic to deselect the memory after the transfer is complete. The transfer requests (output from the AHB/AXI interface) and the sequence in the TX command FIFO generated by the XIP block are illustrated in the following figure.


Figure 32-3. Read Sequence with Continuous Merging Disabled



If enabled, the XIP block does not set the last bit in a TX_COUNT or RX_COUNT command for the data phase of a transfer. This causes the memory interface logic to keep the memory device selected (select signals stays active 0 while no clocks are generated).


When a new transfer of the same (read or write) type as the previous is requested from one of the XIP interfaces and its start address is a continuation of the last data phase of the previous transfer, then the XIP interface continues the transfer with a TX_COUNT or RX_COUNT command. This skips the overhead of new command, address, mode, and dummy cycles as illustrated in the following figure.

Figure 32-4. Read Sequence with Continuous Merging Enabled - Merging Always Possible




When a new transfer is requested, which is of a different (read or write) type or its start address is not a continuation of the previous one, then the XIP interface generates a DESELECT command and later starts the new transfer again with the required TX commands for command, address, mode, and/or dummy cycles as illustrated for the last transfer in the following figure.

Figure 32-5. Read Sequence with Continuous Merging Enabled - Merging Possible Once



To avoid keeping the memory selected for a long time while not doing any transfer (which may cause a higher power consumption) per memory device, a continuous transfer merge timeout in CLK_mem cycles can be specified. This timeout value can be 1, 16, 256, 4096, or 65536 CLK_mem cycles. If the timeout occurs before the next transfer is requested the memory device is deselected and a later transfer always starts with command, address, mode, and/or dummy cycles. This is illustrated in the following figure.

Figure 32-6. Read Sequence with Continuous Merging Enabled - merging always possible, but continuous transfer merge timeout occurs




The XIP block contains a 16-bit counter, which starts counting when the last byte of the previous transfer has been written to the TX data FIFO or read from the RX data FIFO (data FIFOs not shown in Figure 32-6). Note that due to the asynchronous clock domain transfer of the commands in the TX command FIFO, the actual remaining time the memory is selected can differ by a few cycles from the specified timeout value.

Additional to the continuous merge timeout there is also a total transfer timeout. This is used for RAM devices requiring refresh cycles. The value needs to be derived from the RAMs maximum transaction length time (tCMS) minus the time of transferring 16-byte data block (data granularity of the XIP ports).

The total timeout period counting is done in the XIP block using CLK_mem cycles. It starts when the first command of a new (not merged) transfer is written to the TX command FIFO causing the interface logic to select the memory. After this period the memory device is deselected as shown in the following figure.

Figure 32-7. Read Sequence with Continuous Merging Enabled
 - merging always possible, but total transfer merge timeout occurs



32.4.2 Off-chip Interfaces

32.4.2.1 Clock Polarity and Phase

SMIF acts as a master for SPI and HyperBus applications. SPI requires the definition of clock polarity and phase (while HyperBus does not). In SPI SDR (single data rate) mode, SMIF supports a single clock polarity and phase configuration:

- Clock polarity (CPOL) is 0: the base value of the clock SPIHB_CLK is "0".
- Clock phase (CPHA) is 0: driving of data is on the falling edge of the clock SPIHB_CLK; capturing of data is specified by CLOCK_IF_RX_SEL bit in the CTL register.

The above configuration is also known as SPI configuration 0 and is supported by SPI memory devices.

32.4.2.2 Specifying Memory Devices

SMIF requires that the memory devices are defined. It supports up to four memory devices. Each device is defined by a set of MMIO registers. The MMIO definition includes:

- The device base address and capacity. The MMIO ADDR register specifies the memory device's base address in the Traveo II address space and the MMIO MASK register specifies the memory device size and capacity. If a memory device is not present or is disabled, the MMIO ADDR and MASK registers specify a memory device with 0 B capacity. Typically, the device address regions in the Traveo II address space are non-overlapping (except dual-quad configuration), to ensure that the activation of select signals is mutually exclusive.
- The device data signal connections (as described in the next section).
- The definition of a read transfer to support XIP mode.
- The definition of a write transfer to support XIP mode.

Each memory device uses a dedicated device select signal: memory device 0 uses SPIHB_SEL[0], memory device 1 uses SPIHB_SEL[1], and so on. In other words, there is a

fixed, one-to-one connection between memory device, MMIO register set, and select signal connection. For example, memory device 0 uses MMIO register set 0 and select signal SPIHB_SEL[0].

In XIP mode, the XIP AHB-Lite bus transfer address is compared with the device region. If the address is within the device region, the device select signal is activated. If a XIP AHB-Lite bus transfer address is within multiple regions (this is possible if the device regions overlap in dual-quad configuration only), all associated device select signals are activated. This overlap enables XIP in dual-quad SPI mode: the command, address, and mode byte can be driven to two quad SPI devices simultaneously.

■ Dual-quad SPI

In XIP mode, dual-quad SPI mode requires the DIV2 field in the ADDR_CTL register of the selected memory devices to be set to '1'. When this field is '1', the transfer address is divided by 2 and the 'divided by 2' address is provided to the memory devices. Each memory device contributes a 4-bit nibble for each 8-bit byte. However, both memory devices are quad SPI memories with a byte interface. Therefore, the memory transfer size must be a multiple of 2 and the memory transfer address must be 2-byte aligned (must also be a multiple of 2).

When the XIP transfer size or address for a read access is not a multiple of 2, then the memory transfer size is extended and/or the memory transfer address is aligned as needed. Then, only the relevant read byte(s) are used and the non-relevant byte(s) are discarded.

Examples:

- An XIP read access to address offset "0" with a length of 1 byte is extended to a memory read access to address offset "0" with a length of 2 bytes.
- An XIP read access to address offset "1" with a length of 1 byte is extended to a memory read access to address offset "0" with a length of 2 bytes.

- An XIP read access to address offset "1" with a length of 2 bytes is extended to a memory read access to address offset "0" with a length of 4 bytes.

When the XIP transfer size or address for a write access is not a multiple of 2, then no memory transfer is done and an XIP_ALIGNMENT_ERROR interrupt cause is set.

The XIP_ALIGNMENT_ERROR interrupt cause is set under the following conditions (in XIP mode and when ADDR_CTL.DIV2 is '1'):

- The transfer address of a write access is not a multiple of 2.
- The transfer size of a write access is not a multiple of 2.

■ Octal SPI DDR

In XIP mode, for octal SPI DDR configuration the SIZE3 field in the ADDR_CTL register needs to be set to '1' or '3' causing a 2-byte or 4-byte address generation. The DIV2 field in the ADDR_CTL register of the selected memory device needs to be set to '0' if the memory expects a byte address (typical case), but can be set to '1' if the memory expects a 16-bit word address. However, the octal SPI memory in DDR configuration has a 16-bit word-based protocol; that is, the smallest addressable item in octal SPI DDR configuration is a 2-byte word. Therefore, the memory transfer size must be a multiple of 2 and the memory transfer byte address must be 2-byte aligned (must also be a multiple of 2).

When the XIP transfer size or address for a read access is not a multiple of 2, then the memory transfer size is extended and/or the memory transfer address is aligned as needed. Then only the relevant read bytes are used and the non-relevant bytes are discarded (examples shown above for dual-quad SPI configuration also apply here).

When the XIP transfer size or address for a write access is not a multiple of 2, then the behavior depends on the memory write byte masking capability. If the memory supports write byte masking by driving RWDS (DQS) = '1' (specified by setting RWDS_EN bit to '1' in WR_DUMMY_CTL register) then the memory transfer size is extended and/or the memory transfer address is aligned (same as for read accesses) and the RWDS (DQS) signal is driven to '1' for the non-relevant byte(s) to avoid that they get written. If the memory does not support write byte masking (specified by setting RWDS_EN bit to '0' in the WR_DUMMY_CTL register), then no memory transfer is done and an XIP_ALIGNMENT_ERROR interrupt cause is set.

The XIP_ALIGNMENT_ERROR interrupt cause is set under the following conditions (in XIP mode):

- The transfer address of a write access is not a multiple of 2 and memory write byte masking is not supported (RWDS_EN bit is set to '0' in the WR_DUMMY_CTL register).

- The transfer size of a write access is not a multiple of 2 and memory write byte masking is not supported (RWDS_EN bit is set to '0' in the WR_DUMMY_CTL register).

■ HyperBus

In XIP mode, for HyperBus configuration the SIZE3 field in the ADDR_CTL register needs to be set to '7' causing a 5-byte address generation with HyperBus protocol (including reserved bits in transaction address fields).

The DIV2 field in the ADDR_CTL register of the selected memory device is ignored (does not matter). However, because the HyperBus is a 16-bit word-based protocol, the XIP byte address is always divided by 2 to generate a HyperBus word address. The smallest addressable item of a HyperBus memory is a 2-byte word. Therefore, the memory transfer size must be a multiple of 2 and the memory transfer byte address must be 2-byte aligned (must also be a multiple of 2).

When the XIP transfer size or address for a read access is not a multiple of 2, then the memory transfer size is extended and/or the memory transfer address is aligned as needed. Then, only the relevant read bytes are used and the non-relevant bytes are discarded (examples shown above for dual-quad SPI configuration also apply here).

When the XIP transfer size or address for a write access is not a multiple of 2, then the memory transfer size is extended and/or the memory transfer address is aligned (same as for read accesses) and the RWDS signal is driven to '1' for the non-relevant bytes to avoid them from being written. In HyperBus configuration RWDS_EN bit in the WR_DUMMY_CTL register must be set to '1' (indicating the byte write masking capability of HyperBus memories); otherwise, an XIP_ALIGNMENT_ERROR interrupt cause is set for unaligned write accesses (same as for octal SPI DDR configuration).

32.4.2.3 Connecting SPI Memory Devices

Memory device I/O signals (SCK, CS#, SI/IO0, SO/IO1, IO2, IO3, IO4, IO5, IO6, IO7) are connected to the SMIF I/O signals (SPIHB_CLK, SPIHB_SEL[3:0] and SPIHB_DATA[7:0]). Not all memory devices provide the same number of I/O signals.

Table 32-1. Memory Device I/O Signals

Memory Device	I/O Signals
Single SPI memory	SCK, CS#, SI, SO. This memory device has two data signals (SI and SO).
Dual SPI memory	SCK, CS#, IO0, IO1. This memory device has two data signals (IO0, IO1).
Quad SPI memory	SCK, CS#, IO0, IO1, IO2, IO3. This memory device has four data signals (IO0, IO1, IO2, IO3).
Octal SPI memory	SCK, CS#, IO0, IO1, IO2, IO3, IO4, IO5, IO6, IO7. This memory device has eight data signals (IO0, IO1, IO2, IO3, IO4, IO5, IO6, IO7).

Table 32-1 illustrates that each memory has a single clock signal SCK, a single (low active) select signal (CS#), and multiple data signals (IO0, IO1, ...).

Each memory device has a fixed select signal connection (to SPIHB_SEL[3:0]).

Each memory device has programmable data signal connections (to SPIHB_DATA[7:0]): the MMIO DATA_SEL[1:0] field in CTLi register specifies how device

data signals are connected. This information is used by SMIF to drive data on the correct SPIHB_DATA[] outputs and capture data from the correct SPIHB_DATA[] inputs. If multiple device select signals are activated, the same data is driven to all selected devices simultaneously.

Not all data signal connections are legal/supported. Supported connections are dependent on the type of memory device.

Table 32-2. Data Signal Connections.

CTL.DATA_SEL[1:0]	Single SPI Device	Dual SPI Device	Quad SPI Device	Octal SPI Device
0	SPIHB_DATA[0]=SI SPIHB_DATA[1]=SO	SPIHB_DATA[0]=IO0 SPIHB_DATA[1]=IO1	SPIHB_DATA[0]=IO0 ... SPIHB_DATA[3]=IO3	SPIHB_DATA[0]=IO0 ... SPIHB_DATA[7]=IO7
1	SPIHB_DATA[2]=SI SPIHB_DATA[3]=SO	SPIHB_DATA[2]=IO0 SPIHB_DATA[3]=IO1	Illegal	Illegal
2	SPIHB_DATA[4]=SI SPIHB_DATA[5]=SO	SPIHB_DATA[4]=IO0 SPIHB_DATA[5]=IO1	SPIHB_DATA[4]=IO0 ... SPIHB_DATA[7]=IO3	Illegal
3	SPIHB_DATA[6]=SI SPIHB_DATA[7]=SO	SPIHB_DATA[6]=IO0 SPIHB_DATA[7]=IO1	Illegal	Illegal


Memory devices can:

- Use shared data signal connections.
- Use dedicated data signal connections. This reduces the load on the data lines, which allows for faster signal level changes. This in turn allows for a faster I/O interface.

Note that dual-quad SPI mode requires dedicated data signals to enable read and/or write data transfer from and to two quad SPI devices simultaneously.

Figure 32-8 illustrates memory device 0, which is a single SPI memory with data signals connections to SPIHB_DATA[1:0].

Figure 32-8. Single SPI Device 0 Connected to SPIHB_DATA[1:0]



The Traveo II pin layout for example, may make it desirable to connect a memory device to specific data lines. Figure 32-9 illustrates memory device 0, which is a single SPI memory with data signals connections to SPIHB_DATA[7:6].

Figure 32-9. Single SPI Device 0 Connected to SPIHB_DATA[7:6]




Figure 32-10 illustrates memory device 1, which is a single SPI memory with data signals connections to SPIHB_DATA[5:4].

Figure 32-10. Single SPI Device 0 Connected to SPIHB_DATA[5:4]




Figure 32-11 illustrates memory devices 0 and 1, both of which are single SPI memories. Each device uses dedicated data signal connections. The device address regions in the Traveo II address space must be non-overlapping to ensure that the activation of SPIHB_SEL[0] and SPIHB_SEL[1] are mutually exclusive.

Figure 32-11. Single SPI Device 0 to SPIHB_DATA[1:0], Single SPI Device 1 Connected to SPIHB_DATA[7:6]




Figure 32-12 illustrates memory devices 0 and 1, both of which are single SPI memories. Both devices use shared data signal connections. The device address regions in the Traveo II address space must be non-overlapping to ensure that the activation of SPIHB_SEL[0] and SPIHB_SEL[1] are mutually exclusive. Note that this solution increases the load on the data lines, which may result in a slower I/O interface.

Figure 32-12. Single SPI Device 0 to SPIHB_DATA[1:0], Single SPI Device 1 Connected to SPIHB_DATA[1:0]




Figure 32-13 illustrates memory device 0, which is a quad SPI memory with data signals connections to SPIHB_DATA[7:4].

Figure 32-13. Quad SPI Device 0 Connected to SPIHB_DATA[7:4]




Figure 32-14 illustrates memory devices 0 and 1, device 0 is single SPI memory and device 1 is quad SPI memory. Each device uses dedicated data signal connections. The device address regions in the Traveo II address space must be non-overlapping to ensure that the activation of SPIHB_SEL[0] and SPIHB_SEL[1] are mutually exclusive.

Figure 32-14. Single SPI Device 0 to SPIHB_DATA[1:0], Quad SPI Device 1 Connected to SPIHB_DATA[7:4]




Figure 32-15 illustrates memory devices 0 and 1, device 0 is single SPI memory and device 1 is quad SPI memory. Both devices use shared data signal connections. The device address regions in the Traveo II address space must be non-overlapping to ensure that the activation of SPIHB_SEL[0] and SPIHB_SEL[1] are mutually exclusive.

Figure 32-15. Single SPI Device 0 to SPIHB_DATA[1:0], Quad SPI Device 1 Connected to SPIHB_DATA[3:0]




Figure 32-16 illustrates memory devices 0 and 1, both of which are quad SPI memories. Each device uses dedicated data signal connections. The device address regions in the Traveo II address space are the same to ensure that the activation of SPIHB_SEL[0] and SPIHB_SEL[1] are the same (in XIP mode). This is known as a dual-quad configuration: during SPI read and write transfers, each device provides a nibble of a byte.

Figure 32-16. Quad SPI Device 0 to SPIHB_DATA[3:0], Quad SPI Device 1 Connected to SPIHB_DATA[7:4]





Figure 32-17 illustrates memory device 0, which is an octal SPI memory with data signal connections to SPIHB_DATA[7:0].

Figure 32-17. Octal SPI Device 0 to SPIHB_DATA[7:0]



32.4.2.4 Connecting HB Memory Devices

Memory device I/O signals (SCK, SCK#, CS#, RWDS, IO0, SO/IO1, IO2, IO3, IO4, IO5, IO6, IO7) are connected to the SMIF I/O signals (SPIHB_CLK, SPIHB_CLK_inv, SPIHB_SEL[3:0], SPIHB_rwds and SPIHB_DATA[7:0]). Not all memory devices use the SCK# but can be operated single-ended.

Table 32-3. Memory Device I/O Signals.

Memory Device	I/O Signals
HyperBus memory	SCK, CS#, IO0, IO1, IO2, IO3, IO4, IO5, IO6, IO7. This memory device has eight data signals (IO0, IO1, IO2, IO3, IO4, IO5, IO6, IO7).


Table 32-3 illustrates that each HyperBus memory has a single clock signal SCK and optional inverted clock signal SCK#, a single (low active) select signal (CS#), a single read-write-data_strobe RWDS, and eight data signals (IO0, IO1, ...).

Each memory device has a fixed select signal connection (to SPIHB_SEL[3:0]).

Each memory device has programmable data signal connections (to SPIHB_DATA[7:0]): the MMIO CTL.DATA_SEL[1:0] field specifies how device data signals are connected. This information is used by SMIF to drive data on the correct SPIHB_DATA[] outputs and capture data from the correct SPIHB_DATA[] inputs. Because HyperBus devices use all eight data signals, the only valid setting for DATA_SEL[1:0] is 0.

Figure 32-18 illustrates memory device 0, which is a single HyperBus memory with data signals connections to SPIHB_DATA[7:0].

Figure 32-18. HyperBus SPI Device 0 to SPIHB_DATA[7:0]



32.4.2.5 SPI Data Transfer

SPI data transfer uses most-significant-byte (MSB) first data transfer. This means that for a byte B, consisting of bits b7, b6, ..., b0, bit b7 is transferred first, followed by bit b6, and so on. For dual, quad, dual-quad, and octal SPI transfers, multiple bits are transferred per cycle. For a single SPI device and device data signal connections to SPIHB_DATA[1:0] (SPIHB_SEL is 0), [Table 32-4](#) summarizes the transfer of a byte B.

Table 32-4. Single Data Transfer

Cycle	Data Transfer
0	For a write transfer: b7 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b7 is transferred on SPIHB_DATA[1] and SO/IO1.
1	For a write transfer: b6 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b6 is transferred on SPIHB_DATA[1] and SO/IO1.
2	For a write transfer: b5 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b5 is transferred on SPIHB_DATA[1] and SO/IO1.
3	For a write transfer: b4 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b4 is transferred on SPIHB_DATA[1] and SO/IO1.
4	For a write transfer: b3 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b3 is transferred on SPIHB_DATA[1] and SO/IO1.
5	For a write transfer: b2 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b2 is transferred on SPIHB_DATA[1] and SO/IO1.
6	For a write transfer: b1 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b1 is transferred on SPIHB_DATA[1] and SO/IO1.
7	For a write transfer: b0 is transferred on SPIHB_DATA[0] and SI/IO0. For a read transfer: b0 is transferred on SPIHB_DATA[1] and SO/IO1.

Note that in single SPI data transfer, the SPIHB_data signals are unidirectional: in the table, SPIHB_DATA[0] is exclusively used for write data connected to the device SI input signal and SPIHB_DATA[1] is exclusively used for read data connected to the device SO output signal.

For a dual SPI device and device data signal connections to SPIHB_DATA[1:0] (SPIHB_SEL is 0), [Table 32-5](#) summarizes the transfer of a byte B.

Table 32-5. Dual Data Transfer

Cycle	Data Transfer
0	b7, b6 are transferred on SPIHB_DATA[1:0] and IO1, IO0.
1	b5, b4 are transferred on SPIHB_DATA[1:0] and IO1, IO0.
2	b3, b2 are transferred on SPIHB_DATA[1:0] and IO1, IO0.
3	b1, b0 are transferred on SPIHB_DATA[1:0] and IO1, IO0.

For a quad SPI device and device data signal connections to SPIHB_DATA[3:0] (SPIHB_SEL is 0), [Table 32-6](#) summarizes the transfer of a byte B.

Table 32-6. Quad Data Transfer

Cycle	Data Transfer
0	b7, b6, b5, b4 are transferred on SPIHB_DATA[3:0] and IO3, IO2, IO1, IO0.
1	b3, b2, b1, b0 are transferred on SPIHB_DATA[3:0] and IO3, IO2, IO1, IO0.

For a octal SPI device and device data signal connections to SPIHB_DATA[7:0] (SPIHB_SEL is 0), [Table 32-7](#) summarizes the transfer of a byte B.

Table 32-7. Octal Data Transfer

Cycle	Data Transfer
0	b7, b6, b5, b4, b3, b2, b1, b0 are transferred on SPIHB_DATA[7:0] and IO7, IO6, IO5, IO4, IO3, IO2, IO1, IO0.

In dual-quad SPI mode, two quad SPI devices are used.

- The first device (the device with the lower device structure index) should have device data signal connections to SPIHB_DATA[3:0] and must be connected to SPIHB_SEL[0].
- The second device (the device with the higher device structure index) should have device data signal connection to SPIHB_DATA[7:4] and must be connected to SPIHB_SEL[0].

The “command” and “data” phases of the SPI transfer use different width data transfers:

- The command, address, and mode byte use quad SPI data transfer.
- The read data and write data use octal data transfer. Each device provides a nibble of each data byte: the first device provides the lower nibble and the second device provides the higher nibble.

[Table 32-8](#) summarizes the transfer of a read data and write data byte B.


Table 32-8. Dual-quad SPI Mode, Octal Data Transfer

Cycle	Data Transfer
0	b7, b6, b5, b4 are transferred on SPIHB_DATA[7:4] and 2nd device IO3, IO2, IO1, IO0. b3, b2, b1, b0 are transferred on SPIHB_DATA[3:0] and 1st device IO3, IO2, IO1, IO0.

32.4.2.6 SPI - Putting it All Together

Devices 0 and 1 are used to implement dual-quad SPI mode. Both devices are 1MB/8Mb devices; that is, the address requires 3 bytes. Device 0 has device data signal connections to SPIHB_DATA[3:0] and device 1 has device data signal connections to SPIHB_DATA[7:4].

Figure 32-19. Dual-Quad SPI Mode by Connecting Quad SPI Device 0 to SPIHB_DATA[3:0], Quad SPI Device 1 Connected to SPIHB_DATA[7:4]



General Settings

The clock settings are set for SDR timing.

```
MMIO_SMIF_CTL = (1UL << 31) // ENABLED
| (3 << 12) // CLOCK_IF_RX_SEL: "spi_clk_in" (feedback clock) for SDR
| (0 << 10) // DDR_CAPTURE_CYCLE, not used
| (0 << 9) // INT_CLOCK_DL_ENABLED, not used
| (0 << 8) // INT_CLOCK_DEL_TAP_ENABLED, not used
| (0 << 4) // CLOCK_IF_TX_SEL: "clk_if_tx_div_inv" for DDR
| (0 << 0); // MMIO_MODE: select MMIO/XIP mode
```

For dual-quad SPI mode, the AHB-Lite bus transfer address is divided by two. Cryptography and write functionality are disabled:

```
DEV0_ADDR = CPUSS_SMIF_BASE;
DEV0_MASK = 0xfff00000; // MASK: 1 MB region
DEV0_CTL = (1UL << 31) // ENABLED
| (0 << 28) // TOTAL_TIMEOUT_EN
| (0 << 16) // TOTAL_TIMEOUT
| (0 << 15) // MERGE_EN
| (0 << 12) // MERGE_TIMEOUT
| (0 << 8) // DATA_SEL: spi_data[3:0]
| (0 << 4) // CRYPTO_EN
| (0 << 0); // WR_EN
DEV0_ADDR_CTL = (1 << 8) // DIV2: enabled
| ((3-1) << 0); // SIZE: 3 B address


DEV1_ADDR = CPUSS_SMIF_BASE;
DEV1_MASK = 0xfff00000; // MASK: 1 MB region
DEV1_CTL = (1UL << 31) // ENABLED
| (0 << 28) // TOTAL_TIMEOUT_EN
| (0 << 16) // TOTAL_TIMEOUT
```

```

| (0 << 15)          // MERGE_EN
| (0 << 12)          // MERGE_TIMEOUT
| (2 << 8)           // DATA_SEL: spi_data[7:4]
| (0 << 4)           // CRYPTO_EN
| (0 << 0);          // WR_EN
DEV1_ADDR_CTL = (1 << 8) // DIV2: enabled
| ((3-1) << 0)); // SIZE: 3 B address
  
```

For XIP read transfers, the “0xeb” command/instruction is used (Figure 32-20 illustrates a two byte transfer from devices 0 and 1 in dual-quad SPI mode).

Figure 32-20. 0xeb Instruction, Instruction 1-bit/cycle; Address, Mode, Data 4-bits/cycle



The definition of a read transfer is as follows:

```

DEV0_RD_CMD_CTL = (1UL << 30) // PRESENT
| (0 << 16) // WIDTH: single data transfer
| 0xeb; // CODE
DEV0_RD_ADDR_CTL = (2 << 16); // WIDTH: quad data transfer
DEV0_RD_MODE_CTL = (1UL << 31) // PRESENT
| (2 << 16) // WIDTH: quad data transfer
| 0x00; // CODE
DEV0_RD_DUMMY_CTL = (1UL << 30) // PRESENT
| ((4-1) << 0); // SIZE: 4 dummy cycles
DEV0_RD_DATA_CTL = (3 << 16); // WIDTH: octal data transfer
  
```

Note that the command uses single data transfer, the address and mode byte use quad data transfer, and the read data byte uses octal data transfer. All transaction fields use SDR (single data rate) mode.

32.4.2.7 SPI- Slave Select Signal During Power up

Typically, SPI device datasheets specify that the chip select (CS#, which is SPIHB_SEL[]) must follow Vcc applied to the device. This can be achieved by adding a weak pull-up on slave select pin at board level.

32.4.2.8 HB Data Transfers

HyperBus is 2-byte addressed. Therefore, data transfers start with the most-significant byte (MSB) followed by the least-significant byte (LSB) of an addressed 2-byte word.

This means that for 2 bytes B1 and B0, consisting of bits b15, b14, ..., b0, bit b15 to b8 is transferred first, followed by bit b7 to bit b0. [Table 32-9](#) summarizes the transfer of 2 bytes B1 and B0 for a single HyperBus device.

Table 32-9. Single data transfer.

Half Cycle	Data Transfer
0	b15 to b8 are transferred on SPIHB_DATA[7:0].
1	b7 to b0 are transferred on SPIHB_DATA[7:0].

■ HB variable initial latency

HyperRAM memory devices have the option of a variable initial latency. If the memory is doing a refresh cycle a double initial latency is signaled using an active RWDS indicator during the command/address cycles. This is supported in the SMIF.

To enable the HyperRAM variable initial latency mode in MMIO mode bit 16 of the DUMMY_COUNT command in the TX command FIFO needs to be set to '1'.

Note: The SMIF TX interface must consider that the HyperBus latency cycle definition includes the last address cycle while the dummy cycles specified by bits 4:0 of the DUMMY_COUNT commands in the TX command FIFO do not include that.

Example: The single HyperRAM latency count may be six cycles, so the doubled HyperRAM latency cycles count for a refresh cycle is 12 cycles.

The dummy cycle count specified in the DUMMY_COUNT command excludes the last address cycle. That is, bits 4:0 of the DUMMY_COUNT command need to be set to '4' (defining five dummy cycles). If the variable initial latency mode is enabled and the RWDS refresh indicator is active, the SMIF TX interface needs to double the latency cycles. It needs to set the dummy cycle count to 11 $((4+2)*2 - 1)$.

To enable the HyperRAM variable initial latency mode in XIP mode the PRESENT2 field must be set to '2' in the related RD/WR_DUMMY_CTL registers.

If enabled, the SMIF XIP block sets bit 16 of the DUMMY_COUNT command in the TX command FIFO to '1'.

■ HB page boundary crossing latency

HyperFlash memory devices may require page boundary crossing latency cycles.

In today's HyperFlash devices they only apply at the first page boundary crossing or more precisely when the second half-page boundary is crossed. However, future HyperFlash devices may need them for subsequent page boundary crossings. Also, future HyperFlash devices may have a different number of sub pages per page (other than 2).

The presence and number of page boundary crossing latency cycles depend on the latency count and the start

address of the burst transaction. The following two tables show examples of memories with eight words (16 bytes) per sub page, two sub pages per page, and a latency count of 11, 16, or 20 cycles (depending on the frequency).

First Boundary Crossing During Linear Read (Latency Count = 16 Clocks)																																									
Target Address	Clock Cycle After CS# Goes Low																																								
	0	1	2	3	...	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39													
0	CA0	CA1	CA2	Bus Turnaround + Initial Latency		D0	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21														
1					D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	D16	D17	D18	D19	D20	D21															
2					D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	D16	D17	D18	D19	D20	D21															
3					D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	D16	D17	D18	D19	D20	D21															
4					D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	D16	D17	D18	D19	D20	D21															
5					D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	D16	D17	D18	D19	D20	D21															
6					D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	D16	D17	D18	D19	D20	D21															
7					D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	X	D16	D17	D18	D19	D20	D21															
8					D8	D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	D24	D25	D26	D27	D28	D29															
9					D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	D24	D25	D26	D27	D28	D29															
10					D10	D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	D24	D25	D26	D27	D28	D29															
11					D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	D24	D25	D26	D27	D28	D29															
12					D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	D24	D25	D26	D27	D28	D29															
13					D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	D24	D25	D26	D27	D28	D29															
14					D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	D24	D25	D26	D27	D28	D29															
15					D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	X	D24	D25	D26	D27	D28	D29															
16					D16	D17	D18	D19	D20	D21	D22	D23	D24	D25	D26	D27	D28	D29	D30	D31	D32	D33	D34	D35	D36	D37															
	-	-	1	2	...	16	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-													
	Latency Count																																								

First Boundary Crossing During Linear Read (Latency Count = 20 Clocks)																																															
Target Address	Clock Cycle After CS# Goes Low																																														
	0	1	2	3	...	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43																			
0	CA0	CA1	CA2	Bus Turnaround + Initial Latency		D0	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	D16	D17																				
1					D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	D16	D17																					
2					D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	D16	D17																					
3					D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	X	D16	D17																					
4					D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	X	X	D16	D17																					
5					D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	X	X	X	D16	D17																					
6					D6	D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	X	X	X	X	D16	D17																					
7					D7	D8	D9	D10	D11	D12	D13	D14	D15	X	X	X	X	X	X	X	X	X	X	X	D16	D17																					
8					D8	D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	D24	D25																					
9					D9	D10	D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	D24	D25																					
10					D10	D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	D24	D25																					
11					D11	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	X	D24	D25																					
12					D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	X	X	D24	D25																					
13					D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	X	X	X	D24	D25																					
14					D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	X	X	X	X	D24	D25																					
15					D15	D16	D17	D18	D19	D20	D21	D22	D23	X	X	X	X	X	X	X	X	X	X	X	D24	D25																					
16					D16	D17	D18	D19	D20	D21	D22	D23	D24	D25	D26	D27	D28	D29	D30	D31	X	X	X	X	D32	D33																					
	-	-	1	2	...	20	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-		
	Latency Count																																														

In general, the page boundary crossing latency can be calculated by:

```

if ((page_size - base_latency) < Start_Addr & (sub_page_size - 1))
{ ((Start_Addr & (sub_page_size - 1)) - page_size + base_latency) }
else
{ 0 }

```

The SMIF supports the page boundary crossing latencies.

In MMIO mode this is completely under software control (RX_DATA_COUNT and DUMMY_COUNT commands must be written into the TX command FIFO as required by the memory).

In XIP mode the page boundary crossing behavior is defined by the RD_BOUND_CTL register. This only applies when continuous transfer merging is enabled (otherwise, the longest possible transaction only reads 16 bytes = 8 words and therefore, cannot cross a page boundary).

The following variables can be specified:

- ❑ The base latency cycle count (latency cycle count of the memory, which may be different from the dummy cycle count specified in the RD_DUMMY_CTL register due to the last address cycle not included)
- ❑ The size of a memory sub page

- ❑ The number of sub pages (banks) per page
- ❑ The page boundary crossing latency is only generated at the first page boundary crossing (the first time a SUB_PAGE_NRth sub page boundary is crossed). For example, with two sub pages per page when the second sub page boundary is crossed. A page boundary crossing latency can also be generated at the first and subsequent page boundary crossings (every time a SUB_PAGE_NRth sub page boundary is crossed). For example, with two sub pages per page when the second, fourth, sixth, ... sub page boundary is crossed.

32.4.2.9 HB - Putting it All Together

One device is used to implement HB mode. HB devices require 6 bytes for command and address including reserved bits. Device 0 has device data signal connections to SPIHB_DATA[7:0].

Figure 32-21. HyperBus Device Connection





Figure 32-22. HyperBus Transfer Waveform



■ General settings

The clock settings are set for DDR timing with read-write-data-strobe (RWDS).

```
MMIO_SMIF_CTL = (1UL << 31) // ENABLED
               | (6 << 12) // CLOCK_IF_RX_SEL: "spi_rwds" for DDR with RWDS
               | (0 << 10) // DDR_CAPTURE_CYCLE, not used
               | (0 << 9)  // INT_CLOCK_DL_ENABLED, not used
               | (0 << 8)  // INT_CLOCK_DEL_TAP_ENABLED, not used
               | (1 << 4)  // CLOCK_IF_TX_SEL: "clk_if_tx_inv_div" for DDR
               | (0 << 0); // MMIO_MODE: select MMIO/XIP mode

MMIO_SMIF_DELAY_TAP_SEL = (0 << 28) // BIT7, not used
                          | (0 << 24) // BIT6, not used
                          | (0 << 20) // BIT5, not used
                          | (0 << 16) // BIT4, not used
                          | (0 << 12) // BIT3, not used
                          | (0 << 8)  // BIT2, not used
                          | (0 << 4)  // BIT1, not used
                          | (8 << 0); // BIT0, used for RWDS delay (for all bits)
```

Note that the delay tap setting is based on calibration by software

■ MMIO mode

The definition of a read transfer in MMIO mode is as follows (write commands to TX command FIFO via TX_CMD_FIFO_WR):

```

1st command word:  (0 << 24)          // TX command
                   | ((1 << 2) << 20) // slave 2
                   | (0 << 19)         // NOT the last command of transfer
                   | (1 << 18)         // DDR mode
                   | (3 << 16)         // WIDTH: 8 bits/cycle
                   | 0xa000           // CMD (HB Read/MEM/Lin. opcode), word address[31:19]
2nd command word:  (0 << 24)          // TX command
                   | ((1 << 2) << 20) // slave 2
                   | (0 << 19)         // NOT the last command of transfer
                   | (1 << 18)         // DDR mode
                   | (3 << 16)         // WIDTH: 8 bits/cycle
                   | 0x4000           // word address[18:3]
3rd command word:  (0 << 24)          // TX command
                   | ((1 << 2) << 20) // slave 2
                   | (0 << 19)         // NOT the last command of transfer
                   | (1 << 18)         // DDR mode
                   | (3 << 16)         // 8 bits/cycle
                   | 0x0000           // Reserved bits + word address[2:0]
4th command word:  (3 << 24)          // DUMMY_COUNT command
                   | (0 << 17)         // RWDS (write mask) output generation
                   | (0 << 16)         // variable latency based on RWDS refresh// indicator
                   | ((15-1) << 0)    // (LC-1) = 15 dummy cycles (16 HB LC)
5th command word:  (2 << 24)          // RX_COUNT command
                   | (3 << 16)         // 8 bits/cycle
                   | ((4-1) << 0)     // 4 Cycles (8 bytes)
  
```

The data words can be read from RX data FIFO via RX_DATA_MMIO_FIFO_RD2/RX_DATA_MMIO_FIFO_RD4 registers.

■ XIP mode

The definition of a read transfer in XIP mode is as follows:

```

DEV0_ADDR          = CPUSS_SMIF_BASE;
DEV0_MASK          = 0xf0000000;      // MASK: 256 MB region
DEV0_CTL           = (1UL << 31)      // ENABLED
                   | (0 << 28)        // TOTAL_TIMEOUT_EN
                   | (0 << 16)        // TOTAL_TIMEOUT
                   | (0 << 15)        // MERGE_EN
                   | (0 << 12)        // MERGE_TIMEOUT
                   | (0 << 8)         // DATA_SEL: spi_data[7:0]
                   | (0 << 4)         // CRYPTO_EN
                   | (0 << 0));       // WR_EN
DEV0_ADDR_CTL      = (0 << 8)        // DIV2: disabled
                   | (7 << 0));       // SIZE: 5 B address w/ HB protocol
DEV0_RD_CMD_CTL    = (1UL << 30)      // PRESENT
                   | (1 << 18)        // DDR mode
                   | (3 << 16)        // WIDTH: 8 bits/cycle (octal data transfer)
                   | 0xa0;           // CODE HB (Read/MEM/Lin. transfer opcode)
DEV0_RD_ADDR_CTL   = (1 << 18)        // DDR mode
                   | (3 << 16));       // WIDTH: 8 bits/cycle (octal data transfer)
DEV0_RD_MODE_CTL   = (0 << 31)        // NOT PRESENT
DEV0_RD_DUMMY_CTL  = (1UL << 30)      // PRESENT
                   | ((15-1) << 0)); // SIZE: (LC-1) = 15 dummy cycles (16 HB LC)
DEV0_RD_DATA_CTL   = (1 << 18)        // DDR mode
                   | (3 << 16));       // WIDTH: 8 bits/cycle (octal data transfer)
  
```


Note that the RX (or TX) count is loaded directly from the AHB/AXI bus.

32.4.3 AXI Interface

The SMIF provides an XIP AXI interface. It is intended to be used by the M7 CPU.

Figure 32-23 gives an overview of the SMIF AXI interface.

Figure 32-23. AXI Slave Interface Block



The SMIF AXI interface contains four FSMs.

- AXI front-end FSM, responsible for
 - Allocating the buffer slots
 - Writing write transaction data (from AXI write data channel) into the buffer and forwarding write transaction control information to buffer (write transaction control FSM)
 - Initiating block read transactions to the serial memory interface (including translation of wrapping read bursts to incrementing read bursts)
- AXI back-end FSM, responsible for
 - Reading read transaction data from the buffer and forwarding that to the AXI read data channel (including backward translation of incrementing read bursts to wrapping read bursts)
 - Generating the AXI write data response
- Buffer write transaction control FSM, responsible for
 - Reading write transaction data from the buffer and forwarding them to the serial memory interface
 - Initiating block write transactions to the serial memory interface (including translation of wrapping write bursts to incrementing write bursts)
- Buffer read transaction control FSM, responsible for
 - Writing read transaction data from the serial memory interface to the buffer

The SMIF AXI interface contains an SRAM-based buffer of 576 bytes. The buffer is used to hold outstanding transaction data and (in case of writes accesses) write strobe information. The buffer has four slots of 128 bytes data and 16 bytes write strobes each. Each slot is further divided into blocks of 8 bytes of data and 1 byte strobe, which is the maximum size of one AXI transfer (one beat of an AXI burst). For every block a related valid flag exists. For write transactions the communication between AXI front-end FSM and buffer write transaction control FSM as well as for read transactions the communication between AXI back end FSM and buffer read transaction control FSM is done via these valid flags.

32.4.3.1 Support for Multiple Outstanding Transactions

Multiple outstanding transactions are supported to make optimum use of the bandwidth given by the SMIF and not waste time for SMIF internal logic. This means, the transactions are pipelined and includes the following aspects:

- While one read or write transaction is ongoing at the serial memory interface, the following transaction (when available from the AXI interface) is already prepared.
- If the AXI read data channel is temporarily stalled, the previously read data is stored and the serial memory interface already serves the next read or write transaction.

At least three outstanding transactions should be supported: One read transaction is finished and needs to be temporarily stored (due to stalled AXI read data channel), a second read or write transaction is served by the serial memory interface and a third read or write transaction is currently prepared to be served immediately after the second one.

These three outstanding transactions are rounded up to the next power-of-two number which is '4'. Therefore, a buffer that provides storage for four AXI transactions is used.

Note: The number of at least three outstanding transactions here represents only the minimum number of outstanding transactions the SMIF XIP pipeline and therefore the AXI buffer needs to support. However, due to the AXI address channel input FIFO and the AXI read/write response channel output FIFOs, the maximum total outstanding transactions (including these waiting AXI channel FIFO entries) can be higher.

The SMIF AXI interface does not reorder transactions. The transactions are processed in the same order as they occur at the (merged) AXI address channel.

■ Bandwidth

The SMIF can achieve a maximum possible memory interface bandwidth of 400 Mbyte/s when using an 8-bit wide memory interface with double data rate (DDR) at a maximum speed of 200 MHz (HyperBus or octal SPI with DDR). With the ability of merging continuous transfers and its pipelined operation the SMIF is designed to support this bandwidth and not generate stall cycles. However, to limit the area overhead for FIFOs in the AXI interface and XIP block, the maximum bandwidth is only supported for AXI data streams consisting of transactions of at least 16 bytes. For transactions of only 8 bytes or smaller the streaming performance is not optimized. The assumption is that bandwidth-critical AXI masters will use outstanding transactions with a total transaction size of at least 16 bytes (for example, CM7 cache line fills are 32-byte bursts).

After initial memory protocol overhead cycles (command, address, and latency cycles) the full theoretically possible bandwidth of 400 Mbyte/s can be achieved as long as a continuous (linear) data stream is read or written.

For non-continuous transactions, the bandwidth is significantly lower because the overhead cycles need to be sent repeatedly.

32.4.3.2 Wrapping Bursts

HyperBus memories and many SPI memories support wrapping bursts. Additionally, for SPI memories, a memory control register needs to be written whenever there is a change from wrapping to incrementing bursts or vice versa. Also, for SPI and HyperBus memories, a memory control register needs to be written whenever there is a change in the wrapping boundary. This is not possible in XIP mode (and does not fit into the concept of the XIP mode).

32.4.3.3 Fixed Bursts

Fixed bursts to a memory address region is uncommon, but they are supported to be AXI compliant.

For read transactions the SMIF performs the read transfer to the memory only once but its AXI interface provides the same data multiple times as requested by the length of the fixed burst. For write transactions the SMIF performs a write transfer to the memory only for the last AXI write transfer.

32.4.3.5 Write Strobes

The AXI protocol allows write byte masking, that is, write transactions with any combination of write strobes. Therefore, any combination of actual bytes to be written is possible within a transaction. The CM7 store buffer uses it to merge multiple store instructions to normal memory as a performance optimization.

Example 1:

■ CM7 code:

```
MOV r0, #0x4000
STR r1, [r0, #0xC]; Store a word at 0x400C
```

■ This can result in the following write transaction

- address 0x4008, one transfer, size 8 bytes (asize = 3), write strobes (wstrb) set to 0xF0

■ Here only the write strobes for the higher 4 bytes are enabled within an 8-byte transfer.

Example 2:

■ CM7 code:

```
MOV r0, #0x4000
STRH r1, [r0, #0x18]; Store a halfword at 0x4018
STR r2, [r0, #0xC] ; Store a word at 0x400C
STMIA r0, {r4-r7} ; Store four words at 0x4000
STRB r3, [r0, #0x1D]; Store a byte at 0x401D
```

■ These instructions can result in the following write transaction:

- address 0x4000, four transfers, size 8 bytes (asize = 3), write strobes (wstrb) set to 0xFF, 0xFF, 0x00, and 0x23

■ The four store instructions are merged to a single AXI transaction by the CM7 store buffer. The second store instruction is completely skipped (because its target address is overwritten by the third store instruction). The write strobes are not adjacent and even a transfer without any write strobe set is generated.

HyperBus and some octal SPI RAM devices support byte masking with RWDS or DQS signals. When writing data to an external HyperBus or octal SPI RAM, this can be used to directly translate any AXI write transaction to a single memory burst.

However, other SPI devices do not support byte masking. To support such RAM devices an AXI write transaction with non-adjacent write strobes should be split into multiple SPI memory bursts.

32.4.3.4 Splitting Bursts to Smaller Blocks

The SMIF has multiple XIP interfaces (at least one AHB and one AXI interface). A priority-based arbitration is done between the interfaces. To not block high-priority transactions (from a latency critical master) by lower priority long transactions (from a high-bandwidth master), interrupting lower priority transactions is preferred.

Longer transactions are therefore split to multiple blocks, which can be potentially interrupted by the arbiter. The block size is 16 bytes. This matches the block size of the cache subsector fetch or pre-fetch in the AHB interface and the block size of an AES encryption/decryption.

These RAM devices play almost no role in the market, but for serial RAM devices only HyperBus and octal SPI RAM devices are relevant. Therefore, AXI write byte masking for SPI RAM devices without byte masking is not supported in SMIF (when RWDS_EN bit is set to '0' in WR_DUMMY_CTL register). This means an error response is generated when the memory device does not support write byte masking and not all of the AXI write strobes are enabled according to the transfer size.

However, such RAM devices can still be used with the CM7 as strongly-ordered or device memory. The SMIF address space is usually located in a memory region with a normal default memory type, but the CM7 allows it to override default memory types. That means, if an SPI RAM device without byte masking is used, the MPU must be configured to override the memory type of that RAM device region from Normal to Strongly-ordered or Device memory. This prevents CM7 to merge multiple store instructions and ensures that all AXI write channel write strobe bits are enabled according to the transfer size. Then the above examples look as follows in case of a Strongly-ordered or Device memory region:

Example 1:

■ CM7 code:

```
MOV r0, #0x4000
STR r1, [r0, #0xC]; Store a word at 0x400C
```

■ This results in the following write transaction

- address 0x400C, one transfer, size 4 bytes (asize = 2), write strobes (wstrb) set to 0xF0

■ Here all four write strobes are enabled within a 4-byte transfer.

Example 2:

■ CM7 code:

```
MOV r0, #0x4000
STRH r1, [r0, #0x18]; Store a halfword at 0x4018
STR r2, [r0, #0xC]; Store a word at 0x400C
STMIA r0, {r4-r7}; Store four words at 0x4000
STRB r3, [r0, #0x1D]; Store a byte at 0x401D
```

■ This results in the following five write transactions:

- address 0x4018, one transfer, size 2 bytes (asize = 1), write strobes (wstrb) set to 0x03
- address 0x400C, one transfer, size 4 bytes (asize = 2), write strobes (wstrb) set to 0xF0
- address 0x4000, two transfers, size 4 bytes (asize = 2), write strobes (wstrb) set to 0x0F, 0xF0
- address 0x4008, two transfers, size 4 bytes (asize = 2), write strobes (wstrb) set to 0x0F, 0xF0
- address 0x401D, one transfer, size 1 byte (asize = 0), write strobes (wstrb) set to 0x01

■ The store instructions are causing individual AXI transactions. No skipping or merging is done. All write strobes are set according to the transfer size.

It is expected that other AXI masters (such as a graphic block) do not use write byte masking.

32.4.4 Triggers

The SMIF has two level-sensitive triggers:

- `tr_tx_req` is associated with the TX data FIFO.
- `tr_rx_req` is associated with the RX data FIFO.

If the SMIF is enabled (ENABLED is set to '1' in the CTL register) and MMIO operation mode is selected (XIP_MODE is set to '0' in the CTL register), the trigger functionality is enabled. The trigger functionality is defined as follows:

- The TRIGGER_LEVEL field in TX_DATA_FIFO_CTL register specifies a number of FIFO entries. The `tr_tx_req` trigger is active when the number of used TX data FIFO entries is smaller than or equal to the specified number; that is, USED4 in TX_DATA_FIFO_STATUS register $\leq$ TRIGGER_LEVEL.
- The TRIGGER_LEVEL field in RX_DATA_FIFO_CTL register specifies a number of FIFO entries. The `tr_rx_req` trigger is active when the number of used RX data FIFO entries is greater than the specified number; that is, USED4 in RX_DATA_FIFO_STATUS register $>$ TRIGGER_LEVEL.

32.4.5 Interrupts

The SMIF has a single interrupt output. This interrupt has three interrupt causes:

- TR_TX_REQ in INTR register. This interrupt cause is activated in MMIO mode, when the `tr_tx_req` trigger is activated.
- TR_RX_REQ in INTR register. This interrupt cause is activated in MMIO mode, when the `tr_rx_req` trigger is activated.
- XIP_ALIGNMENT_ERROR in INTR register. This interrupt cause is activated in XIP mode, when dual-quad SPI mode or octal SPI DDR/HyperBus mode without memory write byte masking is selected and the AHB-Lite/AXI bus address is not a multiple of 2 or the requested transfer size is not a multiple of 2. This interrupt cause identifies erroneous behavior in dual-quad SPI mode (the selected device DIV2 field is set to '1' in the ADDR_CTL register), Octal SPI DDR mode, or HyperBus mode.

This interrupt cause is activated in XIP mode when the selected memory device does not support write byte masking (RWDS_EN is set to '0' in WR_DUMMY_CTL register) and an AXI transfer occurs with not all of the AXI write strobes (wstrb) enabled according to the transfer size (assize).

- TX_CMD_FIFO_OVERFLOW in INTR register. This interrupt cause is activated in MMIO mode, on an AHB-Lite write transfer to the TX command FIFO (TX_CMD_FIFO_WR) with not enough free entries available.
- TX_DATA_FIFO_OVERFLOW in INTR register. This interrupt cause is activated in MMIO mode, on an AHB-

Lite write transfer to the TX data FIFO (TX_DATA_FIFO_WR1, TX_DATA_FIFO_WR2, TX_DATA_FIFO_WR4) with not enough free entries available.

- RX_DATA_FIFO_OVERFLOW in INTR register. This interrupt cause is activated in MMIO mode, on an AHB-Lite read transfer from the RX data FIFO (RX_DATA_FIFO_RD1, RX_DATA_FIFO_RD2, RX_DATA_FIFO_RD4) with not enough entries available.

32.4.6 Monitor Signals

The SMIF has five monitor signals. These signals reflect active SPI transfers:

- The monitor_smif_SPIHB_SEL[i] (i = 0, 1, 2, 3) signal is active '1' during SPI transfers to memory device i (SPIHB_select_out[i] is '0'). In other words, the monitor signals are the logical inverse of the SPI select signals.
- The monitor_smif_SPIHB_select_any signal is the logical OR of the monitor_smif_SPIHB_SEL[] signals; that is, it is active/'1' when any of the monitor_smif_SPIHB_SEL[] signals is active/'1'.

The monitor signals are driven by dedicated flip-flops that are driven by the negative edge of the transmitter clock, CLK_if_tx_inv.

The monitor signals are connected to an energy profiler, which accumulates the duration of active SPI transfers. The energy profiler and its associated software use the SPI transfer activity as a proxy for the amount of energy consumed by the SMIF.

32.5 Supply Rails and Power Domains

32.5.1 Power Modes

Active, Sleep, DeepSleep, and Hibernate are the different power modes defined in SRSS. The following table describes the status of SMIF during different power modes.

Table 32-10. SRSS Power Modes

System Power Mode	Description
ACTIVE/SLEEP	Active and Sleep are standard Arm-defined power modes, supported by the Arm CPUs and ISA. Active: In this mode, CPUs will execute code, and all the logic and memories are powered on. Sleep: In this mode, CPUs will not execute code and the clock is stopped. All the logic and memories are powered on. It is identical to Active power mode from the peripheral point of view.
LPACTIVE/LPSLEEP	LPActive/LPSleep are similar to Active/Sleep, except the current is limited and clocks may run at a lower frequency (and some functions are not available or limited.) Transitions between Active and LPActive (or Sleep and LPSleep) are initiated by SRSS PWR_ - CONTROL.LOW_POWER control bit.
DEEPSLEEP	DeepSleep is a lower power mode where high-frequency clocks are disabled. Active power domain is powered off (vccact power supply is off). CPUs and most MMIO state is retained (through retention flops). System SRAM retains its data. Flash and ROM memories are powered off.
HIBERNATE	Hibernate is an even lower power mode than DeepSleep, but on wakeup the CPUs (and all peripherals) go through a full reset and firmware reboot. Note: This power mode is more similar to Stop power mode in M0S8 platform.
XRES	XRES is the state of the device when external reset is applied.
OFF	OFF state is the state of the device when no power is applied to it.

The SMIF is an Active functionality. In DeepSleep power mode, the retention MMIO registers are retained. Note that the cache and the AXI interface buffer are not retained in DeepSleep power mode. When exiting DeepSleep power mode, the cache and buffer are reset.

When entering DeepSleep power mode, it is desirable to put the external memory devices in low-power mode as well (if this mode is supported by the devices). Similarly, when exiting DeepSleep mode, the external memory devices should exit their low-power mode. Entering and exiting low-power mode is device-specific. This is supported through the MMIO mode. This means that if SMIF is in XIP mode, a change to MMIO is required before entering the device low-power mode.

32.6 Sub Block Descriptions

32.6.1 Address Space

SMIF has three AHB-Lite slave interfaces:

- Fast and slow XIP interfaces have a shared design time configurable address space. This address space supports XIP_MODE mode of operation and is (partially) populated by the external devices.
- The MMIO interface has a 4 Kbytes address space. This address space supports MMIO_MODE of operation. This address space includes all the MMIO registers (which also provides access to the TX and RX FIFOs).

The following sections describe the address spaces in more detail.

32.6.1.1 XIP Address Space

AHB-Lite transfers to the XIP address space either access the cache or are translated on-the-fly into SPI transfers to the external device. SMIF exposes an address space located at Traveo II address 0x6000:0000. The XIP address capacity is design time configurable:

- The capacity is 2^n bytes, with n in the range [16, 32]. This allows for a minimum capacity of 64K bytes and a maximum capacity of 4 Gbytes.

The cache SRAM memory is used to cache read data (write data is not cached).

An access outside external device spaces result in an AHB-Lite error.

The location of the external devices in the XIP address space is programmable. Each external device $i = 0, 1, 2, 3$ (up to four external devices are supported) has an associated set of MMIO device registers that specify their location and size in the XIP address space:

- ADDR bit field in ADDR register specifies the device location within the XIP address space. The device location should be a multiple of the device capacity.
- MASK bit field in ADDR register specifies the device capacity. The device capacity is 2^m bytes, with m in the range [8, n] (with n specifying the XIP address capacity).

For example, for a 16-Mbyte XIP address space from 0x6000:0000 to 0x60FF:FFFF and a 64-Kbyte device at 0x6001:0000 to 0x6001:FFF, the MMIO device registers are programmed as follows:

- ADDR[23:8] bits in ADDR register is set to 0x0100 (location in XIP address space)
- MASK[23:8] bits in MASK register is set to 0xFF00 (device capacity)

For dual-quad SPI mode, it is required to program the same MMIO device register values for the two external devices that are connected in parallel to the SMIF I/O signal interface.

Write support to external devices is programmable. This is to support nonvolatile devices that do not support write accesses directly, but require a dedicated programming operation:

- WR_EN bit is set to '0' in CTL register: write accesses are not supported. An XIP write transfer results in an AHB-Lite bus error. Typically used for nonvolatile devices without write support.
- WR_EN bit is set to '1' in CTL register: write accesses are supported. Typically used for SRAMs.

32.6.1.2 MMIO Address Space

AHB-Lite transfers to the MMIO address space access the MMIO registers. The MMIO registers include registers to access the FIFOs.

Whereas the XIP address space supports highly efficient read and write access to external devices (through on-the-fly translation of AHB-Lite transfers into SPI transfers), the MMIO address space provides flexibility in the construction of SPI transfers.

32.6.2 TX and RX FIFOs

The SMIF has two TX FIFOs and one RX FIFO. These FIFOs provide an asynchronous clock domain transfer between CLK_mem logic and CLK_if_tx/CLK_if_rx memory interface logic. The memory interface logic is completely controlled through the TX and RX FIFOs. Additionally, SMIF has an RX data MMIO FIFO, which is used only in MMIO mode and which is logically an extension of the RX data FIFO enabling an easy-to-use RX data handling in software.

- The TX command FIFO transmits memory commands to the memory interface logic.
- The TX data FIFO transmits write data to the memory interface transmit logic.
- The RX data FIFO receives read data from the memory interface receive logic.

32.6.2.1 TX Command FIFO

The FIFO consists of eight 27-bit entries. Each entry holds a command. A memory transfer consists of a series of commands. In other words, a command specifies the phase of a memory transfer. Five different types of commands are supported:

- TX command. A memory transfer must start with a TX command. The TX command includes one or two bytes to be transmitted over the memory interface. The TX command specifies the width of the data transfer (single, dual, quad, or octal data transfer). The TX command specifies the data transfer mode (SDR or DDR). The TX command specifies if the command is for the last phase of the memory transfer (explicit "last command" indication). The TX command specifies which of the four external devices are selected (multiple devices can be selected simultaneously); that is, the device selection as encoded by the TX command is used for the complete memory transfer.

The number of bytes included in the TX command depends on the data transfer width and the data transfer mode (SDR or DDR). Two bytes per TX command are transmitted for 8-bit width and DDR (HyperBus/Octal SPI with DDR), when 2 bytes per cycle are transmitted by the memory interface. This way a throughput bottleneck at the TX command FIFO is avoided. In other cases, only one byte per TX command is transmitted to allow byte granularity.

- **TX_COUNT** command. The TX_COUNT command relies on the TX data FIFO to provide the bytes to be transmitted over the memory interface. The TX_COUNT command specifies the number of memory data units to be transmitted. For SPI (except octal SPI with DDR) one memory data unit is a byte; for octal SPI with DDR and HyperBus one memory data unit is a 2-byte word. The TX_COUNT command specifies the width of the data transfer. The TX_COUNT command specifies the data transfer mode (SDR or DDR). The TX command specifies if the command is for the last phase of the memory transfer (explicit “last command” indication).
- **RX_COUNT** command. The RX_COUNT command relies on the RX data FIFO to accept the bytes that are received over the memory interface. The RX_COUNT command specifies the number of memory data units to be received. For SPI (except octal SPI with DDR) one memory data unit is a byte; for octal SPI with DDR and HyperBus one memory data unit is a 2-byte word. The RX_COUNT command specifies the width of the data transfer. The RX_COUNT command specifies the data transfer mode (SDR or DDR). The TX command specifies if the command is for the last phase of the memory transfer (explicit “last command” indication).
- **DUMMY_COUNT** command. The DUMMY_COUNT command specifies a number of dummy cycles. Dummy cycles are used to implement the turn-around (TAR) time in which the memory master changes from a transmitter driving the data lines to a receiver receiving on the same data lines. The DUMMY_COUNT command specifies the number of dummy cycles. The DUMMY_COUNT command specifies if the variable latency mode for HyperRAM (indicated by an active RWDS input) is enabled causing the double number of dummy cycles. The DUMMY_COUNT command specifies whether the RWDS output signal should be driven starting in the last dummy cycle until deselection. The DUMMY_COUNT command never constitutes the last phase of the memory transfer (implicit not “last command” indication); that is, it needs to be followed by another command.
- **DESELECT** command. The Deselect command causes the memory interface transmit logic to finish a transfer and deselect the memory device. The Deselect command always constitutes the last phase of the memory transfer (implicit “last command” indication).

Together, the five command types can be used to construct any SPI or HyperBus (HB) transfer.

The TX command FIFO is used by both the memory interface transmit and receive logic. This ensures lockstep operation.

The software can read the number of used TX command FIFO entries through the USED[2:0] bit field in the MMIO TX_CMD_FF_STATUS register.

The software can write to the TX command FIFO through the MMIO TX_CMD_FIFO_WR register. If software attempts to write an entry of a full TX command FIFO, the BLOCK bit field in the MMIO CTL register specifies the behavior:

- If BLOCK is 0, an AHB-Lite bus error is generated.
- If BLOCK is 1, the AHB-Lite write transfer is extended until an entry is available.

32.6.2.2 TX Data FIFO

The FIFO consists of eight 18-bit entries. Each entry holds a memory data unit (one or two bytes) to be transmitted over the memory interface. Additionally, each entry can hold a byte mask (which is used to mask bytes in HyperBus write transactions). A FIFO TX_COUNT command specifies the number of data units to be transmitted (bytes or 2-byte words); that is, the number of TX data FIFO entries to be used.

The TX data FIFO is used by the memory interface transmit logic.

The software can read the number of used TX data FIFO entries through the USED[3:0] bit field in the MMIO TX_DATA_FF_STATUS register.

In MMIO mode, the software can write to the TX data FIFO through the MMIO TX_DATA_FIFO_WR1, TX_DATA_FIFO_WR1ODD, TX_DATA_FIFO_WR2, and TX_DATA_FIFO_WR4 registers:

- The TX_DATA_FIFO_WR1 register supports write of a single byte to the FIFO.
 - For data transfer width smaller than eight or SDR mode a single byte is written to one FIFO entry.
 - For data transfer width of eight and DDR mode a single byte is written to the low byte of the FIFO entry; the high byte is masked.
- The TX_DATA_FIFO_WR1ODD register supports write of a single byte to the FIFO. This register provides the functionality to write a single byte for a 16-bit word-based memory interface (HyperBus).
 - For data transfer width smaller than eight or SDR mode this register is not used.
 - For data transfer width of eight and DDR mode a single byte is written to the high byte of the FIFO entry; the low byte is masked.
- The TX_DATA_FIFO_WR2 register supports a write of two bytes to the FIFO.
 - For data transfer width smaller than eight or SDR mode a single byte is written to each of two FIFO entries.
 - For data transfer width of eight and DDR mode two bytes are written to one FIFO entry.
- The TX_DATA_FIFO_WR4 register supports a write of four bytes to the FIFO.

- ❑ For data transfer width smaller than eight or SDR mode a single byte is written to each of four FIFO entries.
- ❑ For data transfer width of eight and DDR mode two bytes are written to each of two FIFO entries.

The MMIO interface logic gets the information of the transfer width and transfer rate mode (SDR/DDR) from the previous TX COUNT or RX COUNT command written to the TX command FIFO. Therefore, the related TX COUNT command must be written before writing the transmit data to the TX DATA FIFO.

If software attempts to write more bytes than available entries in the TX data FIFO, the MMIO CTL.BLOCK field specifies the behavior:

- If BLOCK is 0, an AHB-Lite bus error is generated.
- If BLOCK is 1, the AHB-Lite write transfer is extended until the required entries are available.

32.6.2.3 RX Data FIFO

The FIFO consists of sixteen 16-bit entries. Each entry holds up to 2 bytes, which are received over the memory interface. For SDR capturing, only the lower byte is used; for DDR capturing both bytes are used. For SDR capturing with an interface width of less than eight (single, dual, or quad SPI) only 1, 2, or 4 LSBs of the lower byte are used. For DDR capturing with an interface width of less than eight (single, dual, or quad SPI) only 1, 2, or 4 LSBs of both bytes are used.

A FIFO RX_COUNT command specifies the number of data units to be received (bytes or 2-byte words); that is, the number of RX data FIFO entries to be used.

The RX data FIFO is used by both the memory interface transmit and receive logic. This may appear unusual because the memory interface transmit logic does not receive bytes. However, the memory interface transmit logic is responsible for generating the memory interface clock SPIHB_CLK_out, and this clock should not be generated (should be kept low/0) when the RX data FIFO is full. Therefore, although the memory transmit logic does not receive bytes, it should keep track of the RX data FIFO state.

Software can read the number of used RX data FIFO entries through the USED[3:0] bit field in the MMIO RX_DATA_FF_STATUS register.

In MMIO mode, software can read from RX data FIFO through the MMIO RX_DATA_FIFO_RD1, RX_DATA_FIFO_RD2, and RX_DATA_FIFO_RD4 registers.

- The RX_DATA_FIFO_RD1 register supports a read of a single byte from the FIFO.
 - ❑ For data transfer width smaller than eight or SDR mode a single byte is read from one FIFO entry.

- ❑ For data transfer width of eight and DDR mode a single byte is read from the low byte of the FIFO entry, the high byte is discarded.

- The RX_DATA_FIFO_RD2 register supports a read of two bytes from the FIFO.

- ❑ For data transfer width smaller than eight or SDR mode a single byte is read from each of two FIFO entries.
- ❑ For data transfer width of eight and DDR mode two bytes are read from one FIFO entry.

- The RX_DATA_FIFO_RD4 register supports a read of four bytes from the FIFO.

- ❑ For data transfer width smaller than eight or SDR mode a single byte is read from each of four FIFO entries.
- ❑ For data transfer width of eight and DDR mode two bytes are read from each of two FIFO entries.

The MMIO interface logic gets the information of the transfer width and transfer rate mode (SDR/DDR) from the previous TX COUNT or RX COUNT command written to the TX command FIFO. Therefore, the related RX COUNT command must be written before reading the receive data from the RX DATA FIFO.

If the software attempts to read more bytes than available in the RX data FIFO, the BLOCK bit field in the MMIO CTL register specifies the behavior:

- If BLOCK is 0, an AHB-Lite bus error is generated.
- If BLOCK is 1, the AHB-Lite read transfer is extended until the bytes are available.

32.6.3 Interrupts and Triggers

The SMIF has a single interrupt line. The following interrupt causes are associated with this interrupt:

- TR_TX_REQ triggers when tr_tx_req is activated. This interrupt is generated only in MMIO mode.
- TR_RX_REQ triggers when tr_rx_req is activated. This interrupt is generated only in MMIO mode.
- XIP_ADDR_ERROR is activated in XIP mode, if the selected device's DIV2 bit is set to '1' in the ADDR_CTL register and the AHB-Lite bus transfer address is not a multiple of 2. This interrupt is generated only in XIP mode.
- TX_CMD_FIFO_OVERFLOW is activated in MMIO mode on an AHB-Lite write transfer to the TX command FIFO (TX_CMD_FIFO_WR) with not enough free entries available.
- TX_DATA_FIFO_OVERFLOW is activated in MMIO mode on an AHB-Lite write transfer to the TX data FIFO (TX_DATA_FIFO_WR1, TX_DATA_FIFO_WR2, TX_DATA_FIFO_WR4) with not enough free entries available.

- RX_DATA_FIFO_UNDERFLOW is activated in MMIO mode on an AHB-Lite read transfer from the RX data FIFO (RX_DATA_FIFO_RD1, RX_DATA_FIFO_RD2, RX_DATA_FIFO_RD4) with not enough entries available.

The SMIF has two triggers that allow for DW/DMA controller functionality:

- tr_tx_req is activated when the TX data FIFO has 1, 2, 4, or 8 free/available entries. This is a state trigger, and will be (automatically) deactivated when data elements are written into the FIFO.
- tr_rx_req is activated when the RX data FIFO has 1, 2, 4, or 8 used/available entries. This is a state trigger, and will be (automatically) deactivated when data elements are read from the FIFO.

The triggers are activated only in MMIO mode of operation.

32.6.4 Cache

To improve XIP performance, the XIP AHB-Lite interface has a cache. The cache is defined as follows:

- 4-KB capacity.
- Read-only cache. Write transfers bypass the cache.
- Four-way set associative with an LRU replacement scheme. This cache design provides a better hit rate than a direct mapped cache design at the same cache capacity.
- Sequential cache design. The cache tag functionality is performed before the cache data access. A sequential cache design has lower power consumption than a parallel cache design.

- 256 B line/sector, with sixteen 16 B subsectors each. For a 4-KB capacity, this results in a total of 16 lines, distributed over four sets. The subsector design allows for low overhead tag information, as the 16 subsectors in a line/sector share the tag and only have dedicated valid bits.

For each read transfer, the cache tag structure is evaluated before the cache data structure is accessed. The subsector design results in a relatively low number of 16 lines. The 16 associated tags are implemented in flip-flops. The cache data structure is implemented using one 128-bit wide SRAM memory.

Each cache set has an associated 6-bit LRU field, which keeps track of the access history (from the least recently used to most recently used) of the lines in the set.


Each cache line has an associated cache tag and 16 valid bits: one valid bit for each subsector in the cache line. The cache tag identifies the location of the line in system memory.

- The address bits that identify a byte in a cache line are not part of the cache tag (byte address bits 7 down to 0).
- The address bits that identify a cache set are not part of the cache tag (byte address bits 9 and 8).
- The address bits that are not part of the SMIF XIP address region are not part of the tag.

The above omissions of address bits result in small tags. As a result, the cache tag structure can be evaluated quickly.

Figure 32-24 gives an overview of the cache design.

Figure 32-24. Cache Organization



Read transfers that “hit” are processed by the cache.

Read transfers that “miss” result in a XIP SPI read transfer. A miss results in a 16 B (subsector) refill. The cache data structure is updated with the 16 B of refilled data. Two cases are considered:

- The refilled data is a subsector of a resident cache line. This case refills the data to the cache way that is used by the resident cache line. The subsector's valid field is set to '1' (the valid fields of all other subsectors in the cache line remain unchanged).
- The refilled data is not a subsector of a resident cache line. This case refills the data to the cache way that is identified by the LRU scheme. The cache line address bits are updated, and the subsector's valid field is set to '1' (the valid fields of all other subsectors in the cache line are set to '0'). Note that this case replaces a resident cache line.

The cache has a LRU replacement scheme. Each cache set has an associated 6-bit LRU field:

- LRU[5]: '1' when way 0 is less recently used than way 1, '0' otherwise.
- LRU[4]: '1' when way 0 is less recently used than way 2, '0' otherwise.
- LRU[3]: '1' when way 0 is less recently used than way 3, '0' otherwise.
- LRU[2]: '1' when way 1 is less recently used than way 2, '0' otherwise.
- LRU[1]: '1' when way 1 is less recently used than way 3, '0' otherwise.
- LRU[0]: '1' when way 2 is less recently used than way 3, '0' otherwise.

Although six bits allow for $2^6 = 64$ bit patterns, only $4 \times 3 \times 2 \times 1 = 24$ bit patterns are legal LRU representations. The LRU set information is reset to all '1' or 0b111111, representing a set in which way 0 is less recently used than way 1, which is less recently used than way 2, which is less recently used than way 3. In this case, the line in way 0 is replaced when a new line is brought into the set. A line is made the most recently line of its set, when it is brought into the set, or when its line data is used because of an AHB-Lite read transfer.

A write to an address in the read-only cache, invalidates the associated cache subsector, but does not affect the LRU.

If ENABLED bit is set to '1' in the CA_CTL register, the cache is enabled and if ENABLED is set to '0', the cache is disabled.

If PREF_EN bit is set to '1' in the CA_CTL register, prefetching is enabled and if PREF_EN bit is set to '0', prefetching is disabled. If prefetch is enabled, a cache miss results in a 16 B (subsector) refill for the missing data and a 16 B prefetch for the next sequential data (independent of whether this data is already in the cache or not). The data of

the 16 B prefetch is stored in a temporary buffer and only copied into the cache when a future read transfer “misses” in the cache and requires the buffered data.

For debug purposes, the tag and 16 valid bits of a cache line are readable through MMIO registers. The LRU information of a cache set is readable through MMIO registers.

32.6.5 Arbitration

The SMIF provides two AHB-Lite slave interfaces to CPUSS (one fast interface and one slow interface) or one AHB-Lite slave interface (slow interface) and one AXI slave interface. These AHB or AXI interfaces can generate XIP requests to the external memory devices.

32.6.6 Cryptography

In XIP mode, a cryptography component supports on-the-fly encryption for write data and on-the-fly decryption for read data. The use of on-the-fly cryptography is determined by a device's CRYPTO_EN bit field in the MMIO CTL register. In MMIO mode, the cryptography component is accessible through a MMIO register interface to support off-line encryption and decryption.

The rationale for this component is as follows: data is encrypted in the external memory devices and is not encrypted in the device. Therefore, SPI read and write data transfers require decryption and encryption functionality respectively. By storing data encrypted in the external memory devices (possibly nonvolatile devices), leakage of sensitive data is addressed.


Encryption and decryption are based on the AES-128 forward block cipher: advanced encryption standard block cipher with a 128-bit key. KEY[127:0] is a secret key, which is programmed into the MMIO CRYPTO_KEY3 to CRYPTO_KEY0 registers. These MMIO registers are software write-only: a software read returns “0”. By applying AES-128 with KEY[127:0] on a plaintext PT[127:0], we get ciphertext CT[127:0].

In XIP mode, the XIP address is used as the plaintext PT[]. The resulting ciphertext CT[] is used on-the-fly and not software accessible. The XIP address is extended with the MMIO CRYPTO_INPUT3 to CRYPTO_INPUT0 registers.

In MMIO mode, the MMIO CRYPTO_INPUT3 to CRYPTO_INPUT0 registers provide the plaintext PT[]. The resulting ciphertext CT[] is provided through the MMIO CRYPTO_OUTPUT3 to CRYPTO_OUTPUT0 registers.


Figure 32-25 illustrates the functionality in XIP and MMIO modes.

Figure 32-25. Cryptography in XIP and MMIO Mode



In XIP mode, the resulting ciphertext CT[] is XOR'd with the SPI transfer's read data or write data. Note that the AES-128 block cipher is on the address of the data and not on the data itself. For SPI read transfers, this means that as long as the latency of the SPI transfer's read data is longer than the AES-128 block cipher latency, the on-the-fly decryption does not add any delay. Figure 32-26 illustrates the complete XIP mode functionality.

Figure 32-26. Cryptography in XIP Mode



The SMIF supports on-the-fly cryptography. As a result, the external memory content is encrypted.

The encryption/decryption uses the AES block cipher with a 128-bit key in counter mode (CTR mode).

The CTR mode requires a nonce and a counter. The nonce is provided by a programmable SMIF MMIO register, and the system interconnect bus address is used as the counter (the lower 4 bits of the bus address are not used). The nonce and counter values are concatenated to provide the input to the block cipher.

The on-the-fly cryptography provides confidentiality for constant/read-only data in the external SPI memory devices. However, confidentiality is less for non-constant/write data. This is explained as follows. Consider an address A for which the block cipher output is AES(A). First, we store the

plain data p_0 as cipher data $c_0 = \text{AES}(A) \wedge p_0$. Next, we store the plain data p_1 as cipher data $c_1 = \text{AES}(A) \wedge p_1$.

Both c_0 and c_1 can be observed on the device interface. We know that $c_0 \wedge c_1 = \text{AES}(A) \wedge p_0 \wedge \text{AES}(A) \wedge p_1 = p_0 \wedge p_1$. If there are statistical information on plain data p_i samples (for example, the plain data "0" is frequently written), we can deduce p_i , if we have enough cipher data c_i samples available.

If the SMIF on-the-fly cryptography is used for write data transfers, the nonce should be changed between write transfers to the same address to ensure confidentiality.

For dynamic data storage in the SMIF area with fixed CRYPTO_INPUT and KEY, SMIF block encryption is vulnerable to known-plaintext attacks (KPA).

Therefore, content in the SMIF memory area:

- for code storage, read-only and constant, the current scheme is good
- for data storage, the user can change CRYPTO_INPUT and KEY if necessary
- For sensitive data, the user can use mxcrypto for real AES encryption/decryption

32.6.6.1 Mapping of Bytes

The plain text address, cipher text address, plain text data, and cipher data are represented as 128-bit values: PA[127:0], CA[127:0], PD[127:0], CD[127:0]. Each 128-bit value consists of sixteen bytes. The mapping between bytes and the 128-bit vector value V[127:0] is as follows:

- Vector byte 0: V[7:0]
- Vector byte 1: V[15:8]
- ...
- Vector byte 15: V[127:120]

Given a Traveo II address SoC_A[31:0], we derive an external device address A[m-1:0] = {SoC[m-1:4], 4'b0000}. The mapping between data bytes is as follows:

- Plain text byte 0 (PD[7:0]) is located at Traveo II address {SoC_A[31:4], 4'b0000} and cipher text byte 0 (CD[7:0]) is located at external device address {A[m-1:4], 4'b0000}.
- Plain text byte 1 (PD[15:8]) is located at Traveo II address {SoC_A[31:4], 4'b0001} and cipher text byte 1 (CD[15:8]) is located at external device address {A[m-1:4], 4'b0001}.
- ...

Plain text byte 15 (PD[127:120]) is located at Traveo II address {SoC_A[31:4], 4'b1111} and cipher text byte 15 (CD[127:120]) is located at external device address {A[m-1:4], 4'b1111}.

32.6.6.2 Software and MMIO Register Requirements

To ensure maximum protection of the XIP encryption key KEY[127:0], the following MMIO requirements should be met:

- The trusted software writes the encryption key in CRYPTO_KEY0 to CRYPTO_KEY3.
- The hardware provides “write only” access to the encryption key in CRYPTO_KEY0 to CRYPTO_KEY3. Software always reads CRYPTO_KEY0 to CRYPTO_KEY3 as “0” (in both XIP_MODE and MMIO_MODE).
- Software can read CRYPTO_RESULT0, to CRYPTO_RESULT3 in MMIO_MODE. It reads CRYPTO_RESULT0 to CRYPTO_RESULT3 as “0” in XIP_MODE.

To ensure that DAP test accesses cannot access decrypted data from the XIP address space, the following protection mode requirements should be met (these are the same requirements that apply for the SRAM memories in the CPUSS).

Table 32-11. XIP Address Space Protection

Mode	Usage	XIP Address Space
VIRGIN	CPU, DW/DMA	Yes
	DAP	Yes (no write in privileged execution mode)
OPEN	CPU, DW/DMA	Yes
	DAP	Yes (no write in privileged execution mode)
PROTECTED	CPU, DW/DMA	Yes
	DAP	No
KILL	CPU, DW/DMA	Yes
	DAP	No
BOOT	CPU, DW/DMA	Yes
	DAP	No

Note: The XIP address space is a user memory (no privileged memory is present).

32.6.6.3 Cryptography Support for AXI Interface

The AXI interface in XIP mode can generate the next read or write transfer request while the previous read or write transfer is currently processed by the memory interface logic. The exact time of this next transfer request is

- whenever the next transfer is available at the AXI interface and
- the XIP block is ready to consume the next transfer request; that is, does not need the information (especially the address) of the previous transfer anymore.

To support data encryption/decryption for such outstanding transactions, the SMIF crypto block contains:

- An input FIFO (with a depth of one entry), which acts as a buffer to keep the address (used as part of the plain text) stable for an encryption.
- An output FIFO (with a depth of two entries), which stores the encrypted address of the first transfer while the encrypted address for the next transfer is already calculated.

For read transfers, this allows the crypto block to calculate the encrypted address of the next transfer while the previous memory transfer is currently in the data phase or even while the memory latency cycles are generated. This way the memory latency time can be used to calculate not only the current but also the next address encryption. This ensures that no delay is added by on-the-fly decryption for CM7 cache line fills (2 x 16-byte read transfers), even when using a fast HyperBus memory device with merged transfers (only 1 time the address and latency cycles).

32.6.7 Serial Memory Interface Logic

32.6.7.1 TX and RX Logic

The memory interface logic is implemented as two independent pieces:

- **The memory interface transmit logic.** This logic operates on the interface transmitter clocks `clk_if_tx_div`,

`clk_if_tx_data_out`, and `clk_if_tx_int`. This logic is responsible for:


- Generating and driving the `sphb_clk_out` clock.
- Driving the low active select signals, `sphb_select_out[3:0]`.
- Driving the outgoing data on `sphb_data_out[7:0]`.

- **The memory interface receiver logic.** This logic operates on the interface receiver clocks `clk_if_rx_capture_ff`, `clk_if_rx_capture`, `clk_if_rx_capture_del[]`, and `clk_if_rx_int`. This logic captures incoming data on `sphb_data_in[7:0]` based on the following capture schemes:

- (Legacy) Output/feedback clock-based data capture.
- RWDS-based data capture.

The different capture schemes are described in [32.6.7.3 Data Capture](#). [Figure 32-27](#) shows an overview of the interface logic; the details are discussed in the following sections.

Figure 32-27. TX/RX Interface to External Memory Devices



The TX command FIFO is used by both the memory interface transmit and receive logic. This ensures lockstep operation. Both pieces of logic are implemented as state machines that are driven by the TX command FIFO entries.

Note that the memory interface transmit logic clock generation of `spihb_clk_out` affects the memory interface receiver logic clocks. If `spihb_clk_out` is turned off, the memory interface receiver logic clocks will turn off.

The memory interface clock `sphb_clk_out` may turn off during a memory transfer in the following situations:

- The memory transfer produces read data (data from the memory device) and the RX data FIFO is full. To detect this situation, the TX interface transmitter logic needs to have access to the RX data FIFO level.
- The memory transfer consumes write data (data for the memory device) and the TX data FIFO is empty.

32.6.7.2 Flow Control

Flow control prevents overflow and underflow for both directions – TX (SMIF writes to memory) and RX (SMIF reads from memory).

- TX underflow:
 - SMIF stops memory clock until TX cmd/data FIFO has data available
- TX overflow:
 - No flow control mechanism is provided at the memory interface. The memory device prevents TX overflow.
 - Flash devices receive data in WriteBuffer at bus speed or use word write. This excludes overflow.
 - RAM writes the RAM array at the speed of external bus, which excludes overflow.
- RX overflow:
 - SMIF stops memory clock.
- RX underflow:
 - HyperBus/OSPI memories stop RWDS/DQS. The TX interface logic needs to know how many memory interface clock cycles are required to generate without listening to RWDS/DQS. A synchronization from RX to TX interface clock domain causes a latency, which leads to overclocking (providing more memory clock cycles than needed). This leads to a potential mismatch between the number of cycles and therefore between the number of bytes/words read from memory and the number of bytes/words used in the SMIF. This creates an issue for any read side effects in the memory such as Bus CRC generation (potentially added to SMIF in the future).

Therefore, the SMIF generates the number of latency cycles for initial and page boundary crossing based on MMIO registers (reflecting the memory requirements depending on memory interface clock frequency).

The only exception is the variable initial latency for HyperRAM based on the RWDS refresh indicator. This signal can be captured safely with the TX interface clock at the end of the address phase. This is because at that time this signal is stable for at least three memory interface clock cycles (per HyperBus protocol).

- All other SPI memories have no flow control.

32.6.7.3 Data Capture

SMIF supports data capturing with the SMIF output or output feedback clock for SDR and DDR timing:

- (Legacy) output/feedback clock-based data capture

Capture scheme supported by the SMIF:

- Internal clock-based data capture
- RWDS-based data capture

Table 32-12 shows the (DLP/RWDS) capture schemes required for serial memory.

Table 32-12. Transfer Types and Capture Schemes

Memory Bus	Feature	Loading	SMC		SMIF Features for M7 Core	
			Support for Capture via	Frequency MHz	Support for Capture via	Frequency MHz
HyperBus	HB	DDR	RWDS	200	RWDS	100t
xSPI, x=2,4,8 8-8-8	OSPI	DDR	RWDS /DLP	200	RWDS/DLP	100
8-8-8		SDR	RWDS /DLP	200	RWDS/DLP	100
4-4-4	2x/Quad	DDR	N/A	N/A	DLP	100
4-4-4		SDR	N/A	N/A	DLP	100
4-4-4	Quad	DDR	DLP	100	DLP	100
4-4-4		SDR	DLP	200	DLP	100
2-2-2	Dual	DDR	DLP	100	DLP	100
2-2-2		SDR	DLP	200	DLP	100
1-8-8	Octal I/O	DDR	N/A	N/A	RWDS/DLP	100
1-8-8		SDR	N/A	N/A	RWDS/DLP	100
1-4-4	2xQuad I/O	DDR	N/A	N/A	DLP	100
1-4-4		SDR	N/A	N/A	DLP	100
1-4-4	Quad I/O	DDR	DLP	100	DLP	100
1-4-4		SDR	DLP	200	DLP	100
1-2-2	Dual I/O	DDR	DLP	100	DLP	100
1-2-2		SDR	DLP	200	DLP	100
1-1-8	Octal data	DDR	RWDS/DLP	200	RWDS/DLP	100
1-1-8		SDR	RWDS/DLP	200	RWDS/DLP	100
1-1-4	2xQuad data	DDR	N/A	N/A	DLP	100
1-1-4		SDR	N/A	N/A	DLP	100
1-1-4	Quad data	DDR	DLP	100	DLP	100
1-1-4		SDR	DLP	200	DLP	100
1-1-2	Dual data	DDR	DLP	100	DLP	100
1-1-2		SDR	DLP	200	DLP	100
SPI	SPI	DDR	DLP	100	DLP	100
1-1-1		SDR	DLP	200	DLP	100

In the following sections, the principles of the three capture schemes and their capture timing are discussed.

Output/Feedback Clock-based Capture

The output/feedback clock-based capture scheme uses the memory output clock (spihb_clk_out), the inverted memory output clock, the memory output feedback clock (spihb_clk_in), or the inverted memory output feedback clock as the capture clock (clock selection CLOCK_IF_RX_SEL[3:0] in the CTL register).

To allow sampling later and to adjust the sample time with a finer granularity, the option of using the delay line for delaying the output or feedback clock has been added to this capture scheme. The delay line is already available for the RWDS capture scheme.


For SDR timing, any of the clock selections above can be used as capture clock (depending on the delay of the RX data). The data is directly captured in the RX data FIFO.

For DDR timing, the inverted version of the selected capture clock (memory output clock, memory output feedback clock, or an inverted and/or delayed version) also needs to be used as capture clock. The data driven in the first memory half-cycle is

captured in the memory interface RX logic with this inverted selected clock intermediately. The data driven at the second memory half-cycle is directly captured by the RX data FIFO (as data low part); this and the first capture data are handed over to the RX data FIFO (as data high part). The RX data FIFO is 16-bit wide.

Figure 32-28 illustrates the output/feedback clock-based capture scheme.

Figure 32-28. Output/Feedback Clock-based Capture Scheme



The output/feedback clock-based capture scheme includes an RX interface FSM, which tracks the memory protocol and generates a 'capture enable' signal for storing the received data in the RX data FIFO. This RX interface FSM runs with the same clock, `clk_if_rx_out_fback`.

RWDS-based Capture

In the RWDS-based capture scheme, the RWDS signal is used as a clock to capture the input data. The RWDS signal edges are aligned with the start of the data valid time in the memory device.

The SMIF contains a clock multiplexer, which selects one of the delay line taps based on `INT_CLOCK_DELAY_TAP_SEL0/1` registers. This allows it to calibrate the delay.

The calibration works as follows:

- The software can use a known data sequence located at an arbitrary address of the memory to "paint" the data eye. A flash image is known by the customer, in case of the RAM the data has to be written before. Even in the case that there is no known data sequence, software can read one safely by configuring a low-memory clock frequency and a late delay line tap (to meet hold timing).
- Now, software reads the data sequence for each delay tap setting using the SMIF MMIO mode. For some adjacent delay line taps it will receive the correct data.

Note that the selectable delay per delay line tap can also be used in the RWDS-based capture scheme. Software can select that the delay per delay line tap in the `CTL_DELAY_LINE_SEL` register. The delay line with the most adjacent delay line taps resulting in correct received data is preferred.

- Then, it configures the middle tap of all the taps providing the correct data.

Figure 32-29. RWDS-based Capture Scheme





Figure 32-29 shows the timing of the RWDS-based capture scheme.

Figure 32-30. Capturing Data with RWDS



Note that the RWDS signal (sphi_b_rwds) is not only used to capture data from the memory device, and in HyperBus mode when writing to and reading from the HyperBus RAM.

32.7 SMIF Registers

Table 32-13. List of SMIF Registers

Register	Name	Description
SMIFx_CTL	Control	Control register
SMIFx_STATUS	Status	Busy status of AHB cache, AXI interface, cryptography, XIP device interface, and other logic in SMIF.
SMIFx_DELAY_TAP_SEL	RX clock delay tap select	Shift the strobe signal into the data eye.
SMIFx_TX_CMD_FIFO_STATUS	Transmitter command FIFO status	Number of entries used in the TX command FIFO.
SMIFx_TX_CMD_FIFO_WR	Transmitter command FIFO write	Command data
SMIFx_TX_DATA_FIFO_CTL	Transmitter data FIFO control	Determines when the TX data FIFO tr_tx_req trigger is activated.
SMIFx_TX_DATA_FIFO_STATUS	Transmitter data FIFO status	Number of entries used in the TX data FIFO.
SMIFx_TX_DATA_FIFO_WR1	Transmitter data FIFO write 1	Write a single byte to the TX FIFO.
SMIFx_TX_DATA_FIFO_WR2	Transmitter data FIFO write 2	Write two simultaneous bytes to the TX FIFO.
SMIFx_TX_DATA_FIFO_WR4	Transmitter data FIFO write 4	Write four simultaneous bytes to the TX FIFO.
SMIFx_TX_DATA_FIFO_WR1ODD	Transmitter data FIFO write	Write a single byte to the TX FIFO for odd byte addresses when using a word-based memory protocol (Hyperbus/Octal SPI).
SMIFx_RX_DATA_MMIO_FIFO_CTL	Receiver data MMIO FIFO control	Determines when RX data FIFO tr_rx_req trigger is activated.
SMIFx_RX_DATA_MMIO_FIFO_STATUS	Receiver data MMIO FIFO status	Number of entries used in the RX data MMIO FIFO.
SMIFx_RX_DATA_FIFO_STATUS	Receiver data FIFO status	Number of entries used in the RX data FIFO.
SMIFx_RX_DATA_MMIO_FIFO_RD1	Receiver data MMIO FIFO read 1	Provides a single byte from the RX data MMIO FIFO.
SMIFx_RX_DATA_MMIO_FIFO_RD2	Receiver data MMIO FIFO read 2	Provides two bytes from the RX data MMIO FIFO.
SMIFx_RX_DATA_MMIO_FIFO_RD4	Receiver data MMIO FIFO read 4	Provides four bytes from the RX data MMIO FIFO.
SMIFx_RX_DATA_MMIO_FIFO_RD1_SILENT	Receiver data MMIO FIFO silent read	This register is similar to RX_DATA_MMIO_FIFO_RD1, with the exception that the entry is not removed from the FIFO.
SMIFx_SLOW_CA_CTL	Slow cache control	The slow interface is used for the CPUSS CM0+, Crypto, and DataWire components.
SMIFx_SLOW_CA_CMD	Slow cache command	Cache and prefetch buffer invalidation.
SMIFx_CRYPT0_CMD	Cryptography command	Starts an encryption operation.
SMIFx_CRYPT0_INPUT0	Cryptography input 0	Four bytes of the plaintext PT[31:0].
SMIFx_CRYPT0_INPUT1	Cryptography input 1	Four bytes of the plaintext PT[63:32].
SMIFx_CRYPT0_INPUT2	Cryptography input 2	Four bytes of the plaintext PT[95:64].
SMIFx_CRYPT0_INPUT3	Cryptography input 3	Four bytes of the plaintext PT[127:96].
SMIFx_CRYPT0_KEY0	Cryptography key 0	Four bytes of the key KEY[31:0].
SMIFx_CRYPT0_KEY1	Cryptography key 1	Four bytes of the key KEY[63:32].
SMIFx_CRYPT0_KEY2	Cryptography key 2	Four bytes of the key KEY[95:64].
SMIFx_CRYPT0_KEY3	Cryptography key 3	Four bytes of the key KEY[127:96].
SMIFx_CRYPT0_OUTPUT0	Cryptography output 0	Four bytes of the ciphertext CT[31:0].
SMIFx_CRYPT0_OUTPUT1	Cryptography output 1	Four bytes of the ciphertext CT[63:32].
SMIFx_CRYPT0_OUTPUT2	Cryptography output 2	Four bytes of the ciphertext CT[95:64].
SMIFx_CRYPT0_OUTPUT3	Cryptography output 3	Four bytes of the ciphertext CT[127:96].
SMIFx_INTR	Interrupt	Indicates the interrupt causes.
SMIFx_INTR_SET	Interrupt set	Write with '1' to set the corresponding bit in interrupt request register.

Table 32-13. List of SMIF Registers

Register	Name	Description
SMIFx_INTR_MASK	Interrupt mask	Mask bit for the corresponding bit in interrupt request register.
SMIFx_INTR_MASKED	Interrupt masked	Reflects a bitwise 'AND' between the interrupt request and mask registers.
SMIFx_DEVICEy_CTL	Device control	Control register
SMIFx_DEVICEy_ADDR	Device region base address	Device region base address
SMIFx_DEVICEy_MASK	Device region mask	Device region mask
SMIFx_DEVICEy_ADDR_CTL	Address control	Address control
SMIFx_DEVICEy_RD_CMD_CTL	Read command control	Read command control
SMIFx_DEVICEy_RD_ADDR_CTL	Read address control	Read address control
SMIFx_DEVICEy_RD_MODE_CTL	Read mode control	Read mode control
SMIFx_DEVICEy_RD_DUMMY_CTL	Read dummy control	Read dummy control
SMIFx_DEVICEy_RD_DATA_CTL	Read data control	Read data control
SMIFx_DEVICEy_RD_BOUND_CTL	Read bound control	Read bound control
SMIFx_DEVICEy_WR_CMD_CTL	Write command control	Write command control
SMIFx_DEVICEy_WR_ADDR_CTL	Write address control	Write address control
SMIFx_DEVICEy_WR_MODE_CTL	Write mode control	Write mode control
SMIFx_DEVICEy_WR_DUMMY_CTL	Write dummy control	Write dummy control
SMIFx_DEVICEy_WR_DATA_CTL	Write data control	Write data control

Note that overwriting the same value on each register has different effects; this is explained in the register map by the software access attributes. For SMIF registers, the following access restrictions must be mentioned:


- All status registers are not SW-writable.
- INTR is SW-clear and HW-set (or set by writing '1' to INTR_SET).
- Read INTR_SET will return the value of INTR.
- Other registers are normal and can be overwritten with the same value.

33. SDHC Host Controller



The secure digital high capacity (SDHC) host controller in Traveo II allows interfacing with embedded multimedia card (eMMC)-based memory devices, secure digital (SD) cards, and secure digital input output (SDIO) cards. [Figure 33-1](#) illustrates a typical application using the SDHC block.

Figure 33-1. Typical SDHC Application



33.1 Features

- Complies with eMMC 5.1, SD 6.0, and SDIO 4.10 standards
- Supports host controller interface (HCI) 4.2 shared by eMMC and SD
- SD interface supports 1-bit and 4-bit bus interfaces, and the following speed modes. The specified data rate is for a 4-bit bus.
 - 3.3-V signal voltage: Default speed (12.5 MB/s at 25 MHz) and high speed (25 MB/s at 50 MHz)
- eMMC interface supports 1-bit, 4-bit, and 8-bit bus interfaces, and the following speed modes. The specified data rate is for an 8-bit bus.
 - Legacy (26 MB/s at 26 MHz), high-speed SDR (52 MB/s at 52 MHz), and high-speed DDR (104 MB/s at 52 MHz)
- Supports three DMA modes – SDMA, ADMA2, and ADMA3 – through a dedicated DMA engine
- Provides 1KB SRAM for buffering up to two 512-byte blocks
- Provides I/O interfaces for functions such as card detection and mechanical write protection


33.1.1 Features Not Supported

The SDHC block does not support the following features.

- SD/SDIO operation in UHS-II mode
- Command queuing engine (CQE)
- eMMC boot operation in dual data rate mode
- Read wait operation by DAT[2] signaling in an SDIO card
- Suspend/resume operation in an SDIO card
- Interrupt input pins for embedded SD systems
- SPI protocol mode of operation
- SD UHS-I mode using 1.8-V signal voltage: SDR, SDR25, SDR50, and DDR50

33.2 Block Diagram

Figure 33-2. SDHC Block Diagram



The SDHC block supports all three interfaces – SD, SDIO, and eMMC. It does not have the eMMC reset signal, which is not a mandatory signal for eMMC operation. The AHB master interface helps to transfer data to and from the system memory and the AHB slave interface provides access to the configuration registers. The register set comprises the standard SD host controller interface (HCI) registers as specified in the SD Specifications Part A2 SD Host Controller Standard Specification. These registers are described in the *Traveo II Body Controller High Registers TRM*. The DMA engine handles direct data transfer between the SDHC logic and system memory. It supports SDMA, ADMA2, and ADMA3 modes based on the configuration.

The SDHC block complies with the following standards. Refer to the specifications documents for more information on the protocol and operations.

- SD Specifications Part 1 Physical Layer Specification Version 6.00
- SD Specifications Part A2 SD Host Controller Standard Specification Version 4.20
- SD Specifications Part E1 SDIO Specifications Version 4.10
- Embedded Multi-Media Card (eMMC) Electrical Standard 5.1

33.3 Clocking

Table 33-1 lists the different clocks used in the SDHC block. While configuring the clock for SDHC make sure that $CLK\_SLOW \geq CLK\_GRx \geq CLK\_CARD$.

Table 33-1. Clocks in SDHC

Source	SDHC Clock	Function
CLK_SLOW	Core SDHC Clock	Used for core SDHC functions including the packet buffer SRAM; it is sourced from the slow clock (CLK_SLOW); it must be $\geq$ AHB slave clock.
	AHB Master Interface Clock	Used by the AHB master interface; it is sourced from the slow clock (CLK_SLOW); it must be $\geq$ AHB slave clock.
CLK_GRx	AHB Slave Interface Clock	Used by the AHB slave interface; it is clocked by the PERI group clock (CLK_GRx); it must be $\geq$ CLK_CARD. The group clock is derived from the PERI clock (CLK_PERI) using a divider. Because this divider can remain at the default value of '1' for most applications, CLK_PERI can be considered as CLK_GRx for SDHC. See Clocking System chapter on page 197 for information on CLK_GRx and CLK_PERI.
CLK_HFx	Base Clock/Card Clock	Used for sourcing the SD/eMMC interface clock (CLK_CARD); it is derived from CLK_HFx; it must be set to 100 MHz to be compatible with the Capabilities register. See 33.3.2 Base Clock (CLK_HFx) Configuration for details. See High-Frequency Root Clocks on page 206 to know which CLK_HFx drives an SDHC instance.
	Timer Clock	Used for command and data timeout functions; it is derived from CLK_HFx.

33.3.1 Clock Gating

All the clocks except the slave interface clock can be gated internally to enter standby mode (See [Power Modes on page 629](#)). In standby mode, you can also stop the clocks externally if required. The slave clock cannot be gated because it is used for wakeup logic (see [Interrupts to CPU on page 629](#)) during the standby mode. The card clock is gated by clearing the SDHCx_CORE_CLK_CTRL_R.SD_CLK_EN bit and other clocks are gated by clearing the SDHCx_CORE_CLK_CTRL_R.INTERNAL_CLK_EN bit. See [Clock Setup on page 635](#) for the sequence to be followed while modifying this bit.

33.3.2 Base Clock (CLK_HFx) Configuration

The SDHCx_CORE_CAPABILITIES1_R register has a read-only field (BASE_CLK_FREQ) to indicate the base clock frequency so that an SD HCI-compatible driver can easily configure the divider for the required bus speed. This value is set to 0x64 (100 MHz) and hence CLK_HFx must be set to 100 MHz. If this compatibility is not required, CLK_HFx can be set to any value. See [33.3.4 Timeout \(TOUT\) Configuration](#).

33.3.3 Card Clock (SDCLK) Configuration

The SDCLK or card clock frequency is set by configuring the 10-bit divider in SDHCx_CORE_CLK_CTRL_R and selecting the 10-bit divided clock mode by clearing the SDHCx_CORE_CLK_CTRL_R.CLK_GEN_SELECT bit. The default value of this bit is zero. The SDHCx_CORE_CLK_CTRL_R.UPPER_FREQ_SEL field holds the upper two bits (9:8) and the SDHCx_CORE_CLK_CTRL_R.FREQ_SEL field holds the lower eight bits (7:0) of the divider. Base clock frequency is sourced from CLK_HFx as explained in [Table 33-1](#). SDCLK frequency is equal to base clock frequency when the divider value is zero.

$$SDCLK \text{ Frequency} = \text{Base Clock Frequency} / (2 \times 10\text{-bit divider value})$$

These fields are set automatically, based on the selected Bus Speed mode, to a value specified in one of the preset registers when SDHCx_CORE_HOST_CTRL2_R.PRESET_VAL_ENABLE is set. The preset registers are selected according to [Table 33-2](#).

33.3.4 Timeout (TOUT) Configuration

An internal timer is used for command and data timeouts. The timeout value is specified through the SDHCx_CORE_TOUT_CTRL_R.TOUT_CNT register field. The timer clock (TMCLK) frequency indicated by the SDHCx_CORE_CAPABILITIES1_R.TOUT_CLK_FREQ and SDHCx_CORE_CAPABILITIES1_R.TOUT_CLK_UNIT read-only fields is '1' MHz. Timer clock is derived by dividing the CLK_HFx, which means that CLK_HFx must be set to 100 MHz to be compatible with the Capabilities register.

33.4 Bus Speed Modes

The SDHC block can operate in either SD/SDIO mode or in eMMC mode. The SDHC block operates in eMMC mode when the SDHCx_CORE_EMMC_CTRL_R.CARD_IS_EMMC bit is set; otherwise, it operates in SD/SDIO mode. The speed mode is configured through SDHCx_CORE_HOST_CTRL1_R and SDHCx_CORE_HOST_CTRL2_R registers as shown in Table 33-2. The SDHCx_CORE_HOST_CTRL2_R.UHS2_IF_ENABLE bit should remain at its default value of zero because the block does not support UHS-II mode. The card clock must be configured according to the selected speed mode through the SDHCx_CORE_CLK_CTRL1/2_R register. See 33.3.3 Card Clock (SDCLK) Configuration for more information.

Table 33-2. Bus Speed Mode Configuration

Bus Speed Mode	SDHCx_CORE_HOST_CTRL1_R Field	SDHCx_CORE_HOST_CTRL2_R Fields		Selected Preset Register
	HIGH_SPEED_EN	SIGNALING_EN	UHS_MODE_SEL	
SD Default Speed (DS)	0	0	Don't care	PRESET_DS_R
SD High Speed (HS)	1	0	Don't care	PRESET_HS_R
eMMC Legacy	Don't care	0	000b	PRESET_SDR12_R
eMMC High Speed	Don't care	0	001b	PRESET_SDR25_R
eMMC High Speed DDR	Don't care	0	100b	PRESET_DDR50_R

33.5 Power Modes

The block can operate during active and sleep system power modes. It does not support deep sleep mode and cannot wake up from events such as card insertion and removal when the system is in deep sleep. Only the SDHCx_WRAP_CTL.ENABLE register is retained when the system enters deep sleep mode and the SRAM is switched off to save power. Make sure that no AHB traffic (such as register read/write and DMA operation) is present, the SD/SDIO/eMMC bus interface is idle, and no data packets are pending in the packet buffer SRAM when the system transitions into deep sleep mode.

33.5.1 Standby Mode

The block can be put into standby mode to save power during the active and sleep system power modes by turning off the clocks. See Clock Gating on page 628 for details. The block can detect wakeup interrupts (see Interrupts to CPU on page 629) in standby mode.

33.6 Interrupts to CPU

The block provides two interrupt signals to CPUSS:

- Wakeup Interrupt Signal – Triggered on events such as card insertion, removal, and SDIO card interrupt. This interrupt source cannot wake up the system from deep sleep mode and is provided so that a host driver can take appropriate action on those events. For example, resuming operation from standby mode on card insertion. See 33.5.1 Standby Mode for details. As card insertion and removal is not applicable to an embedded device, wakeup interrupt should not be used in this case. However it can still be used for SDIO card interrupt.

- General Interrupt Signal – Triggered on all other events, in either normal conditions or error conditions.

A host driver must not enable the wakeup and general interrupt signals at the same time.

To use only the wakeup interrupt signal, clear the SDHCx_CORE_NORMAL_INT_STAT_R and SDHCx_CORE_NORMAL_INT_SIGNAL_EN_R registers, and then set the enable bits of the required wakeup events in the SDHCx_CORE_WUP_CTRL_R and SDHCx_CORE_NORMAL_INT_STAT_EN registers.

To use only the general interrupt signal, clear the SDHCx_CORE_WUP_CTRL_R and SDHCx_CORE_NORMAL_INT_STAT_R registers. Then, set the required bits in SDHCx_CORE_NORMAL_INT_SIGNAL_EN_R and SDHCx_CORE_NORMAL_INT_STAT_EN registers.

These interrupts remain asserted until the CPU clears the interrupt status through one of the status registers – SDHCx_CORE_NORMAL_INT_STAT_R and SDHCx_CORE_ERROR_INT_STAT_R.

The SDIO card interrupt status bit, SDHCx_CORE_NORMAL_INT_STAT_R.CARD_INTERRUPT, is a read-only bit. The host driver may clear the SDHCx_CORE_NORMAL_INT_STAT_EN_R.CARD_INTERRUPT_STAT_EN bit before servicing the SDIO card interrupt and may set this bit again after all interrupt requests from the card are cleared to prevent inadvertent interrupts.

Following is the list of registers used in interrupt configuration.

Table 33-3. Interrupt Control Registers

Register	Description
SDHCx_CORE_WUP_CTRL_R	Enables or disables different wakeup interrupts. Host driver must maintain voltage on the SD bus by setting SDHCx_CORE_PWR_CTRL_R.SD_BUS_PWR_VDD1 bit for these interrupts to occur. These interrupts cannot wakeup the device from deep sleep.
SDHCx_CORE_NORMAL_INT_STAT_R	Reflects the status of wakeup interrupts and non-error general interrupts. It also has a bit to indicate whether any of the bits in SDHCx_CORE_ERROR_INT_STAT_R is set.
SDHCx_CORE_ERROR_INT_STAT_R	Reflects the status of general interrupts that are triggered by error conditions.
SDHCx_CORE_NORMAL_INT_STAT_EN_R	Provides mask bits for wakeup interrupts and non-error general interrupts.
SDHCx_CORE_ERROR_INT_STAT_EN_R	Provides mask bits for general interrupts that are triggered by error conditions.
SDHCx_CORE_NORMAL_INT_SIGNAL_EN_R	Setting any of these bits to '1' enables interrupt generation for wakeup interrupts and non-error general interrupts.
SDHCx_CORE_ERROR_INT_SIGNAL_EN_R	Setting any of these bits to '1' enables interrupt generation for general interrupts that are triggered by error conditions.
SDHCx_CORE_FORCE_ERROR_INT_STAT_R	Forces an error interrupt to occur when the corresponding bit is set.

33.6.1 SDIO Interrupt

The SDIO interrupt function is supported on SDHCx_CARD_DAT_3TO0_[1] line (SD pin 8). See the SDIO specifications for details on this feature. The SDHCx_CORE_NORMAL_INT_STAT_EN_R.CARD_INTERRUPT_STAT_EN bit and the SDHCx_CORE_NORMAL_INT_SIGNAL_EN_R.CARD_INTERRUPT_SIGNAL_EN bit must be set to enable this interrupt. To use this interrupt as wakeup interrupt, use SDHCx_CORE_WUP_CTRL_R.WUP_CARD_INT instead of SDHCx_CORE_NORMAL_INT_SIGNAL_EN_R.

33.7 I/O Interface

SDHC block provides the signals shown in [Table 33-4](#), which can be routed to pins through the I/O subsystem (IOSS). Refer to the [I/O System chapter on page 246](#) to configure the I/Os, and the device datasheet for specific pins available for each signal. SDHC also supports SDIO interrupt on DAT[1] line (SDHCx_CARD_DAT_3TO0_[1]). The output signals must be configured in strong drive mode, bi-directional signals in strong drive with the input buffer ON, and the input pins in high-impedance mode when an external pull-up resistor is available; otherwise, they must be configured in internal pull-up mode. Input buffer must be enabled for the input pins. The drive mode of the DAT lines must be set to high impedance after card removal. See [Card Detection on page 633](#) for details. In addition to configuring the drive mode and HSIOM registers in IOSS, the SDHCx_CORE_GP_OUT_R register must be configured to enable the required signals. See [Table 33-4](#). The SDHCx_CARD_DETECT_N and SDHCx_CARD_MECH_WRITE_PROT should be connected to ground if an eMMC or an embedded SDIO device is connected.

Table 33-4. I/O Signal Interface


Signal	Function	Register Configuration
SDHCx_CLK_CARD	Clock output	SDHCx_CORE_GP_OUT_R.CARD_CLOCK_OE
SDHCx_CARD_CMD	Command (bi-directional)	Always enabled
SDHCx_CARD_DAT_3TO0_[3:0]	Data (bi-directional)	SDHCx_CORE_HOST_CTRL1_R.DAT_XFER_WIDTH
SDHCx_CARD_DAT_7TO4_[3:0]	Data (bi-directional)	SDHCx_CORE_HOST_CTRL1_R.EXT_DAT_XFER
SDHCx_CARD_DETECT_N	Card detect signal input, Active low	SDHCx_CORE_GP_OUT_R.CARD_DETECT_EN
SDHCx_CARD_MECH_WRITE_PROT	Mechanical write protect signal input, Active low	SDHCx_CORE_GP_OUT_R.CARD_MECH_WRITE_PROT_EN
SDHCx_IF_PWR_EN	Card interface power enable output	SDHCx_CORE_GP_OUT_R.CARD_IF_PWR_EN_OE

33.8 Packet Buffer SRAM

SRAM that is internal to the SDHC block is used as a packet buffer to store data packets while carrying out data transfer to and from the card. The size of the SRAM is 1KB to support buffering of two 512 bytes blocks. As write and read transfers to the cards do not occur simultaneously, a single shared buffer is used for read and write operations. During the data transfer command handshake, the read/write bit of the command register is sampled and stored. This internal bit defines whether the SDHC is in read or write mode.

Figure 33-3 shows how data flows from the card interface to the AHB master interface through the packet buffer for a card read transfer. Received data from the card interface is written into packet buffer. When one block of data is received, DMA starts transmitting that data to the system by reading it from the packet buffer. For a card write transfer, data flows in the reverse direction. DMA writes data into a packet buffer that is subsequently read by the card interface logic. DMA and card interface logic can work simultaneously because read and write to packet buffer can be interleaved. For card read, DMA can send out the previous block while card interface logic is receiving the current block. For card write, DMA can write the current block into packet buffer while card interface logic is sending out the previous block.

Figure 33-3. Data Flow in a Read Transfer



33.8.1 Packet Buffer Full/Empty

When the packet buffer becomes full in card read, the clock to the card is stopped to prevent the card from sending the next data block. When packet buffer is empty, data block is not sent. In both cases, card interface logic is idle. SDHC does not support SDIO Read Wait signaling through DAT[2]. Therefore, the I/O command (CMD52) cannot be performed during a multiple read cycle because the card clock is stopped.


33.9 DMA Engine

The DMA engine handles data transfer between SDHC and system memory. Following are the features of this unit:

- Supports SDMA, ADMA2, and ADMA3 modes based on the configuration.
- The same DMA engine is used to interleave data transfer and descriptor fetch. This enables new task descriptor fetches (for CMD44 and CMD45) while DMA is moving data during task execution (for CMD46 and CMD47).
- Prefetches data for back-to-back eMMC write commands.
- Writes back the received data packets to system memory.

Figure 33-4 shows the data flow between the host driver and SD bus. The host driver can transfer data using either a programmed I/O (PIO) method in which the internal buffer is accessed through the buffer data port (SDHCx_CORE_BUF_DATA_R) register or using any of the defined DMA methods. PIO mode is much slower and burdens the processor. Do not use the PIO mode for large transfers.

Figure 33-4. Data Flow



DMA supports both single block and multi-block transfers. The control bits in the block gap control (SDHCx_CORE_BGAP_CTRL_R) register is used to stop and restart a DMA operation. SDMA mode is used for short data transfer because it generates interrupts at page boundaries. These interrupts disturb the CPU to reprogram the new system address. Only one SD command transaction can be executed for every SDMA operation.

ADMA2 and ADMA3 are used for long data transfers. They adopt scatter gather algorithm so that higher data transfer speed is available. The host driver can program a list of data transfers between system memory and SD card to the descriptor table. ADMA2 performs one read/write SD command operation at a time. ADMA3 can program multiple read/write SD command operation in a descriptor table.

In SDMA and ADMA2 modes, writing the SDHCx_CORE_CMD_R register triggers the DMA operation. In ADMA3 mode, writing SDHCx_CORE_ADMA_ID_LOW_R register triggers the DMA operation.

SD mode commands are generated by writing into the following registers – system address (SDHCx_CORE_SDMASA_R), block size (SDHCx_CORE_BLOCKSIZE_R), block count (SDHCx_CORE_BLOCKCOUNT_R), transfer mode (SDHCx_CORE_XFER_MODE_R), and command (SDHCx_CORE_CMD_R). When SDHCx_CORE_HOST_CTRL2_R.HOST_VER4_EN = 0, SDMA uses SDHCx_CORE_SDMASA_R as system address register and hence Auto CMD23 cannot be used with SDMA because this register is assigned for Auto CMD23 as the 32-bit block count register. When SDHCx_CORE_HOST_CTRL2_R.HOST_VER4_EN = 1, SDMA uses SDHCx_CORE_ADMA_SA_LOW_R as system address register and SDHCx_CORE_SDMASA_R is


reassigned to 32-bit block count and hence SDMA may use Auto CMD23.

To use the 32-bit block count register when SDHCx_CORE_HOST_CTRL2_R.HOST_VER4_EN = 1, it must be programmed with a non-zero value and the value of the 16-bit block count register SDHCx_CORE_BLOCKCOUNT_R must be zero. Refer to the respective specifications documents listed in [Block Diagram on page 627](#) to learn more about the DMA operation.

33.10 Initialization Sequence

Figure 33-5 shows the sequence for initializing SDHC to work with SD/SDIO/eMMC cards. Subsequent sections describe each step. After initialization, SDHC is ready to communicate with the card. Refer to the corresponding specifications document for information on other sequences such as card initialization and identification, changing bus speed mode, signal voltage switch procedure, transaction generation, and error recovery.

Figure 33-5. SDHC Programming Sequence



33.10.1 Enabling SDHC

Ensure CLK_GRx is configured to be greater than or equal to CLK_CARD and is running. Then, follow the sequence in [Figure 33-6](#) to enable the block. The internal clock can also be enabled later during clock setup. It must be enabled to

detect card insertion or removal through general interrupts when SDHC is not in standby mode. See [33.10.2 Card Detection](#) for details.

Figure 33-6. SDHC Enable Sequence





Figure 33-7. Card Status Check Sequence




33.10.2 Card Detection

Check if the card is already inserted by following the sequence shown in [Figure 33-7](#). This step is required for a removable card. After the card is detected, the host driver can supply power and clock to the card. If the card is inserted, then proceed with SDHC initialization. To detect card insertion or removal through interrupt when the internal clock is already enabled, follow the sequence shown in [Figure 33-8](#). To detect the card status through interrupt when the internal clock is disabled (when SDHC is in standby mode), the bits in the SDHCx_CORE_WUP_CTRL_R register must be set and the SDHCx_CORE_NORMAL_INT_SIGNAL_EN register must be cleared. See [Interrupts to CPU on page 629](#) for details. To detect SDIO card interrupt on DAT[1] line, a separate bit is provided in these registers, which must be configured.

SDHC clears the SDHCx_CORE_PWR_CTRL_R.SD_BUS_PWR_VDD1 bit when the card is removed and drives the DAT lines low. Therefore, the drive mode of the DAT lines must be changed from strong (with input buffer ON) to HI-Z when the card is removed to keep the lines pulled high. After detecting card insertion, the drive mode must be configured back to strong (with input buffer ON) mode only after SDHCx_CORE_PWR_CTRL_R.SD_BUS_PWR_VDD1 is set to 1.


Figure 33-8. Card Detection Through Interrupt



33.10.3 SDHC Initialization

To initialize SDHC, configure the basic settings as shown in [Figure 33-9](#). This step can also be executed immediately after enabling SDHC.


Figure 33-9. SDHC Setup



33.10.4 Clock Setup

Enable the internal clock followed by the card clock (SD clock) by following the sequence shown in [Figure 33-10](#). The SD clock frequency must be 100 kHz to 400 kHz during the card initialization. See [Card Clock \(SDCLK\) Configuration on page 628](#) for details. SD clock can be started and stopped by toggling the SDHCx_CORE_CLK_CTRL_R.SD_CLK_EN bit. The same sequence excluding the step of enabling the internal clock can be used to change the SD clock frequency. The SD clock must be stopped before changing its frequency. Note that SDHCx_CORE_GP_OUT_R.CARD_CLOCK_OE should have been set to '1' for the card clock to appear on the pin.

Figure 33-10. Clock Setup



* See the Traveo II Registers TRM for details.

33.11 Error Detection

SDHC can detect different types of errors in SD and eMMC transactions. Error is detected in either the command or data portion of the transaction. When an error is detected, SDHCx_CORE_NORMAL_INT_STAT_R.ERR_INTERRUPT bit is set. The exact error can then be identified through the SDHCx_CORE_ERROR_INT_STAT_R register. The Abort command is used to recover from an error detected during data transfer. In addition to these two registers, SDHC has two other error status registers – Auto CMD Error Status (SDHCx_CORE_AUTO_CMD_STAT_R) and ADMA Error Status (SDHCx_CORE_ADMA_ERR_STAT_R). The following table lists the errors detected by SDHC.

Table 33-5. Errors Detected by SHDC

Type	Error
Command Errors	Command Timeout Error Command CRC Error Command End Bit Error Command Index Error Command Conflict Error Response Error
Auto Command Errors	Command not issued by Auto CMD12 Error Auto Command Timeout Error Auto Command CRC Error Auto Command End Bit Error Auto Command Index Error Auto Command Conflict Error Auto CMD response Error
Data Errors	Data Timeout Error Data CRC Error Data End Bit Error ADMA Error Tuning Error

33.12 Register List

Table 33-6. SDHC Register List

Register	Name	Description
SDHCx_WRAP_CTL	Control Register	Enables SDHC
SDHCx_CORE_SDMASA_R	SDMA System Address Register	This register is used to configure a 32-bit block count or an SDMA system address based on the Host Version 4 Enable bit in the Host Control 2 register. This register is applicable to both SD and eMMC modes.
SDHCx_CORE_BLOCKSIZE_R	Block Size Register	This register is used to configure an SDMA buffer boundary and the number of bytes in a data block. This register is applicable to both SD and eMMC modes.
SDHCx_CORE_BLOCKCOUNT_R	16-bit Block Count Register	This register is used to configure the number of data blocks. This register is applicable to both SD and eMMC modes.
SDHCx_CORE_ARGUMENT_R	Argument Register	This register is used to configure the SD/eMMC command argument.
SDHCx_CORE_XFER_MODE_R	Transfer Mode Register	This register is used to control the operation of data transfers for the SD/eMMC mode.
SDHCx_CORE_CMD_R	Command Register	This register is used to provide the information related to a command and a response packet. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_RESP01_R	Response Register 0/1	This register stores 39-8 bits of the Response Field for the SD/eMMC mode.
SDHCx_CORE_RESP23_R	Response Register 2/3	This register stores 71-40 bits of the Response Field for the SD/eMMC mode.
SDHCx_CORE_RESP45_R	Response Register 4/5	This register stores 103-72 bits of the Response Field for the SD/eMMC mode.

Table 33-6. SDHC Register List

Register	Name	Description
SDHCx_CORE_RESP67_R	Response Register 6/7	This register stores 135-104 bits of the Response Field for the SD/eMMC mode.
SDHCx_CORE_BUF_DATA_R	Buffer Data Port Register	This register is used to access the packet buffer. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_PSTATE_REG	Present State Register	This register indicates the present status of the host controller. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_HOST_CTRL1_R	Host Control 1 Register	This register is used to control the operation of the host controller. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_PWR_CTRL_R	Power Control Register	This register is used to control the bus power for the card. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_BGAP_CTRL_R	Block Gap Control Register	This register is used by the host driver to control any operation related to block gap. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_WUP_CTRL_R	Wakeup Control Register	The register wakes up an otherwise idle host driver. It does NOT wake up from DeepSleep.
SDHCx_CORE_CLK_CTRL_R	Clock Control Register	This register controls SDCLK (card clock) in the SD/eMMC mode. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_TOUT_CTRL_R	Timeout Control Register	This register is used to set the Data Timeout Counter value for the SD/eMMC mode according to the timer clock defined by the Capabilities register, while initializing the host controller.
SDHCx_CORE_SW_RST_R	Software Reset Register	This register is used to generate a reset. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_NORMAL_INT_STAT_R	Normal Interrupt Status Register	This register reflects the status of the Normal Interrupt. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_ERROR_INT_STAT_R	Error Interrupt Signal Enable Register	This register enables an interrupt when the Error Interrupt Status Enable is enabled and at least one of the statuses is set to 1. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_NORMAL_INT_STAT_EN_R	Normal Interrupt Status Enable Register	This register enables the interrupt status for Normal Interrupt Status register (SDHCx_CORE_NORMAL_INT_STAT_R) when SDHCx_CORE_NORMAL_INT_STAT_R is set to 1. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_ERROR_INT_STAT_EN_R	Error Interrupt Status Enable Register	This register sets the Interrupt Status for Error Interrupt Status register (SDHCx_CORE_ERROR_INT_STAT_R), when SDHCx_CORE_ERROR_INT_STAT_EN_R is set to 1. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_NORMAL_INT_SIGNAL_EN_R	Normal Interrupt Signal Enable Register	This register is used to select the interrupt status that is indicated to the host system as the interrupt. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_ERROR_INT_SIGNAL_EN_R	Error Interrupt Signal Enable Register	This register is used to select the interrupt status that is notified to the host system as an interrupt. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_AUTO_CMD_STAT_R	Auto CMD Status Register	This register is used to indicate the CMD12 response error of Auto CMD12, and the CMD23 response error of Auto CMD23. This register is valid only when Auto CMD Error is set. This register is applicable to the SD/eMMC mode.

Table 33-6. SDHC Register List

Register	Name	Description
SDHCx_CORE_HOST_CTRL2_R	Host Control 2 Register	This register is used to control how the host controller operates. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_CAPABILITIES1_R	Capabilities 1 Register - 0 to 31	This register provides the host driver with information specific to the host controller implementation.
SDHCx_CORE_CAPABILITIES2_R	Capabilities 2 Register - 32 to 63	This register provides the host driver with information specific to the host controller implementation.
SDHCx_CORE_CURR_CAPABILITIES1_R	Current Capabilities Register - 0 to 31	This register indicates the maximum current capability for each voltage, for VDD1.
SDHCx_CORE_CURR_CAPABILITIES2_R	Maximum Current Capabilities Register - 32 to 63	This register indicates the maximum current capability for each voltage, for VDD2.
SDHCx_CORE_FORCE_AUTO_CMD_STAT_R	Force Event Register for Auto CMD Error Status Register	The register is not physically implemented but is an address at which the Auto CMD Error Status register can be written. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_FORCE_ERROR_INT_STAT_R	Force Event Register for Error Interrupt Status	This register is not physically implemented but is an address at which the Error Interrupt Status register can be written. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_ADMA_ERR_STAT_R	ADMA Error Status Register	This register stores the ADMA state during an ADMA error. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_ADMA_SA_LOW_R	ADMA System Address Register - Low	This register holds the lower 32-bit system address for DMA transfer. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_ADMA_ID_LOW_R	ADMA3 Integrated Descriptor Address Register - Low	This register holds the lower 32-bit integrated descriptor address. This register is applicable to the SD/eMMC mode.
SDHCx_CORE_HOST_CNTRL_VERS_R	Host Controller Version Register	This register is used to indicate the host controller version number.
SDHCx_CORE_MSHC_VER_ID_R	MSHC version Register	This register reflects the current release number.
SDHCx_CORE_MSHC_VER_TYPE_R	MSHC version type Register	This register reflects the current release type.
SDHCx_CORE_MSHC_CTRL_R	MSHC Control Register	This register is used to control the operation of MSHC host controller.
SDHCx_CORE_MBIU_CTRL_R	MBIU Control Register	This register is used to select the valid burst types that the AHB master bus interface can generate.
SDHCx_CORE_EMMC_CTRL_R	eMMC Control Register	This register is used to control the eMMC operation.
SDHCx_CORE_BOOT_CTRL_R	eMMC Boot Control Register	This register is used to control the eMMC boot operation.
SDHCx_CORE_GP_IN_R	General Purpose Input Register	This register is used as a general-purpose input register.
SDHCx_CORE_GP_OUT_R	General Purpose Output Register	This register is used as a general-purpose output register.

34. Audio Subsystem



The Inter-IC Sound Bus (I²S) is a serial bus interface standard used to connect digital audio devices together. The specification is from Philips® Semiconductor (I²S bus specification: February 1986, revised June 5, 1996). In addition to the standard I²S format, the I²S block also supports the Left Justified (LJ) format and the Time Division Multiplexed (TDM) format.

34.1 Features

- Supports standard Philips I²S, LJ, and eight-channel TDM digital audio interface formats
- Supports both master and slave mode operation in all the digital audio formats
- Supports independent operation of Receive (RX) and Transmit (TX) directions
- Supports operating from an external master clock provided through an external IC such as audio codec
- Provides configurable clock divider registers to generate the required sample rates
- Supports data word length of 8-bit, 16-bit, 18-bit, 20-bit, 24-bit, and 32-bit per channel
- Supports channel length of 8-bit, 16-bit, 18-bit, 20-bit, 24-bit, and 32-bit per channel (channel length fixed at 32-bit in TDM format)
- Provides two hardware FIFO buffers, one each for the TX block and RX block, respectively
- Supports both DMA- and CPU-based data transfers

34.2 Architecture

Figure 34-1. I²S Block Diagram




Figure 34-1 shows the high-level block diagram of the I²S block, which consists of two sub-blocks – I²S Transmitter (TX) and I²S Receiver (RX). The digital audio interface format and master/slave mode configuration can be done independently for the TX and RX blocks. In the master mode, the word select (WS) and serial data clock (SCK) are generated by the I²S block in the MCU. In the slave mode, the WS and SCK signals are inputs signals to the MCU, and are generated by the external master device. The I²S block configuration, control, and status registers, along with the FIFO data buffers are accessible through the AHB bus. AHB bus masters such as CPU and DMA can access the I²S registers through the AHB interface. Refer to the device datasheet for information on port pin assignments of the I²S block signals and AC/DC electrical specifications.

34.3 Digital Audio Interface Formats

The I²S block supports the following digital audio interface formats.

- Standard I²S format
- Left Justified format
- Time Division Multiplexed (TDM) format

The TX and RX sub-blocks can be independently configured to support one of the above formats in either master or slave mode. The I2S_MODE bits in the I2Sx_TX_CTL and I2Sx_RX_CTL registers are used to configure the digital audio interface format for the TX and RX blocks respectively. The MS (Master/Slave) bit in the I2Sx_TX_CTL and I2Sx_RX_CTL registers is used to configure the blocks in master or slave mode.

34.3.1 Standard I²S Format


Figure 34-2 shows the timing diagrams for the different word length and channel length combinations in the standard I²S digital audio format. In the standard I²S format, the word select signal (WS) is low for left channel data, and high for right channel data. The WS signal transitions one bit-clock (SCK) early relative to the start of the left/right channel data. All the serial data (SD), WS signal transitions on the falling edge of the SCK signal, and the read operations on the WS and sd lines are usually done on the rising edge of SCK. Therefore, the I²S TX block writes to the serial data (AUDIOSSx_TX_SDO) line on the falling edge of AUDIOSSx_TX_SCK, and the I²S RX block reads the data (AUDIOSSx_RX_SDI) on the rising edge of AUDIOSSx_RX_SCK. The serial data is transmitted most significant bit (MSb) first. Depending on whether the block is in master or slave mode, the WS/SCK signals are either generated by the block (master mode) or input signals to the block (slave mode).

The I²S block supports configurable word length and channel length selection options. The word length for the TX and RX blocks can be configured using the WORD_LEN bits in the I2Sx_TX_CTL and I2Sx_RX_CTL registers, respectively. The channel length for the TX and RX blocks can be configured using the CH_LEN bits in the I2Sx_TX_CTL and I2Sx_RX_CTL registers respectively. The channel length configuration should always be greater than or equal to the word length configuration. Ensure that when the I²S RX block is operated in slave mode, the master TX device ensures that its channel length configuration aligns with the I²S RX block channel length setting. If there is channel length mismatch, the I²S RX block in slave mode will not operate correctly.


In the TX block, when the channel length is greater than the word length, the unused bits can be transmitted either as '0' or '1'. This selection is made using the OVHDATA bit in the I2Sx_TX_CTL register. In the RX block, when the word length is less than 32 bits, the unused most significant bits written to the 32-bit RX FIFO register can either be set to '0' or sign bit extended. This selection is made using the BIT_EXTENSION bit in the I2Sx_RX_CTL register.

Figure 34-2. Standard I²S Format (Word Length and Channel Length Combination Timing Diagrams)


(1) Channel Length = 32-bits



(2) Channel Length = 24-bits



(3) Channel Length = 20-bits



(4) Channel Length = 8-bits




Table 34-1 lists the supported word length and channel length combinations.

Table 34-1. Word Length and Channel Length Combinations

		Word Length					
		8-bit	16-bit	18-bit	20-bit	24-bit	32-bit
Channel Length	32-bit	Valid	Valid	Valid	Valid	Valid	Valid
	24-bit	Valid	Valid	Valid	Valid	Valid	Invalid
	30-bit	Valid	Valid	Valid	Valid	Invalid	Invalid
	18-bit	Valid	Valid	Valid	Invalid	Invalid	Invalid
	16-bit	Valid	Valid	Invalid	Invalid	Invalid	Invalid
	8-bit	Valid	Invalid	Invalid	Invalid	Invalid	Invalid

34.3.2 Left Justified (LJ) Format

Figure 34-3 shows the timing diagrams for the Left Justified interface format using the 32-bit channel length and 32-bit word length configuration as an example. The only differences between the standard I²S and LJ formats are:


- In the standard I²S format, WS signal is low for left channel data and high for right channel data. In the LJ

format, WS signal is high for left channel data and low for right channel data.

- In the standard I²S format, WS signal transitions one bit-clock (SCK) early relative to the start of the channel data (coincides with LSb of the previous channel). In the LJ format, there is no early transition, and the WS signal transitions coincide with the start of the channel data.

Apart from these differences, all the features explained in 34.3.1 Standard I²S Format apply to the LJ format as well.

Figure 34-3. Left Justified Digital Audio Format



34.3.3 Time Division Multiplexed (TDM) Format


Figure 34-4 shows the timing diagrams for the two types of Time Division Multiplexed (TDM) formats supported by the I²S block. The differences between the standard I²S/LJ formats and the TDM format are as follows:

- Standard I²S/LJ formats support only two channels (left/right) per frame, while TDM format supports up to eight channels per frame.
- In the TDM format, channel length for all eight channels is fixed at 32 bits. In the standard I²S/LJ formats, the channel length is configurable. The word length per channel is configurable similar to the standard I²S and the data is also transmitted most significant bit first. Similar to I²S, when the word length per channel is less than the 32-bit channel length for TX block, the OVHDATA bit in the I2Sx_TX_CTL register is used to fill the unused least significant channel data bits with either all zeros or all ones

- In the TDM format, all eight channels of data are always present in a frame, and thus the frame width is fixed at 256 bits. You have the option to configure the number of active channels in a frame by configuring the CH_NR bits in the I2Sx_TX_CTL and I2Sx_RX_CTL registers. In the standard I²S/LJ format, the CH_NR should always be configured for two channels. The number of active channels in the TDM format can be less than or equal to eight channels. The unused (inactive) channels always follow the active channels in a frame. As an example, if CH_NR is set for four channels, CH0 to CH3 are the active channels and CH4 to CH7 are the unused channels. The OVHDATA bit in the I2Sx_TX_CTL register is used to fill the unused channels with either all zeros or all ones.
- The pulse width of the word select (WS) signal in the TDM format can be configured to be either one bit clock (SCK) wide or one channel wide. The selection is made using the WS_PULSE bit in the I2Sx_TX_CTL and I2Sx_RX_CTL registers. The pulse width is fixed to one channel width in the I²S/LJ format.

- Two types of TDM formats are supported. In TDM mode A, the WS rising edge signal to signify the start of frame coincides with the start of CH0 data. In TDM mode B, the WS rising edge signal to signify the start of frame is one bit clock (SCK) early, relative to the start of CH0 data (coincides with the last bit of the previous frame). The selection between the two TDM formats is made using the I2Sx_MODE bits in the I2Sx_TX_CTL and I2Sx_RX_CTL registers.

Figure 34-4. TDM Digital Audio Interface Format



34.4 Clocking Polarity and Delay Options

The I²S block supports configurable clock polarity and delay options to alleviate any timing issues in the system involving PCB signal propagation delays, and delays associated with internal device signal routing.

When the I²S TX block operates in the slave mode, the AUDIOSSx_TX_SCK and AUDIOSSx_TX_WS signals are input signals to the MCU, and the AUDIOSSx_TX_SDO output signal is transmitted off the AUDIOSSx_TX_SCK

falling edge. The AUDIOSSx_TX_SDO signal is sampled by the external master device RX block on the subsequent AUDIOSSx_TX_SCK rising edge. Timing issues arise if the AUDIOSSx_TX_SDO signal reaching the master side RX block does not meet the setup and hold time requirements for input data on the master side. The I²S TX block in the MCU has an option to advance the serial data transmission by 0.5 SCK cycles when the B_CLOCK_INV bit in the I2Sx_TX_CTL register is set. This feature can be used if there are timing issues while operating the I²S TX block in slave mode.

Table 34-2. TX Block Configuration for Master Mode

TX Block Configuration	Clock Polarity Register	Description
	I2Sx_TX_CTL.SCKO_POL	
(1)	0	Serial data is transmitted off the SCK falling edge
(2)	1	Serial data is transmitted off the SCK rising edge

Table 34-3. TX Block Configuration for Slave Mode

TX Block Configuration	Clock Polarity Register	Delay Option Register	Description
	I2Sx_TX_CTL.SCKI_POL	I2Sx_TX_CTL.B_CLOCK_INV	
(1)	0	0	Serial data is transmitted off the SCK falling edge
(2)	0	1	Serial data is transmitted off the SCK rising edge that is 0.5 SCK cycles before (1)
(3)	1	0	Serial data is transmitted off the SCK rising edge
(4)	1	1	Serial data is transmitted off the SCK falling edge that is 0.5 SCK cycles before (3)

Similarly, when the I²S RX block operates in the master mode, the AUDIOSSx_RX_SCK and AUDIOSSx_RX_WS signals are output signals from the MCU, and the AUDIOSSx_RX_SDI signal is transmitted by the external master device on the falling edge of AUDIOSSx_RX_SCK. The I²S RX block samples the AUDIOSSx_RX_SDI signal on the subsequent AUDIOSSx_RX_SCK rising edge. Timing issues arise if the AUDIOSSx_RX_SDI signal reaching the MCU block does not meet the setup and hold time requirements for input data. The I²S RX block has an option to delay the serial data capture by 0.5 SCK cycles when the B_CLOCK_INV bit in the I2Sx_RX_CTL register is

set. This feature can be used if there are timing issues while operating the I²S RX block in master mode.

In addition to these clock delay options, there is also an option to invert the outgoing bit clock (SCK) in master mode by setting the SCKO_POL bit in the I2Sx_TX_CTL and I2Sx_RX_CTL registers. Similarly, in the slave mode, there is an option to invert the incoming bit clock (SCK) by setting the SCKI_POL bit in the I2Sx_TX_CTL and I2Sx_RX_CTL registers.

Refer to the *Traveo II Body Controller High Registers TRM* for detailed description of the B_CLOCK_INV, SCKI_POL, and SCKO_POL register configurations.

Table 34-4. RX Block Configuration for Master Mode

RX Block Configuration	Clock Polarity Register	Delay Option Register	Description
	I2Sx_RX_CTL.SCKO_POL	I2Sx_RX_CTL.B_CLOCK_INV	
(1)	0	0	Serial data is captured by the SCK rising edge
(2)	0	1	Serial data is captured by the SCK falling edge that is 0.5 SCK cycles after (1)
(3)	1	0	Serial data is captured by the SCK falling edge
(4)	1	1	Serial data is captured by the SCK rising edge that is 0.5 SCK cycles after (3)

Table 34-5. RX Block Configuration for Slave Mode

RX Block Configuration	Clock Polarity Register	Description
	I2Sx_RX_CTL.SCKI_POL	
(1)	0	Serial data is captured by SCK rising edge
(2)	1	Serial data is captured by SCK falling edge

34.5 Interfacing with Audio Codecs


The I²S block in the MCU interfaces with an audio codec device based on the choice of codec device and the end application requirements. Some scenarios and the connection diagrams are as follows:

- Codecs with separate WS and SCK signals for the RX and TX directions: To interface with these codecs, the connections between the I²S block and the codec device will be as shown in [Figure 34-5](#) where the I²S TX signals (AUDIOSSx_TX_SCK, AUDIOSSx_TX_WS, AUDIOSSx_TX_SDO) connect to the codec RX signals, and the I²S RX signals (AUDIOSSx_RX_SCK, AUDIOSSx_RX_WS, AUDIOSSx_RX_SDI) connect to the codec TX signals. The direction of SCK (AUDIOSSx_TX_SCK, AUDIOSSx_RX_SCK) and WS (AUDIOSSx_TX_WS, AUDIOSSx_RX_WS) signals

depends on which device is the master and which device is the slave.

- Codecs with common WS and SCK signals for both RX and TX directions: There are two possible configurations to interface these codecs with the MCU as shown in [Figure 34-5](#). In both configurations, the SCK signals (AUDIOSSx_TX_SCK, AUDIOSSx_RX_SCK, codec_sck) are shorted externally. The same goes for the WS signal connections as well (AUDIOSSx_TX_WS, AUDIOSSx_RX_WS, codec_ws). Ensure that only one block is driving the SCK and WS lines. So when the codec acts as the slave device, the I²S RX block should be in the master mode, and the I²S TX block should be in the slave mode (or I²S RX as slave and I²S TX as master). When the codec acts as the master device, both the I²S RX and I²S TX blocks should be in slave mode.

Figure 34-5. Interfacing with Codecs having Common WS and SCK Signals



34.6 Clocking Features


The I²S unit has three clock inputs.

Table 34-6. Clock Inputs

Signal	DESCRIPTION
CLK_SYS_I2S	System clock. This clock is used for the AHB slave Interface, control, status, and interrupt registers, and also clocks the DMA trigger control logic.
CLK_AUDIO_I2S	I ² S internal clock. This clock is used for I ² S transmitter (TX)/receiver (RX) blocks; it is asynchronous with the CLK_SYS_I2S. This clock is connected to the CLK_HFX high-frequency clock in the device. Refer to the Clocking System chapter on page 197 for more details on high frequency clocks.
AUDIOSSx - CLK_I2S_IF	I ² S external clock. This clock is provided from an external I ² S bus host through a port pin. It is used in place of the CLK_AUDIO_I2S clock to synchronize I ² S data to the clock used by the external I ² S bus host.

[Figure 34-6](#) shows the clocking divider structure in the I²S block. In the master mode, the SCK and WS signals are generated either using the CLK_AUDIO_I2S internal clock or the AUDIOSSx_CLK_I2S_IF external clock. Refer to the device datasheet for the port pin assignment of AUDIOSSx_CLK_I2S_IF clock. The CLOCK_SEL bit in the I2Sx_CLOCK_CTL register controls the selection between internal and external clocks.

Figure 34-6. Clocking Divider Structure



There are two stages of clock dividers in the I²S block as follows.

- The first stage clock divider is used to generate the internal I²S master clock (MCLK_SOC). The input clock to the first stage divider is either CLK_AUDIO_I2S or AUDIOSSx_CLK_I2S_IF. The first stage clock divider is configured using the CLOCK_DIV bits in I2Sx_CLOCK_CTL register. Divider values from 1 to 64 are supported.
- The second stage clock divider is used to generate the SCK signals. The input clock is the output from the first stage clock divider. This divider value is fixed at '8' (FTX_SCK = FRX_SCK = FMCLK_SOC/8). The word select (WS) signal frequency depends on the SCK frequency, and the configured channel length value.

When in slave mode, the internal clock (MCLK_SOC) frequency should still be eight times the frequency of the

input serial clock. You must choose the appropriate clock source and the CLOCK_DIV divider value to guarantee this condition is met in the slave mode of operation. Usually, when the I²S block operates in the slave mode, the host sends a master clock which is an integral multiple of the sampling rate. This master clock can be routed to the AUDIOSSx_CLK_I2S_IF port pin. The CLOCK_DIV divider value can then be adjusted to ensure that the MCLK_SOC is eight times the input SCK frequency.

Table 34-7 gives an example of the clock divider settings for operating the I²S block at the standard sampling rates in the standard I²S format. Note that the first stage divider values in the table are the register field values – the actual divider values are one more than the configured register values as explained in the clock divider section. Refer to the device datasheet for details on maximum values of SCK frequency, and the output sampling rates.

 Table 34-7. I²S Divider Values for Standard Audio Sampling Rates in Standard I²S Format

Sampling Rate (SR) (kHz)	WORD_LEN (bits)	SCK (2*WORD_LEN*SR) (MHz)	CLK_HF (MHz)	CLK_HF/SCK (Total Divider Ratio)	CLK_CLOCK_DIV (First Divider)	Second Stage Divider (Fixed at 8)
8	32	0.512	49.152	96	11	8
16	32	1.024	49.152	48	5	8
32	32	2.048	49.152	24	2	8
48	32	3.072	49.152	16	1	8
44.1	32	1.4112	45.1584	32	3	8

34.7 FIFO Buffer and DMA Support

The I²S block has two FIFO buffers – one each for the TX block and RX block, respectively. The ordering format of the channel data in both TX and RX FIFOs depends on the configured digital audio format. This ordering format should be considered when writing to the TX FIFO or reading from the RX FIFO. In the standard I²S and LJ digital audio

formats, the ordering of the data is (L, R, L, R, L, ...) where L refers to the left channel data and R refers to the right channel data. In the TDM format with the number of active channels set to four, the data order will be (CH0, CH1, CH2, CH3, CH0, CH1, CH2, CH3, CH0, ...). If the number of active channels is set to eight, the cycle will repeat after CH0–CH7 data.

Table 34-8. I²S FIFO Buffers

Feature	TX FIFO	RX FIFO
Architecture	256 depth FIFOs for up to 32-bit data elements	
Data Register	I2Sx_TX_FIFO_WR	I2Sx_RX_FIFO_RD
Data format	Right-aligned	Right-aligned Receive data is extended by zeros or the sign-bit
Trigger control register for DMA	I2Sx_TR_CTL.TX_REQ_EN	I2Sx_TR_CTL.RX_REQ_EN
Trigger level	When the TX FIFO has less entries than I2Sx_TX_CTL.TRIGGER_LEVEL, a transmitter trigger event is generated	When the RX FIFO has more entries than I2Sx_RX_CTL.TRIGGER_LEVEL, a receiver trigger event is generated

I²S TX FIFO: The I²S TX block has a hardware FIFO of depth 256 elements where each element is 32-bit wide. In addition to this 256-element FIFO, the I²S block has an internal transmit buffer that can store four 32-bit data to be transmitted. This four-element buffer is used as an intermediary to hold data to be transferred on the I²S bus, and is not exposed to the AHB BUS interface.

The I2Sx_TX_FIFO_CTL register is used for FIFO control operations. The TRIGGER_LEVEL bits in the I2Sx_TX_FIFO_CTL register can be used to generate a transmit trigger event when the TX FIFO has less entries than the value configured in the TRIGGER_LEVEL bits.

The FIFO freeze operation can be enabled by setting the FREEZE bit in the I2Sx_TX_FIFO_CTL register. When the FREEZE bit is set and the TX block is operational (TX_START bit in I2Sx_CMD is set), hardware reads from the TX FIFO do not remove the FIFO entries. Also, the TX FIFO read pointer will not be advanced. Any writes to the I2Sx_TX_FIFO register will increment the TX FIFO write pointer; when the TX FIFO becomes full, the internal write pointer stops incrementing. The freeze operation may be used for firmware debug purposes. This operation is not intended for normal operation. To return to normal operation after using the freeze operation, the I²S must be reset by clearing the TX_ENABLED bit in the I2Sx_CTL register, and then setting the bit again.

The CLEAR bit in the I2Sx_TX_FIFO_CTL register is used to clear the TX FIFO by resetting the read/write pointers associated with the FIFO. Write access to the TX FIFO using the I2Sx_TX_FIFO_WR register is not allowed while the CLEAR bit is set.

The I2Sx_TX_FIFO_STATUS register provides FIFO status information. This includes number of used entries in the TX FIFO and the current values of the TX FIFO read/write pointers. This register can be used for debug purposes. The I²S TX FIFO read pointer is updated whenever the data is transferred from the TX FIFO to the internal transmit buffer. TX FIFO write pointer is updated whenever the data is written to the I2Sx_TX_FIFO_WR register, either through the CPU or the DMA controller.

For TX FIFO data writes using the CPU, the hardware can be used to trigger an interrupt event for any of the FIFO conditions such as TX_TRIGGER, TX_NOT_FULL, and TX_EMPTY. As part of the interrupt handler, the CPU can write to the I2Sx_TX_FIFO_WR register. The recommended method is to write (256 – TRIGGER_LEVEL) words to the I2Sx_TX_FIFO_WR register every time the TX_TRIGGER interrupt event is triggered. In addition, interrupt events can be generated for FIFO overflow/underflow conditions.

For DMA-based TX data transfers, the I²S TX DMA trigger signal (tr_i2s_tx_req) can be enabled by writing '1' to the TX_REQ_EN bit in I2Sx_TR_CTL register. The trigger signal output will become high whenever the TX FIFO has less entries than that configured in the TRIGGER_LEVEL field. The DMA channel can be configured to transfer up to (256 – TRIGGER_LEVEL) words from the applicable source address (such as Flash and SRAM regions). The destination address of the DMA should always be the I2Sx_TX_FIFO_WR register address, with the destination address increment feature disabled in the DMA channel configuration. This FIFO address increment logic is handled internally to adjust the write pointer, and the DMA should not increment the destination address. For more details on DMA

The data in the I2Sx_TX_FIFO_WR is always right-aligned. The I2Sx_TX_FIFO_WR format for different word length configurations is provided in [Figure 34-7](#).

Figure 34-7. I2Sx_TX_FIFO_WR Register Format for Different Word Lengths

	write data format of I2Sx_TX_FIFO_WR																															
	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 24-bit mode									MSb																LSb							
	fixed "0"								23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 20-bit mode									MSb																LSb							
	fixed "0"												19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 18-bit mode									MSb																LSb							
	fixed "0"														17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 16-bit mode									MSb																LSb							
	fixed "0"																15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

I²S RX FIFO write pointer is updated whenever the data is transferred to the RX FIFO from the internal receive buffer. RX FIFO read pointer is updated whenever the data is read from the I2Sx_RX_FIFO_RD register, either through the CPU or the DMA controller. For debug purposes, the I2Sx_RX_FIFO_RD_SILENT register is available, which always returns the top element of the RX FIFO without updating the read pointer.

For RX FIFO data reads using the CPU, the hardware can be used to trigger an interrupt event for any of the FIFO conditions such as RX_TRIGGER, RX_NOT_EMPTY, and RX_FULL. As part of the interrupt handler, the CPU can read from the I2Sx_RX_FIFO_RD register. The recommended method is to read (TRIGGER_LEVEL + 1) words from the I2Sx_RX_FIFO_RD register every time the RX_TRIGGER interrupt event is triggered. In addition, interrupt events can be generated for FIFO overflow/underflow conditions.

For DMA-based RX data transfers, the I²S RX DMA trigger signal (`tr_i2s_rx_req`) can be enabled by writing '1' to the `RX_REQ_EN` bit in the `I2Sx_TR_CTL` register. The trigger signal output will become high whenever the RX FIFO has more entries than that configured in the `TRIGGER_LEVEL` field. The DMA channel can be configured to transfer up to (`TRIGGER_LEVEL + 1`) words to the applicable destination address (such as SRAM regions). The source address of the DMA should always be the `I2Sx_RX_FIFO_RD` register address, with the source address increment feature disabled in the DMA channel configuration. This FIFO address increment logic is handled internally to adjust the read pointer, and the DMA should not increment the source address. For more details on DMA channel configuration, refer to the [Direct Memory Access chapter on page 71](#).

The data in the I2Sx_RX_FIFO_RD is always right aligned. The I2Sx_RX_FIFO_RD format for different word length configurations is provided in [Figure 34-8](#). Note that the unused most significant bits are either set as '0' or sign-bit extended depending on the BIT_EXTENSION bit in the I2Sx_RX_CTL register.

		read data format of I2Sx_RX_FIFO_RD																																			
		31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				
Word Length = 24-bit mode not Bit extension										MSb																				LSb							
		fixed "0"								23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				
Word Length = 20-bit mode not Bit extension										MSb																				LSb							
		fixed "0"																19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
Word Length = 18-bit mode not Bit extension										MSb																				LSb							
		fixed "0"																17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
Word Length = 16-bit mode not Bit extension										MSb																				LSb							
		fixed "0"																15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				

The I²S block has one interrupt output signal that goes to the interrupt controller in the CPU. Refer to the [Interrupts chapter on page 144](#) for details on the vector number of the I²S interrupt and the procedure to configure the interrupt priority, vector address, and enabling/disabling. The I²S interrupt can be triggered under any of the following events – TX_TRIGGER, TX_NOT_FULL, TX_EMPTY, TX_OVERFLOW, TX_UNDERFLOW, TX_WD, RX_TRIGGER, RX_NOT_EMPTY, RX_FULL, RX_OVERFLOW, RX_UNDERFLOW, or RX_WD. Each of the interrupt events can be individually enabled/disabled to generate an interrupt condition. The I2Sx INTR MASK

Interrupts can also be triggered in software by writing to the corresponding bits in I2Sx_INTR_SET register. [Figure 34-9](#) illustrates the interrupt signal generation.

```

graph LR
    A[HW event for an interrupt] --> OR1[OR]
    B[INTR_SET.XXX] --> OR1
    OR1 --> D[D Flip-Flop]
    D -- Q --> INTR_XXX[INTR.XXX]
    INTR_XXX --> AND[AND]
    C[INTR_I2S_MASK.XXX] --> AND
    AND --> OR2[OR]
    D[Other interrupt signals  
(INTR_I2S & INTR_I2S_MASK)] --> OR2
    OR2 --> E[audioss_x_interrupt_i2s_IRQn]
  
```


signal will always be high. A dummy read of the I2Sx_INTR register should be done for the earlier register write to take effect.

34.9 Watchdog Timer

The TX and RX blocks have independent watchdog timers, which can be used to generate an interrupt event if the Word Select (WS) input is idle for more than the configured time period. This feature is available only in the slave mode of operation where the external master drives the WS input lines (AUDIOSSx_TX_WS, AUDIOSSx_RX_WS). This feature can be used to detect any signal transmission issues, master device issues, or if the master has halted communication. If the master drives the same word select signal to both the AUDIOSSx_TX_WS and AUDIOSSx_RX_WS lines, then only one of the watchdog timers can be enabled to cause the interrupt event. Although the following explanation covers TX watchdog, the same explanation applies to RX watchdog as well.

To enable the TX watchdog timer feature, WD_EN bit in the I2Sx_TX_CTL register should be set. The watchdog timer reload value (32-bit timer) is configured by writing to the I2Sx_TX_WATCHDOG register. A value of zero written to the I2Sx_TX_WATCHDOG register will also disable the watchdog timer. Figure 34-10 illustrates the watchdog behavior when the timer is enabled. The timer runs off the CLK_PERI system clock. Refer to the [Clocking System chapter on page 197](#) for details on generation of CLK_PERI. The timer starts running when WD_EN and TX_START bits are set. The timer reload happens either on a rising edge event on AUDIOSSx_TX_WS input signal, or when the timer values reaches zero. When the timer value reaches zero, the TX_WD interrupt event is generated. The TX_WD bit in the I2Sx_INTR_MASK register should be set to enable interrupt generation by the watchdog timer interrupt event. The interrupt event can be cleared by writing '1' to the TX_WD bit in the I2Sx_INTR register.

Figure 34-10. Watchdog Timer Working




34.10 MCLK Output Function

The I²S unit generates MCLK output signal for an external audio codec. The MCLK output signal is generated only when the following conditions are met:

- I2Sx_CTL.TX_ENABLE or RX_ENABLE = 1 (I²S is enabled)
- I2Sx_CLOCK_CTL.CLOCK_SEL = 0 (I²S clock is from internal clock: CLK_HF_x)
- I2Sx_CLOCK_CTL.MCLK_DIV = 0, 1, 2, or 3 (Division ratio: 1, 2, 4, or 8)
- I2Sx_CLOCK_CTL.MCLK_EN = 1 (MCLK output enabled)

Figure 34-11 illustrates the using of MCLK output signal for the external Audio codec.

Figure 34-11. MCLK Output Signal for External Audio Codec



34.11 Register List

Register	Name	Description
I2Sx_CTL	I ² S Control Register	This register is used to control the I ² S block
I2Sx_CLOCK_CTL	I ² S Clock Control Register	This register can configure frequencies for I ² S block and master clock. It should not be changed during communication
I2Sx_CLOCK_STAT	I ² S Clock Status Register	This register indicates the status of the MCLK divider and is used to allow a controlled shutdown of the clock
I2Sx_CMD	I ² S Command Register	This register can control I ² S functions
I2Sx_TR_CTL	I ² S Trigger Control Register	This register can enable DMA for I ² S. It should not be changed during communication
I2Sx_TX_CTL	I ² S Transmitter Control Register	This register can configure the transmitter for I ² S functions. It should not be changed during communication
I2Sx_TX_WATCHDOG	I ² S Transmitter Watchdog Register	This register can configure counter value for the TX watchdog. It should not be changed during the operation
I2Sx_RX_CTL	I ² S Receiver Control Register	This register can configure the receiver for I ² S functions. It should not be changed during communication
I2Sx_RX_WATCHDOG	I ² S Receiver Watchdog Register	This register can configure counter value for the RX watchdog. It should not be changed during the operation
I2Sx_TX_FIFO_CTL	I ² S TX FIFO Control Register	This register can configure TX FIFO for I ² S
I2Sx_TX_FIFO_STATUS	I ² S TX FIFO Status Register	This register displays the status of TX FIFO for I ² S
I2Sx_TX_FIFO_WR	I ² S TX FIFO Write Register	This register stores transmission data for I ² S transfer
I2Sx_RX_FIFO_CTL	I ² S RX FIFO Control Register	This register can configure RX FIFO for I ² S
I2Sx_RX_FIFO_STATUS	I ² S RX FIFO Status Register	This register displays the status of RX FIFO for I ² S
I2Sx_RX_FIFO_RD	I ² S RX FIFO Read Register	This register stores received data for I ² S transfer
I2Sx_RX_FIFO_RD_SILENT	I ² S RX FIFO Silent Read Register	This register stores received data for I ² S transfer
I2Sx_INTR	I ² S Interrupt Register	This register displays the interrupt status
I2Sx_INTR_SET	I ² S Interrupt Set Register	This register can be used to trigger an interrupt for firmware testing.
I2Sx_INTR_MASK	I ² S Interrupt Mask Register	This register controls whether the I ² S interrupt is forwarded to the corresponding processor
I2Sx_INTR_MASKED	I ² S Interrupt Masked Register	Bitwise AND between the interrupt request and mask registers, so firmware can read the status of all mask enabled interrupt causes with a single load operation

Section F: Analog Subsystem




This section encompasses the following chapter:

- [SAR ADC chapter on page 654](#)

Top Level Architecture

Analog System Block Diagram



35. SAR ADC



The Traveo II features a successive approximation register analog-to-digital converter (SAR ADC). The SAR ADC is designed for applications that require a moderate resolution and high data rate. It consists of the following blocks:

- SARADC Core
- SARMUX
- SAR Sequencer
- Diagnostic Reference
- Reference Buffer

SARMUX is an analog multiplexer to connect the signal sources to the ADC input; SARADC core then performs analog-to-digital conversion. A SAR sequencer is responsible for prioritizing the triggers requests, enable the appropriate analog channel, and control the sampling.

A single-ended SAR ADC system is capable of scanning up to 40 analog inputs (32 GPIO and eight internal signals) as shown in the block diagram.


35.1 Features

- SAR ADC Core
 - 12-bit resolution with a maximum sample rate of 1 Msps
- 32 logical channels with the same capabilities
- Each logical channel can select input from
 - 32 analog input pins
 - Diagnostic signals
 - Analog input pins of other ADC units
 - Support for external mux (three select bits)
 - AMUXBUS/A/B
- Scans triggered by timer, software, continuous, pins, or system triggers
 - Multiple ADC units can be triggered by the same trigger to ensure lock-step operation
 - Triggers can be cleared by software
 - Optional debug pause
- Double buffering of output data
- Programmable sample time for each channel
- Programmable post processing options for each channel
 - Sign/zero extension to 16-bit
 - Left/right alignment
 - Averaging: first order accumulate and dump, up to 256 samples
 - Programmable right shift
 - Range detection: below/above threshold, in/out-side range
 - Pulse detection: programmable positive and negative event counters

- Channels can be individual or grouped
 - Flexible grouping: from 32 groups with one channel to one group with 32 channels
- Group scans are dynamically scheduled by the hardware
 - eight priorities, programmable per group
 - Four preemption types: resume, restart, cancel, or finish
 - Optional automatic idle power down
- Interrupt generation
 - Group scan done
 - Group scan done overflow detect
 - Group scan canceled
 - Per channel range detect
 - Per channel pulse detect
 - Per channel pulse/range overflow detect
- Output trigger generation per channel
 - Data ready/completion (each channel can trigger DW transfer)
 - Range violation detected
- Digital and analog calibration available
- Programmable offset and gain calibration
 - Non-intrusive background recalibration
 - Coherent calibration update
- Support for diagnostic measurements including broken wire detection. This includes:
 - ADC sampling capacitor preconditioning feature
 - Selectable current source or sink on selected ADC input while sampling
 - Support for LED diagnostics (see [35.7.1 Trigger Outputs](#) for details)
- On chip temperature sensor and power monitoring

35.2 Block Diagram

Figure 35-1. SAR ADC System Block Diagram



35.3 Operation

SAR conversion begins when a trigger signal is received by the SAR sequencer. The sequencer selects the appropriate analog input for the logical channel to be converted and triggers the analog-to-digital conversion and manages conversion results. If the sequencer is performing a group conversion, it proceeds to the next channel in the group and starts another conversion.

To perform a conversion, the ADC starts with optional preconditioning interval, which pre-charges or discharges the sampling capacitor, if preconditioning is enabled. The appropriate analog switches are then enabled, and the desired input is connected to the ADC. Signal sampling occurs when the ADC core start conversion signal (STC) goes high (from sequencer). STC remains high during the signal sampling window for a user specified number of clocks. At the end of the sampling window, STC goes low starting a 13 to 15 clock conversion cycle. At the end of conversion, the results can be post-processed; for example, averaging, range detection, and pulse detection. Results are

stored in dedicated double-buffered locations for each virtual channel. [Figure 35-2](#) shows a simple block diagram of the ADC core.

Several SARs can operate in lock-step for simultaneous conversion of analog inputs. This configuration is useful for brushless motor control, multi-phase power conversion, and other applications where simultaneous sampling is needed.

The analog input multiplexer (SARMUX) is implemented with two-levels of transmission gate switches. Input selection is performed by addresses stored in the logical channel configuration. Each channel can select an input signal, a diagnostic reference signal, or both.

The standard analog multiplexer has 32 inputs for signals from I/O pins, and supports another eight channels to measure special internal signals such as a bandgap reference voltage, a temperature sensor voltage, power supply pins, and long-reach signals to other input pads through GPIO AMUXBUSA/B signals. One channel is also reserved for motor sensing inputs. The SARMUX can use


an expansion signal to reach other SARMUXs when their ADCs are not in use.

The following sections will include detailed descriptions of various aspects of the ADC system including:

- SAR ADC Core
- SARMUX
- SAR Sequencer
- Triggering and Scheduling
- Output Triggers and Interrupts
- Diagnostic Reference Generator
- Reference Buffer

Note: Do not divide CLK_GR9; this makes CLK_GR9 = CLK_PERI keeping all the clocks coming to SAR ADC at the same frequency. If these clocks are not equal, this can cause GRP_CANCELLED bit to be set incorrectly during ABORT_CANCEL preemption.

Figure 35-2. ADC Core Block Diagram




35.3.1 SAR ADC Conversion Flow

Analog-to-digital conversion starts with a trigger received by the sequencer. All the channels in the group are then sequentially scanned and converted until the end of the group. After the end channel of the group is converted, the group conversion done interrupt is set. The size of the group can vary from one to 31 channels. Figure 35-3 shows the flow of the analog-to-digital conversion, and the important steps are explained here:

1. Analog-to-digital conversion request trigger and the trigger pending bit (PASSx_SARy_TR_PEND.TR_PEND) is set for the first logical channel of the group. The first channel should be triggered to scan the complete group.
2. If the ADC is in power-down state, a latency is added before the first conversion due to power up settling time. This latency is software configurable (PASSx_SARy_CTL.PWRUP_TIME).
3. Sampling the analog input to the mapped logical channel; sample time is configurable for each channel PASSx_SARy_CHz_SAMPLE_CTL.SAMPLE_TIME.
4. Analog-to-digital conversion of the sampled input.
5. Sample and analog-to-digital conversion for the next channels in the group (repeat steps 3 and 4).
6. Conversion of the last channel in the group is completed. The new result is coherently updated in the PASSx_SARy_CHz_RESULT register of all the channels. Group conversion done interrupt is set and trigger pending bit is cleared.
7. The group conversion done interrupt is cleared by writing '1' to the Interrupt Request Register PASSx_SARy_CHz_INTR.
8. Start the next group conversion if there is a trigger or go to power-down state if idle and auto idle power-down PASSx_SARy_CTL.IDLE_PWRDWN is enabled.

Figure 35-3. ADC Conversion Timing Flow



35.3.2 Result Data Format

The result after conversion is stored in the lower 16 bits of a 32-bit register, `PASSx_SARy_CHz_RESULT` (see [Table 35-2](#)). The upper 16 bits store the mirror bits of certain flags. Depending on the configuration the possible result data formats are shown [Figure 35-1](#).

35.3.2.1 Signed/Unsigned Result

Conversion results can be treated as signed or unsigned. The `PASSx_SARy_CHz_POST_CTL.SIGN_EXT` bit is used to set the format. Unsigned is the default and is effectively a 12-bit value zero-extended. Considering the result, signed can be useful when `VREFH/2` is the virtual analog ground. The 12-bit code for a signal at `VREFH/2` is `0x800`. This means `0x800` is considered 0, any value below `0x800` is considered negative, and values above `0x800` are considered positive. Therefore, when 'Signed' is set, the MSb (bit 11) is inverted and sign extended.

35.3.2.2 Alignment

A 12-bit result can either be right-aligned (default) or left-aligned within the lower 16 bits of the `PASSx_SARy_CHz_RESULT` register. This is configured per channel with the `PASSx_SARy_CHz_POST_CTL.LEFT_ALIGN` bit. This feature is sometimes used for fixed point arithmetic. For example, this allows for a 12-bit conversion results to be compared to a 16-bit result from a conversion with averaging.

This post processing step takes the 16-bit output from the Right shift step. When `PASSx_SARy_CHz_POST_CTL.LEFT_ALIGN` is used the 16-bit data is shifted four bits to the left; this is done with the assumption that only the lower 12 of the 16 bits are used. The output from the step is still 16-bit.

Table 35-1. Result Data Format

Alignment	Signed/ Unsigned	Result Register															
		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Right	Unsigned	-	-	-	-	11	10	9	8	7	6	5	4	3	2	1	0
Right	Signed	11	11	11	11	11	10	9	8	7	6	5	4	3	2	1	0
Left	-	11	10	9	8	7	6	5	4	3	2	1	0	-	-	-	-

Table 35-2. Channel Result Register

SARx_CHy_RESULT				
Field	Bits	Access	Default	Description
RESULT	15:0	R	-	Conversion result of the channel. Data is copied from the WORK register after all the enabled channel of the group are sampled.
ABOVE_HI_MIR	28	R	-	Set if the result of range detect was above RANGE_HI or cleared otherwise. Note this is only done for the OUTSIDE_RANGE mode, for all other range detection modes this bit is undefined.
RANGE_IN-TR_MIR	29	R		Mirror bit of INTR_CH_RANGE bit
PULSE_IN-TR_MIR	30	R		Mirror bit of INTR_CH_PULSE bit
VALID_MIR	31	R		Mirror bit of the corresponding bit in RESULT_VALID register

35.3.3 Acquisition/Sample Time

The SAR ADC acquisition consists of two steps. In the first step, the sample window, the analog input signal is sampled on the sampling capacitor in the ADC core and in the second step that voltage value is converted to the corresponding digital code.

To get an accurate conversion the analog input signals need to have sufficient time to charge the sampling capacitor. This time is called the 'sample time'. The right sample time depends on the drive strength of the signal source and the RC delay of the whole signal path, including the chip pin, SARMUX, sample capacitor, and wiring. As a result, the required sample time may be different for each signal.

In this SAR ADC, each channel configuration has its own sample time definition (PASSx_SARy_CHz_SAMPLE_CTL.SAMPLE_TIME). This enables optimizing the sample time for each separate analog signal, which in turn enables optimal use of the ADC resource (and power).

Given the fixed on-chip signal path and maximum allowed current draw from the signal source, the minimum sample time can be calculated. This time needs to be translated to the number of SAR clock cycles. PASSx_SARy_CHz_SAMPLE_CTL.SAMPLE_TIME (Table 35-3) is a 12-bit field and the legal values are [1...4095] (0 will be interpreted as 1). At the SAR clock frequency of say, 20 MHz, this corresponds to a sample window range of 50 ns to 0.2 ms. The recommended minimum sampling time for proper settling of the signal is 412 ns; refer to the device datasheet for the exact minimum sample time requirement. The maximum clock frequency for SAR ADC is 26.7 MHz (80/3 MHz) to achieve the 1 MS/s throughput; the recommended sample time corresponds to ~11 clock cycles at this frequency.

The converter requires 13 to 15 clock cycles to perform successive approximation of sampled voltages and present the conversion results. Basic conversion takes 13 cycles with an extra cycle required if the MSB stretch option is enabled (PASSx_SARy_CTL.MSB_STRETCH bit) and another clock is required if a missing code recovery mode (MCR) is enabled (PASSx_SARy_CTL.HALF_LSB bit). To simplify the use, the SAR sequencer may typically use 15 clocks regardless of options, with unused clocks transparently added to sampling clocks.


Table 35-3. Channel Sample Time Setting Field

SARx_CHy_SAMPLE_CTL				
Field	Field	Field	Field	Field
SAMPLE_TIME	27:16	RW	-	Sample time in ADC clock cycles. Minimum value is 1 (Setting 0 gives same result as 1)

35.4 SARMUX

The SARMUX is the analog multiplexer to route the signal to be converted to the ADC core input. The number of GPIO inputs can be up to 32 but the actual connected pins may vary with the device variants. In addition, it can support up to eight internal signals. The selection of input signal is controlled by the physical address field of each virtual channel `PASSx_SARy_CHz_SAMPLE_CTL.PIN_ADDR`. SARMUX can access and route the analog signal from other SARMUXs through an expansion signal. Figure 35-4 shows a high-level block diagram of the SARMUX.

Figure 35-4. SARMUX Block Diagram



35.4.1 Preconditioning

For functional safety and diagnostics, the SAR sequencer supports preconditioning, which enables broken wire detection by charging or discharging the ADC sampling capacitor before sampling the input signal. The use of this feature is optional and is defined by the `PASSx_SARy_CHz_SAMPLE_CTL.PRECOND_MODE` field (see Table 35-4) of the channel configuration. There are four possible selections:

- OFF – no preconditioning
- VREFL – discharge to VREFL
- VREFH – charge to VREFH
- DIAG – connect to the diagnostic reference output during preconditioning

When DIAG preconditioning is used, the diagnostic reference should be configured to output a reference voltage. Note that for overlap diagnostics, the diagnostic reference needs to be configured to supply an i_{bias} current.

There is one diagnostic reference per ADC with a global (not per channel) configuration; therefore, DIAG preconditioning is mutually exclusive with overlap diagnostics. The duration of preconditioning is configurable through the register `PASSx_SARy_PRECOND_CTL.PRECOND_TIME`. This time is specified in SAR clock cycles.

35.4.2 Overlap Diagnostic

Overlap diagnostics is another functional safety feature. For overlap diagnostics, the diagnostic reference typically sources or sinks a small i_{bias} current. In this case, the diagnostic reference output and the analog input signal are both connected to the ADC sampling capacitor at the same time.

The use of this feature is optional and is defined by the `PASSx_SARy_CHz_SAMPLE_CTL.OVERLAP_DIAG` field of the channel configuration. There are three overlap diagnostics modes:

- OFF – No overlap diagnostics
- HALF – Overlapping diagnostic reference for the first half of the sample window
- FULL – Overlapping diagnostic reference for the full sample window

- FULL – Overlapping diagnostic reference for the full sample window

For FULL, the sample window duration is defined by the `PASSx_SARy_CHz_SAMPLE_CTL.SAMPLE_TIME`.

However, for the HALF overlap diagnostic mode, the `PASSx_SARy_CHz_SAMPLE_CTL.SAMPLE_TIME` defines the duration of only half the sample window.

35.4.3 SARMUX Diagnostics


SARMUX diagnostics is a functional safety feature, used to verify the connection from the selected SARMUX input to ADC sampling capacitor. This is done by connecting only the diagnostic reference output to the selected SARMUX input. Note that this does not disturb the analog input signal (for analog details, see [35.4 SARMUX](#)).

The SARMUX diagnostics mode is a per channel optional feature, which is selected by setting the `PASSx_SARy_CHz_SAMPLE_CTL.OVERLAP_DIAG` field. The diagnostic reference should be configured to provide one of the available reference voltages.

Table 35-4. Preconditioning Mode Selection

SARx_CHy_SAMPLE_CTL				
Field	Bits	Access	Default	Description
PRECOND_-MODE	13:12	RW	-	Select the Preconditioning mode for the channel. 00 - No Preconditioning 01 - Discharge to VREFL 10 - Charge to VREFH 11 - Connect to Diagnostic Reference Output (the diagnostic reference generator must be configured accordingly)

Figure 35-5. SAR Sequencer Block Diagram



35.5 SAR Sequencer

The SAR subsystem is largely autonomous, arbitrating acquisition requests (triggers), performing acquisitions, and optionally post processing results without firmware intervention. This important characteristic enables real-time measurements without loading the CPU. The SAR sequencer optionally generates various interrupts to enable further CPU processing of the results or to handle errors. The SAR sequencer also generates several triggers to enable low-latency handling of a detected error or to have the data picked up by DataWire.

35.5.1 Analog Input Selection

There are up to 32 analog pins connected to the regular SARMUX inputs. The number of analog pins is device-specific (refer to the device datasheet for details). In addition, 10 special analog signals can be selected. Eight are selected through the SARMUX and the remaining two are VREFL and VREFH, which bypass the SARMUX and are directly selected at the ADC core input (see Figure 35-1). These two signals are used for calibration.

One of the eight special analog signal is an on-chip temperature sensor. There is only one temperature sensor,

which is shared between all ADCs. For correct operation, the temperature sensor should not be connected to more than one ADC at any given time.

One of the 42 analog signals can be selected by setting the `PASSx_SARy_CHz_SAMPLE_CTL.PIN_ADDR` field (see Table 35-5) of the respective channel. The selected analog signal will only be connected to the ADC during the sample window. If an acquisition is aborted during the sample window, then it is guaranteed that there will be at least one break-before-make cycle (SAR clock) before the new signal is connected to the ADC.

35.5.2 External Analog Multiplexer

The SAR sequencer supports the use of an external analog mux. This can be used to expand the number of analog signals that can be sampled beyond the number of analog pins. Each channel configuration has its own 3-bit wide external mux select value (`PASSx_SARy_CHz_SAMPLE_CTL.EXT_MUX_SEL` in Table 35-5). This allows up to eight channels to use the same analog input pin with different select values. The three external mux select signals are connected at chip level to GPIO digital outputs.

The per channel external mux enable bit (PASSx_SARy_CHz_SAMPLE_CTL.EXT_MUX_EN) is next to the external mux select. This can be used as a chip select for the external mux select device. It is also connected at chip level to a GPIO digital output. Note this enable is not used as output enable for the GPIO digital output drivers of the external mux select. Additionally, when the

PASSx_SARy_CHz_SAMPLE_CTL.EXT_MUX_EN bit is low, the PASSx_SARy_CHz_SAMPLE_CTL.EXT_MUX_SEL field will be ignored.

Table 35-5. External Analog Multiplexer Select and Enable

SARx_CHy_SAMPLE_CTL				
Field	Bits	Access	Default	Description
EXT_MUX_SEL	10:8	RW	-	External analog multiplexer select bits
EXT_MUX_EN	11	RW	-	External analog mux enable. This can be used as enable (chip select) for the external analog mux (it is not used as enable for the GPIO output driver).

35.5.3 Port Selection

Each ADC is preceded by its own SARMUX, which connects to a distinct set of up to 32 analog pins. This means that ADC1 cannot sample the analog pins connected to ADC2. In some cases, it may be desirable to have one ADC being able to reach all the analog inputs of the chip. To support this use-case the ADC0 channels have an additional PASSx_SARy_CHz_SAMPLE_CTL.PORT_ADDR field.

With this field, ADC0 can be connected to the output of the SARMUXes of the other ADCs. This is done through the 'expansion' bus shown in [Figure 35-4](#).

Note that ADC0 can only use the SARMUX of another ADC if that ADC is disabled (PASSx_SARy_CTL.ADC_EN = 0), while SARMUX for that ADC is enabled (PASSx_SARy_CTL.ENABLED = 1 and PASSx_SARy_CTL.SARMUX_EN = 1).

When ADC0 borrows another SARMUX it may need a longer sample time due to the additional on-chip wiring and connected switches.

35.5.4 Averaging


The SAR sequencer includes basic averaging functionality for every channel. When enabled for a channel the SAR sequencer will do back-to-back acquisitions of the same signal and accumulate the results (after sign extension) in a 20-bit accumulator. This is also referred to as a "first order accumulate and dump" filter.

Averaging is fully configured per channel by the PASSx_SARy_CHz_POST_CTL register (see [Table 35-8](#)). The number of samples averaged is determined by the 8-bit PASSx_SARy_CHz_POST_CTL.AVG_CNT field. The number of samples averaged is AVG_CNT+1, which gives a range of [1...256].

35.5.5 Range Detect

The SAR sequencer supports optional range detection feature. Range detection enables a check against up to two programmable threshold values (see [Table 35-7](#)) without CPU involvement. The result is a fast, fixed latency, response time, which is a critical requirement for some use-cases.

Figure 35-6. Range Detection, Threshold, and Events



Range detection is defined by two 16-bit threshold values and a mode field selecting one of four possible modes. Both the mode (PASSx_SARy_CHz_POST_CTL.RANGE_MODE) and the two thresholds (PASSx_SARy_CHz_RANGE_CTL.RANGE_LO and PASSx_SARy_CHz_RANGE_CTL.RANGE_HI) are configured per channel. The available range detection modes are:

- BELOW_LO ($\text{RESULT} < \text{RANGE\_LO}$)
- INSIDE_RANGE ($\text{RANGE\_LO} \leq \text{RESULT} < \text{RANGE\_HI}$)
- ABOVE_HI ($\text{RANGE\_HI} \leq \text{RESULT}$)
- OUTSIDE_RANGE ($\text{RESULT} < \text{RANGE\_LO}$) || ($\text{RANGE\_HI} \leq \text{RESULT}$)

Range detection uses the 16-bit PASSx_SARy_CHz_RESULT.RESULT from the Left-/Right-Align step. This means that the threshold values need to be in the same format as the PASSx_SARy_CHz_RESULT.RESULT after all the preceding post processing steps (including averaging). The event flag will be set when the range mode condition evaluates to true.

When the event flag is set, the PASSx_SARy_CHz_INTR.CH_RANGE interrupt will be set and a pulse is output on the range violation trigger (see [35.7 Output Triggers and Interrupts](#)).

Note that the range detection results (trigger and interrupt) are in addition to the data result.

However, if pulse detection is also enabled for the channel then neither the range detection results nor the data result will be visible; instead, only the range detect event flag is forwarded to the pulse detect feature.


35.5.6 Pulse Detect

The SAR sequencer supports optional pulse detection. Pulse detection is used to filter the events resulting from range detection. The pulse detection filter counts events to detect 'sufficiently long' high pulses while ignoring 'short enough' low spikes.

The pulse detection filter consists of an 8-bit positive event counter and a 5-bit negative event counter. These event counters decrement and/or reload based on the range detection event. When the event is high it is called a positive event and when low, it is called a negative event. The reload values are per channel pulse detection configuration settings. The positive reload value determines what is considered a 'sufficiently long' high pulse and the negative reload value determines which low spikes are 'short enough' to ignore (which is equivalent to; too long not to ignore).

Note that both the pulse detection filter and averaging are used to filter noise. Only one of these two methods can be used for a channel; these two features are mutually exclusive, and their configuration fields are mapped to the same bits.

Figure 35-7. Pulse Detection



*Overflow Interrupt is generated if the previous pulse detect interrupt is still not cleared

Table 35-6. Channel Sample Control Register

SARx_CHy_SAMPLE_CTL				
Field	Bits	Access	Default	Description
PIN_ADDR	5:0	RW	-	Address of the analog signal (pin) to be sampled by this channel 0..31 - AN0...AN31 32 - Select motor input 33 - Select auxiliary input 34 - AMUXBUSA 35 - AMUXBUSB 36 - Digital power supply (VCCD) 37 - Analog power supply (VDDA) 38 - Bandgap voltage from SRSS 39 - Temperature sensor 40..61 - Reserved 62 - VREFL 63 - VREFH
PORT_ADDR	7:6	RW	-	Select Physical Port. This field is only valid for SAR0 (or ADC0) 00 - SARMUX0 (SAR0 uses its own MUX) 01 - SARMUX1 (SAR0 uses MUX of SAR1) 10 - SARMUX2 (SAR0 uses MUX of SAR2) 11 - SARMUX3 (SAR0 uses MUX of SAR3)

Table 35-7. Channel Range Control Register

SARx_CHy_RANGE_CTL				
Field	Bits	Access	Default	Description
RANGE_LO	15:0	RW	-	Range Detect Low Threshold
RANGE_HI	31:16	RW	-	Range Detect High Threshold

35.5.7 Double Buffer

For each channel the SAR sequencer has two complete sets of registers to hold the acquired data and derived flags. The first set of registers are the working registers (PASSx_SARy_CHz_WORK). The working registers are used to store preliminary results, after post processing, from newly completed channel acquisitions.

The second set of registers are the result registers (PASSx_SARy_CHz_RESULT). When a group scan completes, the contents of the working registers are copied (committed) to the corresponding result registers and the Group Done interrupt is set. A group scan completes when the acquisition for the last channel of the group successfully completes. Also, the bits corresponding to the channels are set in the PASSx_SARy_RESULTy_VALID register.

When the results are copied to the result registers, the working registers are immediately available for the SAR sequencer to start a new group scan (for example, in a continuous trigger). In parallel to the new group scan, the software can process the results of the just completed group

scan. This double buffering maximizes the time that the software has to pick up the results.

Note that software should never use information from the PASSx_SARy_CHz_WORK registers, as that information is not coherent (see [35.5.8 Group Coherency](#)). The PASSx_SARy_CHz_WORK registers are only visible to software to provide the status of a group scan in progress, which may be helpful for debug. The corresponding channel bits are set in the PASSx_SARy_WORK_VALID register as soon as the conversion of a channel is completed and result is stored in the PASSx_SARy_CHz_WORK register.

Note that for both the result and working registers the lower 16 bits contain the data (or pulse-detect counters) and in the upper 16 bits the flags/interrupts are mirrored. Double buffering also ensures that preliminary results from a canceled or restarted group scan are discarded; that is, they are never copied in the result registers and thus are not made available to the software.

Table 35-8. Post Processing Control Register

SARx_CHy_POST_CTL				
Field	Bits	Access	Default	Description
POST_PROC	2:0	RW	-	Post Processing 000 - No Processing 001 - Averaging 010 - Averaging followed by Range Detect 011 - Range Detect 100 - Range Detect followed by Pulse detect 101 - Reserved 110 - Reserved 111 - Reserved
LEFT_ALIGN	6	RW	-	Alignment 0 - Data is right aligned RESULT[11:0], 1 - Data is left aligned RESULT[15:4]

Table 35-8. Post Processing Control Register

SARx_CHy_POST_CTL				
Field	Bits	Access	Default	Description
SIGN_EXT	7	RW	-	Sign Extension of Result 0 - Result is unsigned (0- extended if needed) 1 - Result is signed
AVG_CNT	15:8	RW	-	Either Average Count for Channel or Pulse Positive reload value (if Pulse detection is enabled). Averaging Count for channels that have averaging enabled. A channel will be sampled (AVG_CNT+1) = [1...256] time
SHIFT_R	20:16	RW	-	Either Shift Right (no pulse detection) or Pulse negative reload value (if pulse detection is enabled)
RANGE_-MODE	23:22	RW	-	Range Detection Mode 00 - BELOW_LO (RESULT < RANGE_LO) 01 - INSIDE_RANGE (RANGE_HI > RESULT > RANGE_LO) 10 - ABOVE_HI (RESULT > RANGE_HI) 11 - OUTSIDE_RANGE (RANGE_HI < RESULT or RESULT < RANGE_LO)

35.5.8 Group Coherency

For software, it is important that all the results of a group scan are coherent. Coherent results means that all information for all the channels in one group are guaranteed to have been acquired during the same group scan. The SAR sequencer achieves this coherency by making sure that the copy from PASSx_SARy_CHz_WORK to PASSx_SARy_CHz_RESULT registers, of all the channels in the group, happens on a single clock edge. The information for a group scan includes the following:

- RESULT data
- Data valid flags
- All interrupt flags
- Range detect ABOVE_HI flags
- Pulse detect counters
- Channel done triggers

Note that the range violation trigger is not coherent (however, the range detect interrupt is coherent). The range violation trigger is required to have a low latency in the group trigger; therefore, this trigger is set immediately after the range violation is detected.

35.5.9 Status

The SAR sequencer has several status registers to allow software to observe what it is doing. Most of these registers are only intended for debug purposes. Some status registers can also be used for polling.

The following status registers are available:

- A generic status register (PASSx_SARy_STATUS) that shows:
 - If the ADC is busy or not (BUSY)
 - If not busy why not (PWRUP_BUSY, DBG_PAUSE)

- Or if busy shows with which channel (CUR_CHAN) and the priority (CUR_PRIO) and the preemption (CUR_PREEMPT_TYPE) attributes for that channel
- An averaging status register (PASSx_SARy_AVG_STAT) that shows:
 - Current averaging counter
 - Current value of the accumulator – sum of averaging samples acquired so far
- A register to show which input triggers are currently pending (PASSx_SARy_TR_PEND)
- A group status register (PASSx_SARy_CHz_GRP_STAT) that only gathers copies of bits from other registers (PASSx_SARy_CHz_INTR, PASSx_SARy_TR_PEND)

35.6 Triggering and Scheduling

The automotive SAR sequencer has several specific features required for the automotive market. Most of these unique features are related to how acquisitions are scheduled. For example, this SAR sequencer supports the creation of several signal acquisition groups each with their own trigger and priority, which potentially can preempt each other.

35.6.1 Channel Grouping

All the available channels (up to 32) of the SAR ADC can be grouped. A group consists of several consecutive channels of which only the last channel has the 'PASSx_SARy_CHz_TR_CTL.GROUP_END' flag set (see Table 35-9). The number of channels in a group can be anywhere from one (single channel) to 32. Separate groups can have different number of channels.

The first channel of a group defines which trigger is used for that group.

Notes:

- The first channel, which defines the trigger in the group must be enabled
- The last channel, which defines the end in the group must be enabled

- A channel in the group may be disabled in which case it will be skipped
- A group implicitly ends at the last existing and enabled channel

Table 35-9. Channel Group End Flag Field

SARx_CHy_TR_CTL				
Field	Bit	Access	Default	Description
GROUP_END	11	RW	1	0 - Continue group with next channel 1 - Last channel of a group

35.6.2 Triggers

A trigger for a group will cause the acquisitions, as defined by the configurations of the channels in the group, to be executed. There is one dedicated (one-to-one) trigger input for each channel of each of the ADCs connected to the trigger outputs from corresponding TCPWM.

In addition to the TCPWM triggers there are (4 × Number of SAR instances) generic trigger inputs that are shared between ADC channels. Any five of these generic triggers can be routed to any ADC channel. PASS_SAR_TR_IN_SEL_x register can forward any five of all the available generic triggers (up to 16).

Note that the synchronization of the generic trigger inputs happens at a level before the trigger routing selection. This means that the generic triggers arrive synchronously at each ADC, which is an essential feature that enables the synchronized ADC triggering (and thus lock-step execution) that is required for the motor control.

The trigger for a channel group is selected by the configuration (PASSx_SARy_CHz_TR_CTL.SEL, see [Table 35-10](#)) of the first channel of the group. There are seven possible hardware triggers and a software trigger. The hardware trigger options are:

- TCPWM – one-to-one trigger output from a corresponding TCPWM
- GENERIC0-4 – five generic input triggers routed to this ADC
- CONTINUOUS – this trigger is always high, making the group always triggered or in other words Idle trigger
- OFF – no hardware trigger

A group can be software-triggered by setting the PASSx_SARy_CHz_TR_CMD.START bit. This software trigger can be used even if the group is configured to use a hardware trigger.

Notes:

- Setting the pending bit has priority over clearing, so if the hardware trigger input is still high when a trigger clear is received then the pending bit will remain pending.
- If a new trigger is received while the pending bit is already set, then effectively the new trigger is ignored.
- Only the first channel of a group should ever be triggered.

Table 35-10. Trigger Control Register

SARx_CHy_TR_CTL				
Field	Bits	Access	Default	Description
SEL	2:0	RW	0	Analog-to-digital conversion trigger select for the channel 000 - OFF 001 - TCPWM 010 - GENERIC TRIGGER 0 011 - GENERIC TRIGGER 1 100 - GENERIC TRIGGER 2 101 - GENERIC TRIGGER 3 110 - GENERIC TRIGGER 4 111 - CONTINUOUS

35.6.3 Arbitration, Preemption, and Acquisition Scheduling

When a trigger occurs (pending bit high), then the corresponding group of acquisitions (group scan) needs to be executed. When triggers for multiple groups happen, then arbitration is needed to determine which of the pending group scans will be executed first.

The arbitration of the pending triggers is based on both an explicit and an implicit priority. The explicit priority is set, as a trigger attribute (PASSx_SARy_CHz_TR_CTL.PRIO, see [Table 35-11](#)), by software. There are eight explicit priority levels and priority level 0 is the highest.

The implicit priority is defined by the channel ordering as follows: a pending trigger for a lower channel has a higher priority than a higher channel with the same explicit priority.

When a group scan is ongoing and a new higher priority trigger arrives, then it can cause the preemption of the ongoing lower priority group scan. Whether preemption happens is determined by the scheduler based on the explicit priority level and the trigger preemption type of the ongoing group scan. The trigger preemption type is another trigger attribute (PASSx_SARy_CHz_TR_CTL.PREEMPT_TYPE) set by software.

The trigger preemption type determines both when preemption is allowed and what happens with the

preempted group scan on return – after the preempting group scan is done.

The following four preemption types are available:

- **ABORT_RESUME**
 - Immediately abort the ongoing acquisition and on return resume the group scan starting with the aborted channel.
 - Keep the pending trigger of the aborted group.
- **ABORT_RESTART**
 - Immediately abort the ongoing acquisition and on return restart the group scan from the first channel of the group.
 - Keep the pending trigger of the aborted group.
- **ABORT_CANCEL**
 - Immediately abort the ongoing acquisition and do not return.
 - Clear pending trigger of the aborted group and set the canceled interrupt for the last channel of the aborted group.
- **FINISH_RESUME**
 - Before preempting, complete the ongoing acquisition (including averaging) and on return resume the group scan starting with the next channel.
 - Keep the pending trigger of the aborted group.

[Figure 35-8](#) and [Figure 35-9](#) show the behavior of these preemption types when a high priority group trigger arrives.

Figure 35-8. Preemption Types: A Low Priority Group B Behavior with FINISH_RESUME Preemption Type




Figure 35-9. Preemption Types: A Low Priority Group B Behavior with Different Preemption Types




Table 35-11. Trigger Control Register

SARx_CHy_TR_CTL				
Field	Bits	Access	Default	Description
PRI0	6:4	RW	0	Channel Priority 0 Highest Priority 1 6 7 Lowest Priority
PREEMPT_TYPE	9:8	RW	0	Preemption type of the group 00 - ABORT_CANCEL 01 - ABORT_RESTART 10 - ABORT_RESUME 11 - FINISH_RESUME

35.6.4 Debug Freeze

When enabled, the assertion of the debug freeze trigger prevents the scheduler from starting acquisitions for a new channel. Note that averaging, if started, will complete even if the debug pause trigger is asserted. The SAR ADC system has only one debug freeze trigger. However, there is a separate debug freeze enable for each ADC (PASS_PASS_CTL.DBG_FREEZE_EN).

35.6.5 Auto Idle Power Down

The SAR sequencer can optionally be configured (mask bit in the PASSx_SARy_CHz_INTR_MASK register) to automatically power down the analog when the ADC is idle. When this feature is used, the analog will also automatically power up when a trigger arrives. However, after power-up the analog circuit must settle for some time before it can make accurate acquisitions. The required power-up time needs to be configured by software (PASSx_SARy_CTL.PWRUP_TIME).

35.6.6 Channel Disable/Software Abort

When a group is activated, it is no longer allowed to change the configuration settings of the channels in the group. It is undefined what will happen when this rule is violated. When a group or channel needs to be reconfigured it should be disabled first by clearing PASSx_SARy_CHz_ENABLE.ENABLE. All channels in a group need to be disabled together. The channels in a group should be disabled in order, from first to last. If these rules are violated some undefined output may be produced, but no lock up will occur.

Disabling a channel has the following consequences:

- Immediately abort the acquisition if it happens to be in progress for that channel.
- Clear the pending trigger for the channel (if present).
- Discard preliminary results ('work' flags and data).

35.7 Output Triggers and Interrupts

For each channel, there are two trigger outputs, three channel interrupts, and three group interrupts. In addition, there are two generic output triggers per ADC. Only enabled channels can generate new triggers or interrupts. Interrupts are implemented compliant to the platform rules, which means:

- Each of the interrupts has a corresponding mask bit in the PASSx_SARy_CHz_INTR_MASK register to individually enable or disable that interrupt source.
- Software needs to clear the interrupt by writing a '1' to the corresponding bit in the PASSx_SARy_CHz_INTR register.

All enabled interrupts are consolidated into one interrupt output signal per channel. Note that disabling a channel does not clear already pending triggers or interrupts.

35.7.1 Trigger Outputs

Two trigger outputs can be generated per enabled channel: Channel Done and Range Violation triggers.

35.7.1.1 Channel Done Trigger

This trigger indicates that the data for a channel is available in the result register. This means it is never asserted for a pulse-detect channel. This trigger is intended to be used to trigger a DataWire (DW) channel to pick up the result and copy it to system RAM. There is a one-to-one trigger connection from each ADC channel to a corresponding DW channel. The done trigger can be configured (PASSx_SARy_CHz_TR_CTL.DONE_LEVEL) as a level trigger or a pulse trigger. When triggering the DW, a level trigger is recommended. In this mode, the level trigger will remain asserted until the corresponding data is read; that is, the level trigger is de-asserted as a side effect of reading the data. Level trigger mode also enables channel overflow interrupt detection as described here.

Note that all the channel done triggers of the group only assert when the whole group scan is complete and not immediately after the channel acquisition is complete.

In addition, the done trigger of the last channel of a group can be used as a 'group violation' trigger (PASSx_SARy_CHz_POST_CTL.TR_DONE_GRP_VIO). In this mode, the done trigger is only set if at least one of the channels in the group detect a range violation. If none of the channels in the group have range detection enabled, then this trigger is never set.

35.7.1.2 Range Violation Trigger

This trigger generates a pulse in case the acquisition result for the channel causes a range detect event (see [35.5.5 Range Detect](#)). This trigger is a one-to-one trigger connection from each ADC channel to a corresponding TCPWM channel. This trigger is typically used to 'kill' the TCPWM whenever the ADC acquisition results in a value that is outside the predefined allowable range. Note that range detection will not generate a trigger if pulse detection is also enabled for the channel. Range violation trigger asserts immediately after the channel acquisition is complete; unlike all the other ADC outputs it does not wait until the whole group scan is complete.

In addition to these two triggers, there are two generic triggers routed out to the generic trigger infrastructure. This enables the ADC to trigger another IP, other than DW, on completion of a group conversion.

One common use case for one-to-one trigger connection from ADC channel to a TCPWM channel is LED diagnostics.

In this use case the LED is driven with a pulse generated by a TCPWM and the SAR is used to sample a diagnostic feedback signal from the LED driver to ensure that the LED is operating correctly. If the SAR result is outside a predefined range it needs to immediately stop the TCPWM from driving the LED.

35.7.1.3 Generic Trigger Output

For each SAR ADC, two channel triggers (done trigger or range violation trigger) can be routed to the generic trigger infrastructure. `PASS_SAR_TR_OUT_SEL_x.OUT0_SEL` and `PASS_SAR_TR_OUT_SEL_x.OUT1_SEL` registers can be used to select channel triggers to forward. This enables the use of these triggers outside the dedicated 1-to-1 trigger connections to DW or TCPWM. For this use case, it may be better to configure the channel done trigger as a pulse (two cycles on `CLK_SYS`), to avoid the need for a data read to de-assert the trigger.

35.7.2 Channel Interrupts

Each channel has a dedicated interrupt and associated interrupt registers – `PASSx_SARy_CHz_INTR`, `PASSx_SARy_CHz_INTR_SET`, `PASSx_SARy_CHz_INTR_MASK`, and `PASSx_SARy_CHz_INTR_CLR`. The `PASSx_SARy_CHz_INTR_MASK` register is used to mask the interrupt source register, `PASSx_SARy_CHz_INTR`; only the masked interrupt flags are forwarded to the CPU.

Three types of channel interrupts can be generated for each enabled channel.

35.7.2.1 Range Detect Interrupt

This interrupt is set if the acquisition result for the channel causes a range detect event (see [35.5.5 Range Detect](#)). Note that this interrupt is never set if pulse detection is also enabled for the channel – range and pulse detection interrupts are mutually exclusive.

35.7.2.2 Pulse Detect Interrupt

This interrupt is set if the acquisition result for the channel causes a pulse detection.

35.7.2.3 Channel Overflow Interrupt

The channel overflow interrupt is only set if a new group scan completes while the results from a previous completion have not yet been handled. There are three error situations for the channel that the hardware detects. The overflow interrupt is set when on completion of a group scan one of the following conditions is true:

- the range detect interrupt is enabled and still pending
- the pulse detect interrupt is enabled and still pending
- the channel done trigger is set to LEVEL and still asserted

For the first two cases software should have handled and cleared the interrupts before a new one is set.

Similarly, the channel done trigger should have been cleared by the reading the result. If this is not the case, because DataWire is too slow, then the previous data will be overwritten and thus is lost.

35.7.3 Group Interrupts

These are interrupts that can only be set for the last channel of a group (which must be an enabled channel). There are three group interrupts: Group Done, Group Canceled, and Group Overflow interrupts.

35.7.3.1 Group Done Interrupt

This interrupt is set every time a group scan completes.

35.7.3.2 Group Canceled interrupt

This interrupt can only be set for an enabled group with the `ABORT_CANCEL` preemption type. As explained in [35.6.3 Arbitration, Preemption, and Acquisition Scheduling](#), this interrupt is set when the group scan is aborted due to preemption or if a new trigger arrives but it does not immediately result in starting the corresponding group scan.

35.7.3.3 Group Overflow Interrupt

This interrupt is set when a new group scan completes and the group done interrupt is enabled, and still pending from a previous completion. This is an error situation that occurs when software is too slow to pick up the previous results and clear the group done interrupt.

35.8 Calibration

35.8.1 Analog Calibration

Analog calibration is used to make the actual ADC transfer curve come closer to the Ideal transfer curve. Analog calibration can correct an offset and a gain error – linear errors as shown by the 'Actual curve' as shown in [Figure 35-10](#).

Analog calibration is controlled by the configuration register `PASSx_SARy_ANA_CAL`, which contains an 8-bit analog offset (`PASSx_SARy_ANA_CAL.AOFFSET`) and 5-bit analog gain (`PASSx_SARy_ANA_CAL.AGAIN`) calibration value. Before the ADC is used for acquisitions, these values need to be set to the correct values; that is, the ADC needs to be calibrated. The following steps are recommended to find the optimal analog calibration values¹.

As shown in [Figure 35-10](#), the ideal transfer curve has the following two characteristics:

1. Based on Traveo 1 hardware manual (S6J3300).


- Transition between values 0x000 and 0x001 for VREFL + 0.5LSB input voltage.
- Transition between values 0xFFE and 0xFFF for VREFH – 1.5LSB input voltage.

If this is not the case, then the ADC needs to be calibrated, which can be done by performing the following steps:

- Set the analog gain correction value (PASSx_SARy_ANA_CAL.AGAIN) to '0'.
- Configure a channel to convert VREFL.
- Do several software-triggered acquisitions using different PASSx_SARy_ANA_CAL.AOFFSET values.
- Do this until the AOFFSET value X is found for which the converted value transitions from 0x000 to 0x001.
- Change the channel configuration to convert VREFH.
- Do several software triggered acquisitions using different AOFFSET values.

- Do this until the AOFFSET value Y is found for which the converted value transitions from 0xFFE to 0xFFF.
 - Use averaging when the search approaches the desired target.
- Set AOFFSET to $(X+Y)/2 + 2$.
- Change the channel configuration back to converting VREFL.
- Do several software triggered acquisitions using different AGAIN values.
- Do this until the AGAIN value Z is found for which the converted value transitions from 0x000 to 0x001 (using averaging for the final acquisitions).
- Set AGAIN to Z+1.

Figure 35-11 shows the flow chart to calibrate the ADC.



35.8.2 Alternate Calibration

A potential problem with calibration is that over time the ADC error drifts; for example, due to temperature changes. To counter that problem, periodically re-calibrate the ADC.

Running the calibration algorithm requires experimenting with the global calibration configurations. As a result, all other acquisitions should be paused until the calibration algorithm is finished. Considering the number of acquisitions needed for calibration such a pause is unacceptably long.

To solve that problem, a second set of alternate calibration values are added (PASSx_SARy_ANA_CAL_ALT). Each channel can choose to use the alternate calibration values by setting the PASSx_SARy_CHz_SAMPLE_CTL.ALT_CAL bit. With this in place the periodic recalibration algorithm can quietly run in the background while the main application can keep on running undisturbed using the active calibration values.


35.8.3 Coherent Calibration Update

After the new calibration values are established, the next step is to coherently deploy the new values. Changing the calibration values while an acquisition is in progress will result in undefined results for that acquisition and is therefore not allowed.

Similarly, changing the calibration in the middle of a group scan will lead to incoherent results within that group scan. Due to preemption, it is troublesome to determine if some group scan is still waiting to resume and complete.

The PASSx_SARy_CAL_UPD_CMD.UPDATE bit solves this problem. When this bit is set, the sequencer will wait for the 'right moment' to coherently copy values from the alternate calibration registers to the regular calibration registers. At the same time, the PASSx_SARy_CAL_UPD_CMD.UPDATE bit will also be cleared. The right moment for a coherent calibration update is when the ADC becomes idle, or a 'continuous' triggered group completes. This ensures that all acquisitions within a group scan (even if preempted) are done with the same calibration values.

Figure 35-11. Calibration Flow Diagram



35.9 Temperature Measurement

Traveo II devices are equipped with a built-in temperature sensor to measure the chip temperature. To accurately measure the temperature at runtime, use the reference measurement done during production. This reference data is stored in SFlash along with other calibration data. Refer to the device datasheet for the exact address to read this data. Two types of reference data are stored in SFlash – SORT and CLASS.

SORT data is more accurate because it is measured and taken from the actual die. CLASS data is less accurate (distorted) because it is taken when the silicon is in the package. The temperature measurement is done by the user in the CLASS manner; therefore, CLASS data needs to be fitted using the least square approximation (LSA) algorithm. The CLASS fitting should also be adjusted based on the SORT curvature to compensate the package effect.

35.9.1 Example Measurement Flow

The second order polynomial equation is described as follows.

$$y(x) = a_1 + a_1x + a_2x^2$$

The second order polynomial fitting needs at least three data values to produce three polynomial coefficients. Because CLASS data consists of only two values (25°C and 125°C), one coefficient is taken from the second order polynomial of SORT data based on the assumption that the parabolic curvature of the SORT and CLASS data are the same. It means the quadratic coefficient (a_2) for SORT and CLASS are the same. The location of the parabolic vertex determined by a_1 and parabolic offset by a_0 may be different due to the package effect. Then, the two unknown variables a_1 and a_0 can be found using the two known CLASS data. The V_{bg} data is used to remove the ADC output dependency on the ADC reference voltage. Thus, the user does not have to consider the ADC voltage reference when the algorithm is executed. Figure 35-12 shows the conceptual relation between SORT and CLASS data.

Figure 35-12. SORT and CLASS Relationship



35.9.2 Temperature Sensor Calibration and SFlash Address

The temperature sensor calibration is performed by measuring

- Die temperature using external currents and external ADC
- On-chip diode voltage (VBE) with EPASS ADC
- On-chip bandgap reference (VBG) using EPASS ADC

The calibration is done for

- Two set of supplies (3.3 V and 5 V)
- Three different temperatures (SORT2, SORT3, and CLASS HOT)

Table 35-12. SFLASH DATA SET#0: For Die Temperature

SFlash Address	SFlash Name	Parameter	Ambient Temperature (C)	Description
1700_0654, 1700_0655	EPASS_TEMP_TRIM_TEMP_COLDSORT	TCHIP_S2	~40	On-chip temperature measured using external currents and external ADC
1700_064E, 1700_064F	EPASS_TEMP_TRIM_TEMP_ROOMSORT	TCHIP_S3	~25	
1700_065A, 1700_065B	EPASS_TEMP_TRIM_TEMP_HOTSORT	TCHIP_CHI	~130	

Table 35-13. SFLASH DATA SET#1: VBE and VBG at VDDA = 3.3 V

SFlash Address	SFLASH Name	Parameter	Description
1700_0656, 1700_0657	EPASS_TEMP_TRIM_TEMP_DIODE_COLDSORT	VBE_S2	Temperature sensor diode voltage at COLD using EPASS ADC
1700_0658, 1700_0659	EPASS_TEMP_TRIM_TEMP_VBG_COLDSORT	VBG_S2	Bandgap voltage at COLD using EPASS ADC
1700_0650, 1700_0651	EPASS_TEMP_TRIM_TEMP_DIODE_ROOMSORT	VBE_S3	Temperature sensor diode voltage at ROOM using EPASS ADC
1700_0652, 1700_0653	EPASS_TEMP_TRIM_TEMP_VBG_ROOMSORT	VBG_S3	Bandgap voltage at ROOM using EPASS ADC
1700_065C, 1700_065D	EPASS_TEMP_TRIM_TEMP_DIODE_HOTSORT	VBE_CHI	Temperature sensor diode voltage at HOT using EPASS ADC
1700_065E, 1700_065F	EPASS_TEMP_TRIM_TEMP_VBG_HOTSORT	VBG_CHI	Bandgap voltage at HOT using EPASS ADC

Table 35-14. SFLASH DATA SET#1: VBE and VBG at VDDA = 5 V

SFlash Address	SFlash Name	Parameter	Description
1700_066E, 1700_066F	EPASS_TEMP_TRIM_TEMP_DIODE_COLDSORT_5V	VBE_S2_5V	Temperature sensor diode voltage at COLD using EPASS ADC
1700_0670, 1700_0671	EPASS_TEMP_TRIM_TEMP_VBG_COLDSORT_5V	VBG_S2_5V	Bandgap voltage at COLD using EPASS ADC
1700_066A, 1700_066B	EPASS_TEMP_TRIM_TEMP_DIODE_ROOMSORT_5V	VBE_S3_5V	Temperature sensor diode voltage at ROOM using EPASS ADC
1700_066C, 1700_066D	EPASS_TEMP_TRIM_TEMP_VBG_ROOMSORT_5V	VBG_S3_5V	Bandgap voltage at ROOM using EPASS ADC
1700_0672, 1700_0673	EPASS_TEMP_TRIM_TEMP_DIODE_HOTSORT_5V	VBE_CHI_5V	Temperature sensor diode voltage at HOT using EPASS ADC
1700_0674, 1700_0675	EPASS_TEMP_TRIM_TEMP_VBG_HOTSORT_5V	VBG_CHI_5V	Bandgap voltage at HOT using EPASS ADC

35.9.3 Temperature Calculation

The die temperature can be calculated using the on-chip temperature sensor and EPASS ADC, along with the calibration data stored in the SFlash.

SFlash data entries are classified by the values of analog supply VDDA and I/O supply to ADC VDDIO. The values of the supply are 3.3 V and 5 V. Based on the application, the respective data needs to be fetched from the SFlash. The following three cases should be considered.

Table 35-15. Use Case Table

Case	Condition	SFlash Address	When to use this data
1	VDDA = VDDIO = 3.3 V ± 10%	SET#0 and SET#1	3 V ≤ VDDA ≤ 3.6 V and 3 V ≤ VDDIO ≤ 3.6 V
2	VDDA = VDDIO = 5 V ± 10%	SET#0 and SET#2	4.5 V ≤ VDDA ≤ 5.5 V and 4.5 V ≤ VDDIO ≤ 5.5 V
3	VDDA = VDDIO = 2.7 V ~ 5.5 V	SET#0 and SET#1	VDDA ≠ VDDIO

35.9.3.1 Procedure to Calculate the Temperature

Table 35-16. Example SORT and CLASS Data Format in SFlash

Step	Temp ADC Reading [Code]	VBG ADC Reading [Code]	Temp [°C]
SORT RT (ROOM)	554	737	25
SORT HT (HOT)	357	737	150
SORT LT (COLD)	649	737	-40
CLASS RT (ROOM)	554	737	25
CLASS HT (HOT)	398	737	125

The procedure to generate a second order polynomial fitting to predict the temperature is as follows:

1. Read SORT data from SFlash. Extract the data for -40 °C, 25 °C, and 150 °C.

Note: While reading from SFlash, the data must be fetched according to the instructions per case in [Table 35-16](#).

2. From the SORT data, construct A, x, and b matrices.

$$A = \begin{bmatrix} 1 & T_1 & T_1^2 \\ 1 & T_2 & T_2^2 \\ 1 & T_3 & T_3^2 \end{bmatrix} = \begin{bmatrix} 1 & -40 & 1600 \\ 1 & 25 & 625 \\ 1 & 150 & 22500 \end{bmatrix}, x = \begin{bmatrix} a_{0s} \\ a_{1s} \\ a_{2s} \end{bmatrix}, b = \begin{bmatrix} D_1 \\ D_2 \\ D_3 \end{bmatrix} = \begin{bmatrix} 649 \\ 554 \\ 357 \end{bmatrix}$$

3. Solve the Ax = b equation to get the polynomial coefficients

$$D_{out} = a_{2s}T^2 + a_{1s}T + a_{0s}.$$

$$\begin{bmatrix} 1 & -40 & 1600 \\ 1 & 25 & 625 \\ 1 & 150 & 22500 \end{bmatrix} \begin{bmatrix} a_{0s} \\ a_{1s} \\ a_{2s} \end{bmatrix} = \begin{bmatrix} 649 \\ 554 \\ 357 \end{bmatrix}$$

$$\begin{bmatrix} a_{0s} \\ a_{1s} \\ a_{2s} \end{bmatrix} = \begin{bmatrix} 591.1858 \\ -1.4720 \\ -6.1658 \times 10^{-4} \end{bmatrix}$$

Thus, the second order polynomial fitting for SORT data is given by:

$$D_{out} = -6.1658 \times 10^{-4}T^2 + -1.4720T + 591.1858$$

4. Read the CLASS data from the SFlash. Extract the data for 25 °C and 125 °C.
5. Assume $a_{2c} = a_{2s}$, because the curvature of the parabola is the same.

6. From the CLASS data, construct A, x, and b matrices. Since $a_{2c} = a_{2s}$, the b elements must be subtracted from its correspondent $a_{2c}T^2$.

$$A = \begin{bmatrix} 1 & T_1 \\ 1 & T_2 \end{bmatrix} = \begin{bmatrix} 1 & 25 \\ 1 & 125 \end{bmatrix}, x = \begin{bmatrix} a_{0c} \\ a_{1c} \end{bmatrix}, b = \begin{bmatrix} D_1 \\ D_2 \end{bmatrix} = \begin{bmatrix} 554 - (-6.1658 \times 10^{-4})25^2 \\ 398 - (-6.1658 \times 10^{-4})125^2 \end{bmatrix}$$

7. Solve the $Ax = b$ equation to complete the polynomial coefficients of $D_{out} = a_{2c}T^2 + a_{1c}T + a_{0c}$.

$$\begin{bmatrix} 1 & 25 \\ 1 & 125 \end{bmatrix} \begin{bmatrix} a_{0c} \\ a_{1c} \end{bmatrix} = \begin{bmatrix} 554.385 \\ 407.634 \end{bmatrix}$$

$$\begin{bmatrix} a_{0c} \\ a_{1c} \end{bmatrix} = \begin{bmatrix} 591.073 \\ -1.468 \end{bmatrix}$$

From this step, the second order polynomial fitting for CLASS data is given by:

$$D_{out} = -6.1658 \times 10^{-4}T^2 + -1.468T + 591.073$$

8. Read the V_{bg} value from SORT data.
9. Measure the temperature using the ADC. Read the ADC data ($data_{raw}$) and use this value as an input to find the predicted temperature. For example, let $data_{raw}$ be 500.
10. Measure the V_{bg} using the ADC. Read the ADC data ($V_{bg\_measured}$) and use this value as an input to find the correction factor term. For example, let $V_{bg\_measured}$ be 737.
11. Calculate the correction factor (k) to remove dependency of the ADC on the reference voltage.

$$k = \frac{V_{bg\_fromSFlash}}{V_{bg\_measured}} = \frac{737}{737} = 1$$

12. Scale the $data_{raw}$ with k

$$D = k \times data_{raw} = 1(500) = 500$$

13. Find the predicted temperature from the second order polynomial fitting CLASS data from D

$$\text{PredictedTemperature} = \frac{-b - \sqrt{b^2 - 4a(c - D)}}{2a}$$

Where,

$$a = -6.1658 \times 10^{-4}, b = -1.468, c = 591.073, D = 500$$

For this example, predicted temperature = 60.5 °C


14. Find the nearest calibration temperature to the temperature calculated in step 13. If the temperature calculated in step 13 is close to
- COLD (−40 °C), repeat steps 11 and 12 with VBG_S2_5V and recalculate the temperature using step 13.
 - HOT (150 °C), repeat steps 11 and 12 with VBG_CHI and recalculate the temperature using step 13.
 - ROOM (27 °C), temperature calculated in step 13 is the final temperature.

35.10 Diagnostic Reference Generator

This block provides voltages and currents for chip level use to test internal and external signal connections. Injecting voltages near pads allows testing the continuity of internal routing to the ADC. Provisions also exist to connect reference voltages and currents to pads whether independently or simultaneously with ADC inputs.

It generates selectable output voltages. It also includes the option to add 10- μ A sink or source currents. The voltages can be derived from the supply (VDDA) or can be a buffered version of the bandgap (VBG). The 10- μ A sink/source currents used by the diagnostic reference are generated by the reference buffer described in the next section.

Figure 35-13. Diagnostic Reference Block Diagram



35.10.1 Diagnostic Reference Configuration

Each ADC has one diagnostic reference block (see [Figure 35-13](#)), which consists of:

- An RDAC providing 10 voltage levels from VREFL to VREFH
- Four other voltage references
- A current source and sink function
- An analog mux to select one of these signals

The diagnostic reference block is configured through the DIAG_CTL register. This is a global (per ADC) configuration.

When the diagnostic reference block is not used by any channel it should be disabled, to save power, by clearing the DIAG_EN bit.

- Providing a temperature dependent voltage for on-die temperature sensing.

The buffered bandgap voltage connects to SAR and diagnostic reference generator multiplexer inputs. The current source/sinks are used by the diagnostic reference generators for broken wire detection. The temperature sensor should be allowed to settle for at least 1 μ sec before converting (a sample time of 1 μ s is sufficient) and several measurements with averaging are recommended.

Only one ADC at a time should perform temperature measurements as sampling of another ADC may disturb the temperature sensor output voltage.

The power supply monitoring portion of the reference buffer provides termination, which forms voltage dividers between power supply voltages routed on the AMUXBUSA/B signals and the ADC. Resistors and switches contained in power and ground pins connect these supplies to the AMUXBUSA/B as controlled by the IOSS (HSIOM_MONITOR_CTL register), with the termination of the signals controlled by the reference buffer power supply monitor block. The midpoint of the signal (AMUXBUSA/B) is connected to the SARMUX (internal signals) and can be selected for analog-to-digital conversion by a channel.

35.11 Reference Buffer

The reference buffer contains control logic, a voltage follower amplifier (the same amplifier circuit used in the SRSS to buffer Vbg), a current multiplier and array of current sources and sinks, and a temperature sensor. Additionally, it also provides power monitoring block. In principle it provides four functions:

- Power Supply Monitoring
- Buffering the 0.9-V bandgap signal from the SRSS.
- Scaling 1 μ A (4x 250 nA currents in parallel) from the SRSS to 10 μ A and replicating this current (both source and sink) for diagnostic reference generators.

Figure 35-14. Power Monitoring Block Diagram




Table 35-17. PASS Control register (PASS_CTL)

Field	Bits	Access	Default	Description
SUPPLY_MON_EN_A	0	RW	0	Supply monitor enable for AMUXBUS_A
SUPPLY_MON_LVL_A	1	RW	0	Supply monitor level select for AMUXBUS_A 0 - VRL 1 - VRH
SUPPLY_MON_EN_B	4	RW	0	Supply monitor enable for AMUXBUS_B
SUPPLY_MON_LVL_B	5	RW	0	Supply monitor level select for AMUXBUS_B 0 - VRL 1 - VRH
REFBUF_MODE	22:21	RW	0	The reference needs to be present when using TEMP sensor or diagnostic reference (in addition to SAR.DIAG_CTL.DIAG_EN). Note: Setting this mode is not required for the ADC operation itself. 00 - OFF - No Reference (Disabled) 01 - ON - Reference enabled with buffered Vbg from SRSS. 10 - Reserved 11 - BYPASS - Reference enabled with unbuffered Vbg from SRSS

35.12 Registers

Symbol	Name	Description
PASSx_SARy_CTL	Analog control register	This register controls the power and configuration of SAR ADC instance.
PASSx_SARy_DIAG_CTL	Diagnostic reference control register	This register configures the diagnostic reference generator.
PASSx_SARy_PRECOND_CTL	Preconditioning control register	This register set the time for precondition. The value is set in number of clock cycles.
PASSx_SARy_ANA_CAL	Current analog calibration values	This register stores the value of offset and gain for the ADC core.
PASSx_SARy_DIG_CAL	Current digital calibration values	This register stores the value of digital offset and gain.
PASSx_SARy_ANA_CAL_ALT	Alternate analog calibration values	This register stores the offset and gain; it enables the background calibration.
PASSx_SARy_DIG_CAL_ALT	Alternate digital calibration values	It stores the digital offset and gain and enables the background calibration.
PASSx_SARy_CAL_UPD_CMD	Calibration update command register	Coherently updates the calibration registers with the value from alternate calibration register.
PASSx_SARy_TR_PEND	Channel trigger pending status register	Channel trigger pending status bit is set when the trigger is received for the channel.
PASSx_SARy_WORK_VALID	WORK data valid flag register	This is set when the data in the WORK register of channel is valid or the conversion is successfully completed.
PASSx_SARy_WORK_RANGE	WORK range detect flag register	Channel range detect flag register.
PASSx_SARy_WORK_RANGE_HI	WORK outside range detect flag register	This bit is set when the range violation detected in OUTSIDE_RANGE mode.
PASSx_SARy_WORK_PULSE	Channel pulse detect	Pulse detect flag register.
PASSx_SARy_RESULT_VALID	Channel result data register 'valid' bits	Channel RESULT register data valid flags.
PASSx_SARy_RESULT_RANGE_HI	Channel range above Hi flags	This bit is set when the range violation detected in OUTSIDE_RANGE mode.
PASSx_SARy_STATUS	SAR status register	Reads the status of SAR and the currently scanning channel.
PASSx_SARy_AVG_STAT	Current averaging status	Reads the current value of the accumulator and counter
PASSx_SARy_CHz_TR_CTL	Channel trigger control register	Sets the channel triggers, priority, preempt, and group related configurations.
PASSx_SARy_CHz_SAMPLE_CTL	Channel sample control register	Configure the analog sampling related configurations such as physical pin address, sample time, and preconditioning.
PASSx_SARy_CHz_POST_CTL	Channel post processing control register	Configures the post processing of the converted analog value such as averaging, alignment, and range mode.
PASSx_SARy_CHz_RANGE_CTL	Channel range threshold register	Stores the low and high range thresholds for range detection of the channel.
PASSx_SARy_CHz_INTR	Channel interrupt request register	Channel interrupt request register clears the interrupt request by writing '1'.
PASSx_SARy_CHz_INTR_SET	Channel interrupt set request register	Sets the INTR by writing '1'.
PASSx_SARy_CHz_INTR_MASK	Channel interrupt mask register	Channel interrupt masking register.
PASSx_SARy_CHz_INTR_MASKED	Channel interrupt masked request register	INTR register after the masking.
PASSx_SARy_CHz_WORK	Channel working data register	Stores the conversion result as soon as it is completed for the channel
PASSx_SARy_CHz_RESULT	Channel result data register	Data is copied from the work register after all the channels in the current group are sampled.

Symbol	Name	Description
PASSx_SARy_CHz_GRP_STAT	Channel group status register	Reads the status of the current scanning group.
PASSx_SARy_CHz_ENABLE	Channel enable register	Channel enable/disable, resets trigger, and valid flags immediately if disabled.
PASSx_SARy_CHz_TR_CMD	Channel software trigger	Channel software trigger, triggered by setting '1'. Always reads '0'.
PASS_PASS_CTL	PASS control register	Debugs freeze control of all the ADC instances and reference buffer control.
PASS_SAR_TR_IN_SEL_x	Generic input trigger select register	Selection of five generic input triggers for the ADC.
PASS_SAR_TR_OUT_SEL_x	Generic output trigger select register	Selection of output trigger for the two generic output triggers.

Note: In PASSx_SARy_CHz, 'x' signifies the PASS instance, 'y' signifies the SAR instance, and 'z' signifies the SAR channel. Refer to the device datasheet for the specifications.


Section G: Program and Debug Overview



This section encompasses the following chapters:

- Program and Debug Interface chapter on page 684
- Nonvolatile Memory Programming chapter on page 694
- Flash Boot chapter on page 740

Program and Debug Architecture



36. Program and Debug Interface



The TVII-B-H program and debug interface provides a communication gateway for an external device to perform programming and debugging. The external device can be a Cypress-supplied programmer and debugger, or a third-party device that supports programming and debugging. The serial wire debug (SWD) or the JTAG interface can be used as the communication protocol between the external device and the TVII-B-H device.

36.1 Features

The TVII-B-H program and debug interface has the following features:

- Supports programming and debugging through the JTAG or SWD interface
- CM7 supports ETM/ITM tracing through a 4-bit TPIU, embedded trace buffer (ETB) with 8KB dedicated RAM. CM0+ supports micro trace buffer (MTB) with 4KB dedicated RAM.
- Supports cross-trigger interface (CTI) and cross-trigger matrix (CTM)
- CM0+ supports four hardware breakpoints and two watchpoints. CM7 supports six hardware breakpoints and four watchpoints
- Provides read and write access to all memory and registers in the system while debugging, including the Cortex-M7 and Cortex-M0+ register banks when the core is running or halted

36.2 Functional Description

The debug and access port (DAP) acts as the program and debug interface. The external programmer or debugger, also known as the “host”, communicates with the DAP of the TVII-B-H “target” using either the SWD or JTAG interface. The debug physical port pins communicate with the DAP through the high-speed I/O matrix (HSIOM). See the [I/O System chapter on page 246](#) for details on HSIOM.

The debug infrastructure is organized into the following four groups:

- DAP (provides pin interfaces through which the debug host can connect to the chip)
- Cortex-M0+ core debug components
- Cortex-M7 core debug components, tightly related to the CPU and running at the same frequency
- Other debug infrastructure (includes the low-speed section of the CM7 tracing, the cross-triggering matrix, and the system ROM table)

The debug and trace infrastructure is built mostly using the library of CoreSight compliant components from Arm CoreSight-400 IP. The following are the various debug and trace components.

The following are the various debug and trace components:

- Debug components
 - JTAG and SWD for debug control and access
- Trace source components
 - Micro trace buffer (MTB-M0+) to trace Cortex-M0+ program execution
 - Embedded trace macro (ETM-M7) to trace Cortex-M7 program execution
- Trace sink components
 - ETB for on-chip storage of the trace

- Trace port interface unit (CoreSight-TPIU) to drive the trace information out of the device to an external trace port analyzer (TPA)
- Trace link components
 - ATB (AMBA trace bus) funnel
 - ATB replicator
- Cross-triggering components
 - CTI
 - CTM
- ROM tables

36.2.1 Debug Access Port (DAP)

The DAP consists of a combined SWD/JTAG interface (SWJ) that also includes the SWD listener. The SWD listener decides if the JTAG interface (default) or SWD interface is active. Note that JTAG and SWD are mutually exclusive because they share pins.

The debug port (DP) connects to the DAP bus, which in turn connects to one of three access ports (AP), namely:

- The CM0-AP, which connects directly to the AHB debug slave port (SLV) of the CM0+ and gives access to the CM0+ internal debug components. This also allows access to the rest of the system through the CM0+ AHB master interface. This provides the debug host the same view as an application running on the CM0+. This includes access to the MMIO of other debug components of the Cortex M0+ subsystem. These debug components can also be accessed by the CM0+ CPU, but cannot be reached through the other APs or by the CM7 core.
- The CM7_0-AP/CM7_1-AP, which connects directly to AHB-Lite debug slave port of CM7_0/CM7_1 (DBG) through an AHB2AHB bridge gives access to the CM7_0/CM7_1 internal debug components. The CM7_0-AP/CM7_1-AP also allows access to the rest of the system through the CM7_0/CM7_1 AHB-Lite master interfaces. This provides the debug host the same view as an application running on the CM7_0/CM7_1 core. Additionally, the CM7_0-AP/CM7_1-AP provides access to the debug components in the CM7_0/CM7_1 core through the external peripheral bus (EPB). These debug components can also be accessed by the CM7_0/CM7_1 CPU, but cannot be reached through the other APs or by the CM0+ core.
- The system-AP gives access to the rest of the system through an AHB mux. This allows access to the system ROM table, which cannot be reached any other way. The system ROM table provides the chip ID but is otherwise empty.

36.2.1.1 DAP Security

For security reasons all three APs each can be independently disabled. Each AP disable is controlled by

two MMIO bits. The DAP_CTL.xxx_AP_DISABLE bit (where xxx can be CM0 or CM7_0 or CM7_1 or SYS), can be set during boot, before the debugger can connect, based on eFuse settings. After this bit is set it cannot be cleared.

The second bit, DAP_CTL.xxx_AP_ENABLE (where xxx can be CM0 or CM7_0/CM7_1 or SYS), is a regular read/write bit. This bit also resets to zero and is set to '1' by either the ROM boot code or the flash boot code depending on the life-cycle stage. This feature can be used to block debug access during normal operation, but re-enable some debug access after a successful authentication.

In addition, the system AP is also protected by an MPU. This can be used to give the debugger limited access to the rest of the system. For chip identification, access to the system ROM table should be allowed. If debug access is restored after successful authentication, this MPU needs to be configured to allow authentication requests.

Note: The debug slave interfaces of all the CPUs bypass the internal CPU MPU.

36.2.1.2 DAP Power Domain

Almost all the debug components are part of the Active power domain. The only exception is the SWD/JTAG-DP, which is part of the DeepSleep power domain. This allows the debug host to connect during DeepSleep, while the application is 'running' or powered down. This enables in-field debugging for low-power applications in which the chip is mostly in DeepSleep.

After the debugger is connected to the device, it needs to bring the device to the Active state before any operation. For this, the SWD/JTAG-DP has a register (CTRL/STAT) with two power request bits. The two bits, CDBGPWRUPREQ and CSYSPWRUPREQ, request for debug power and system power respectively. These bits need to remain set for the duration of the debug session.

Note that only the two SWD pins (SWCLKTCK and SWDIOTMS) are operational during the DeepSleep mode - the JTAG pins are only operational in Active mode. The JTAG debug and JTAG boundary scan are not available when the system is in DeepSleep. JTAG functionality is only available after a device power-on-reset.

A system reset (XRES_L pin or AIRCR.SYSRESETREQ) will reset the I/O configuration and cause the host connection to be lost.

36.2.2 ROM Tables

The ROM tables are organized in a tree hierarchy. Each AP has a register that contains a 32-bit address pointer to the base of the root ROM table for that AP. TVII-B-H has four such root ROM tables.

Each ROM table contains 32-bit entries with an address pointer that either points to the base of the next level ROM

table or a leaf debug component. Each ROM table also contains a set of ID registers that hold JEDEC compliant identifiers to identify the manufacturer, part number, and major and minor revision numbers. These IDs are the same for all ROM tables in TVII-B-H. Each ROM table and CoreSight compliant component also contain component identification registers.

36.2.3 Trace

TVII-B-H supports instruction tracing for all the CPUs (design time configurable). The CM7 also includes an optional instrumentation trace macrocell (ITM). The CM0+ trace can only be captured in the MTB; that is, CM0+ tracing is not connected to the trace infrastructure.

The CM7 trace uses the standard AMBA trace bus (ATB) for both ITM and ETM output. The trace streams from the two CM7 CPUs are first combined using ATB funnel components to create a single trace stream. This trace output is then replicated, which enables the traces to go to two trace sink components, namely the trace port interface unit (TPIU) and the embedded trace buffer (ETB).

The ETB is similar to the MTB and captures the trace information in an on-chip SRAM, which can be retrieved by a debugger. The TPIU brings the trace off-chip through asynchronous interface with up to four pins. The trace sources live in the CM7 clock domain and have a width of 8-bit.

36.2.4 Embedded Cross-Triggering

The Arm CoreSight includes embedded cross-triggering (ECT) to communicate events between debug components. These events are particularly useful with tracing and multicore platforms. For example, trigger events can be used to:

- Start or stop both CPUs at (almost) the same time
- Start or stop instruction tracing based on trace buffer being full or not or based on other events

CoreSight uses two components to support ECT, namely a CTI and a CTM, both of which are used in TVII-B-H.

The CTI component interfaces with other debug components, sending triggers back and forth and synchronizing them as needed. The CTM connects several CTIs, thus allowing events to be communicated from one CTI to another.

The TVII-B-H has four CTIs, one for each CPU and one for the trace components in the debug structure. These three CTIs are connected together through the CTM. The CM7_0/CM7_1 CTI is located in the fast clock domain and the other two CTIs and the CTM are all located in the same slow-frequency clock domain. The list of the triggers connected to each CTI are as follows:

CM0+ CTI

- Input triggers:
 - 0 = cm0p.halted // CM0+ is in debug mode
- Output triggers:
 - cm0p.edbgr // CM0+ to enter debug mode
 - 2 = sys.cm0_cti_irq[0] // Interrupt request
 - 3 = sys.cm0_cti_irq[1] // Interrupt request
 - 4 = mtb.tstart // Request MTB to start tracing
 - 5 = mtb.tstop // Request MTB to stop tracing
 - 7 = cm0p.dbgrestart // Request CM0+ to exit debug mode

CM7 CTI

- Input triggers:
 - 0 = cm7.halted // CM7 entered debug mode
 - 1 = cm7.dwtmatch[0] || cm7.dwtmatch[4] // CM7 DWT comparator outputs
 - 2 = cm7.dwtmatch[1] || cm7.dwtmatch[5] // CM7 DWT comparator outputs
 - 3 = cm7.dwtmatch[2] || cm7.dwtmatch[6] // CM7 DWT comparator outputs
 - 4 = etm.eventm[0] // ETM event output
 - 5 = etm.eventm[1] // ETM event output
 - 6 = etm.eventm[2] // ETM event output
 - 7 = etm.eventm[3] // ETM event output
- Output triggers:
 - 0 = cm7.edbgrq // Request CM7 to enter debug mode
 - 1 = sys.cm7_cti_irq[0] // interrupt request TBD which system interrupts?
 - 2 = sys.cm7_cti_irq[1] // interrupt request
 - 3 = etm.events[0] // ETM event input
 - 4 = etm.events[1] // ETM event input
 - 5 = etm.events[2] // ETM event input
 - 6 = etm.events[3] // ETM event input
 - 7 = cm7.dbgrestart // Request CM7 to exit debug mode

TRC CTI

- Input triggers:
 - 0 = cm0p.halted (level) // CM0+ is in debug mode (level)
 - 2 = etb_full // Flag that ETB is full
 - 3 = etb_acqomp complete // Flag trace acquisition complete
 - 4 = cm7_0.halted (level) // CM7_0 is in debug mode (level)
 - 5 = cm7_1.halted (level) // CM7_1 is in debug mode (level)

- ❑ 6 = sys.tr_cti_in[0] // CTI trigger input from system triggers
- ❑ 7 = sys.tr_cti_in[1] // CTI trigger input from system triggers
- Output triggers:
 - ❑ 0 = etb_flushin // Request ETB to flush
 - ❑ 1 = etb_trigin // Request ETB to stop tracing
 - ❑ 2 = tpiu_flushin // Request TPIU to flush
 - ❑ 3 = tpiu_trigin // Request TPIU to stop tracing
 - ❑ 6 = sys.tr_cti_out[0] // CTI trigger output to system triggers
 - ❑ 7 = sys.tr_cti_out[1] // CTI trigger output to system triggers

Note that CoreSight cross-triggering is mostly separate from the peripheral trigger multiplexer. The only connection between the two are the four sys.tr_cti_in/out signals mentioned above.

Note: The CTI registers are only accessible if a debugger is connected. Any code that needs to access CTI registers should first ensure the presence of a debugger; for example, by using the CPUSS_DP_STATUS register.

36.3 Serial Wire Debug (SWD) Interface


The TVII-B-H device supports programming and debugging through the SWD interface. The SWD protocol is a packet-based serial transaction protocol. At the pin level, it uses a single bidirectional data signal (SWDIO) and a unidirectional clock signal (SWDCK). The host programmer always drives the clock line, whereas either the host or the target drives the data line. A complete data transfer (one SWD packet) requires 46 clocks and consists of three phases:

- **Host Packet Request Phase** – The host issues a request to the TVII-B-H target.
- **Target Acknowledge Response Phase** – The TVII-B-H target sends an acknowledgment to the host.
- **Data Transfer Phase** – The host or target writes data to the bus, depending on the direction of the transfer.

When control of the SWDIO line passes from the host to the target, or vice versa, there is a turnaround period (Trn) where neither device drives the line and it floats in a high-impedance (Hi-Z) state. This period is either one-half or one and a half clock cycles, depending on the transition.

Figure 36-1 shows the timing diagrams of read and write SWD packets.

Figure 36-1. SWD Write and Read Packet Timing Diagrams



The sequence to transmit SWD read and write packets are as follows:

1. Host Packet Request Phase: SWDIO driven by the host
 - a. The start bit initiates a transfer; it is always logic 1.
 - b. The “AP not DP” (APnDP) bit determines whether the transfer is an AP access - 1b1 or a DP access - 1b0.
 - c. The “Read not Write” bit (RnW) controls which direction the data transfer is in. 1b1 represents a 'read from' the target, or 1b0 for a 'write to' the target.
 - d. The address bits (A[2:3]) are register select bits for AP or DP, depending on the APnDP bit value.
Note: Address bits are transmitted with the LSb first.
 - e. The parity bit contains the parity of APnDP, RnW, and A[2:3] bits. It is an even parity bit; this means, when XORed with the other bits, the result will be 0. If the parity bit is not correct, the header is ignored by TVII-B-H; there is no ACK response (ACK = 3b111). The programming operation should be aborted and retried again by following a device reset.
 - f. The stop bit is always logic 0.
 - g. The park bit is always logic 1.
2. Target Acknowledge Response Phase: SWDIO driven by the target
 - a. The ACK[0:2] bits represent the target to host response, indicating failure or success, among other results. See [Table 36-2](#) for definitions.
Note: ACK bits are transmitted with the LSb first.
3. Data Transfer Phase: SWDIO driven by either target or host depending on direction
 - a. The data for read or write is written to the bus, LSb first.
 - b. The data parity bit indicates the parity of the data read or written. It is an even parity; this means when XORed with the data bits, the result will be 0. If the parity bit indicates a data error, corrective action should be taken. For a read packet, if the host detects a parity error, it must abort the programming operation and restart. For a write packet, if the target detects a parity error, it generates a FAULT ACK response in the next packet.

According to the SWD protocol, the host can generate any number of SWDCK clock cycles between two packets with SWDIO low. It is recommended to generate three or more dummy clock cycles between two SWD packets if the clock is not free-running or to make the clock free-running in IDLE mode.

The SWD interface can be reset by clocking the SWDCK line for 50 or more cycles with SWDIO high followed by at least two idle cycles.

36.3.1 SWD Timing Details

The SWDIO line is written to and read at different times depending on the direction of communication. The host drives the SWDIO line during the host packet request phase and, if the host is writing data to the target, during the data transfer phase as well. When the host is driving the SWDIO line, each new bit is written by the host on falling SWDCK edges, and read by the target on rising SWDCK edges. The target drives the SWDIO line during the target acknowledge response phase and, if the target is reading out data, during the data transfer phase as well. When the target is driving the SWDIO line, each new bit is written by the target on rising SWDCK edges, and read by the host on falling SWDCK edges.

[Table 36-1](#) and [Figure 36-1](#) illustrate the timing of SWDIO bit writes and reads.

Table 36-1. SWDIO Bit Write and Read Timing

SWD Packet Phase	SWDIO Edge	
	Falling	Rising
Host Packet Request	Host Write	Target Read
Host Data Transfer		
Target ACK Response	Host Read	Target Write
Target Data Transfer		

36.3.2 ACK Details

The acknowledge (ACK) bit-field is used to communicate the status of the previous transfer. OK ACK means that previous packet was successful. For a FAULT status, the programming operation should be aborted immediately. [Table 36-2](#) shows the ACK bit-field decoding details.

Table 36-2. SWD Transfer ACK Response Decoding

Response	ACK[2:0]
OK	3b001
WAIT	3b010
FAULT	3b100
NO ACK	3b111

Details on WAIT and FAULT response behaviors are as follows:

- For a WAIT response, if the transaction is a read, the host should ignore the data read in the data phase. The target does not drive the line and the host must not check the parity bit as well.
- For a WAIT response, if the transaction is a write, the data phase is ignored by the TVII-B-H. However, the host must still send the data to be written to complete the packet. The parity bit corresponding to the data should also be sent by the host.
- For a WAIT response, it means that the TVII-B-H is processing the previous transaction. The host can try for

a maximum of four continuous WAIT responses to see if an OK response is received. If it fails, then the programming operation should be aborted and retried.

- For a FAULT response, the programming operation should be aborted and retried by doing a device reset.


36.3.3 Turnaround (Trn) Period Details

There is a turnaround period between the packet request and ACK phases, as well as between the ACK and data phases for host write transfers, as shown in [Figure 36-1](#). According to the SWD protocol, the Trn period is used by both the host and target to change the drive modes on their respective SWDIO lines. During the first Trn period after the packet request, the target starts driving the ACK data on the SWDIO line on the rising edge of SWDCK. This action ensures that the host can read the ACK data on the next falling edge. Thus, the first Trn period lasts only one-half cycle. The second Trn period of the SWD packet is one and a half cycles. Neither the host nor the TVII-B-H device should drive the SWDIO line during the Trn period.

36.4 JTAG Interface


In response to higher pin densities on microcontrollers, the Joint Test Action Group (JTAG) proposed a method to test circuit boards by controlling the pins on the microcontrollers (and reading their values) via a separate test interface. The solution, later formalized as IEEE Standard 1149.1-2001, is based on the concept of a serial shift register routed across all of the pins – hence the name boundary scan. The circuitry at each pin is supplemented with a multipurpose element called a boundary scan cell. In TVII-B-H devices, most GPIO port pins have a boundary scan cell associated with them (see the GPIO block diagrams in the [I/O System chapter on page 246](#)). The interface used to control the values in the boundary scan cells is called the test access port (TAP) and is commonly known as the JTAG interface. It consists of three signals: test data in (TDI), test data out (TDO), and test mode select (TMS). Also included is a clock signal (TCK) that clocks the other signals. TDI, TMS, and TCK are all inputs to the device and TDO is the output from the device. This interface enables testing multiple devices on a circuit board, in a daisy-chain fashion, as shown in [Figure 36-2](#).

Figure 36-2. JTAG Interface to Multiple Devices on a Circuit Board



[Figure 36-3](#) shows the JTAG interface architecture within each device. Data at TDI is shifted in, through one of several available registers, and out to TDO.

Figure 36-3. JTAG Interface Architecture




The TMS signal controls a state machine in the TAP. The state machine controls which register (including the boundary scan path) is in the TDI-to-TDO shift path, as shown in [Figure 36-4](#). The following terms apply:

- **ir** - the instruction register
- **dr** - one of the other registers (including the boundary scan path), as determined by the contents of the instruction register
- **capture** - transfer the contents of a dr to a shift register, to be shifted out on TDO (read the dr)
- **update** - transfer the contents of a shift register, shifted in from TDI, to a dr (write the dr)

Note: Flash boot configures the JTAG reset input pin to SWJ_TRSTN mode upon reset (refer to the related device datasheet for the pin number). The JTAG TRSTn pin is optional for the JTAG debug protocol; the user application can configure it to some other mode (such as GPIO).

If the debug protocol is JTAG and this pin is configured, the debug session will crash. To avoid this issue, modify the JTAG reset input to GPIO within the scope of the debug script.

Figure 36-4. TAP State Machine



The registers in the TAP are:

- **Instruction** – Typically two to four bits wide, holds the current instruction that defines which data register is placed in the TDI-to-TDO shift path.
- **Bypass** – one bit wide, directly connects TDI with TDO, causing the device to be bypassed for JTAG purposes.
- **ID** – 32 bits wide, used to read the JTAG manufacturer/part number ID of the device.
- **Boundary Scan Path (BSR)** – Width equals the number of I/O pins that have boundary scan cells, used to set or read the states of those I/O pins.

Other registers may be included according to the device manufacturer specifications. The standard set of instructions (values that can be shifted into the instruction register), as specified in IEEE 1149, are:

- **EXTEST** – Causes TDI and TDO to be connected to the BSR. The device is changed from its normal operating mode to a test mode. Then, the device's pin states can be sampled using the capture dr JTAG state. New values can be applied to the pins of the device using the update dr state.
- **SAMPLE** – Causes TDI and TDO to be connected to the BSR, but the device is left in its normal operating mode.

During this instruction, the BSR can be read by the capture dr JTAG state to take a sample of the functional data entering and leaving the device.

- **PRELOAD** – Causes TDI and TDO to be connected to the BSR, but device is left in its normal operating mode. The instruction is used to preload test data into the BSR before loading an EXTEST instruction.

Optional, but commonly available, instructions are:

- **IDCODE** – Causes TDI and TDO to be connected to an IDCODE register.
- **INTEST** – Causes TDI and TDO to be connected to the BSR. While the EXTEST instruction allows access to the device pins, INTEST enables similar access to the corelogic signals of a device

For more information, see the IEEE Standard, available at www.ieee.org.

Note: The Boundary Scan TAP is daisy-chained with the CoreSight DAP in the following order: TDI → Boundary Scan TAP (IR length = 4) → CoreSight DAP (IR length = 4) → TDO. To get boundary scan working, an additional BSDL file is needed that sets the CoreSight DAP into BYPASS mode.

36.5 Programming the TVII-B-H Device

TVII-B-H is programmed using the following sequence. Refer to the *Traveo II MCU Programming Specifications* for complete details on the programming algorithm, timing specifications, and hardware configuration required for programming.

1. Acquire the SWD port in TVII-B-H.
2. Enter the programming mode.
3. Execute the device programming routines such as silicon ID check, flash programming, flash verification, and checksum verification.

36.5.1 SWD Port Acquisition

36.5.1.1 SWD Port Acquire Sequence

The default interface on power-on reset is JTAG; to switch to SWD, use a transition through dormant state.

The first step in device programming is for the host to acquire the target's SWD port. The host performs a device reset by asserting XRES_L pin. After removing the XRES_L signal, the host must send an SWD connect sequence for the device within the acquire window to connect to the SWD interface in the DAP.

The debug access port must be reset using the standard Arm command. The DAP reset command consists of more than 49 SWDCK clock cycles with SWDIO asserted high. The transaction must be completed by sending at least one SWDCK clock cycle with SWDIO asserted low. This sequence synchronizes the programmer and the chip. Read_DAP() refers to the read of the IDCODE register in the debug port. The sequence of line reset and IDCODE read should be repeated until an OK ACK is received for the IDCODE read or a timeout (2 ms) occurs. The SWD port is said to be in the acquired state if an OK ACK is received within the time window and the IDCODE read matches with that of the Cortex-M0+ DAP.

36.5.2 SWD Programming Mode Entry

After the SWD port is acquired, the host must enter the device programming mode within a specific time window. This is done by setting the TEST_MODE bit (bit 31) in the test mode control register (MODE register). The debug port should also be configured before entering the device programming mode. Timing specifications and pseudo code for entering the programming mode are detailed in the *Traveo II MCU Programming Specifications* document.

36.5.3 SWD Programming Routines Executions

When the device is in programming mode, the external programmer can start sending the SWD packet sequence to perform programming operations such as flash erase, flash program, checksum verification, and so on. The programming routines are explained in the [Nonvolatile Memory Programming chapter on page 694](#). The exact sequence of calling the programming routines is given in the *Traveo II MCU Programming Specifications* document.

36.6 Registers

Table 36-3. List of Registers

Register Name	Description
SYSAP_ROM	System Debug Access Port ROM-Table with Cypress Vendor/Silicon ID
CM0P_DWT	Cortex M0+ Data Watchpoint and Trace (DWT) registers
CM0P_BP	Cortex M0+ System Control Space (SCS) Registers
CM0P_SCS	Cortex M0+ BreakPoint (BP) registers
CM0P_ROM	Cortex M0+ CPU Coresight ROM table
CM0P_EXT_ROM	Cortex-M0+ system ROM-Table with Cypress Vendor/Silicon ID
CM0_MTB	Cortex-M0+ MTB Registers
CM0P_MTB_SRAM	Cortex-M0+ MTB SRAM
CM0P_CTI	Cortex M0+ CTI registers
CM7_ITM	Cortex-M7 Instrumentation Trace Macrocell (ITM) Registers
CM7_DWT	Cortex-M7 Data Watchpoint and Trace (DWT) Registers
CM7_FPB	Cortex-M7 Flash Patch and Breakpoint (FPB) Registers
CM7_SCS	Cortex-M7 System Control Space (SCS) Registers
CM7_ETM	Cortex-M7 Embedded Trace Macrocell (ETM) Registers
CM7_CTI	Cortex-M7 CTI Registers
TRC_CTI	System Trace CTI
TRC_ITM_CSTF	System Trace ITM CoreSight Trace Funnel (CSTF) Registers
TRC_ETM_CSTF	System Trace ETM CoreSight Trace Funnel (CSTF) Registers
TRC_ETB_CSTF	System Trace ETB CoreSight Trace Funnel (CSTF) Registers
TRC_ETB	System Trace Embedded Trace Buffer (ETB) Registers
TRC_TPIU	System Trace Coresight Trace Port Interface Unit (TPIU) Registers
CM7_EXT_ROM	Cortex-M7 system ROM-Table with Cypress Vendor/Silicon ID
CM7_ROM	Cortex-M7 CPU CoreSight ROM-table

37. Nonvolatile Memory Programming



Nonvolatile memory programming refers to the programming of flash memory in the Traveo II Body High (TVII-B-H) device. This chapter explains the different functions that are part of device programming, such as erase, write, program, and checksum calculation. Cypress-supplied programmers and other third-party programmers can use these functions to program the TVII-B-H device with the data in an application hex file. They can also be used to perform bootloader operations where the CPU will update a portion of the flash memory.

The TVII-B-H device supports programming through the debug and access port (DAP) of Cortex-M7 (CM7_0 and CM7_1) and Cortex-M0+ CPUs.

37.1 Functional Description


The user software must set up CM0+ IRQ0 and IRQ1 interrupts correctly for system call management. The boot code automatically sets CPUSS_CM0_SYSTEM_INT_CTL0.CPU_INT_VALID bit to 1 and CPUSS_CM0_SYSTEM_INT_CTL0.CPU_INT_IDX [2:0] bits to b'000. Hence, the mapping of System Interrupt 0 (IPC Interrupt Structure 0 interrupt) to CM0+ IRQ0 for system calls is done by boot code and CM0+ IRQ0 is triggered by IPC Interrupt Structure 0 interrupt.

However, the software needs to ensure that CM0+ IRQ0 and IRQ1 are enabled and they are configured with the correct priorities as this is not automatically done by the boot code. Also, the software must ensure that IRQ0 and IRQ1 vector entries in the user CM0+ vector table are identical to the vector entries in the default SROM vector table (addresses 0x00000040 and 0x00000044, respectively). This is achieved by copying values from the SROM vector table to the user vector table during runtime if it is located in RAM; otherwise, hard-coded values need to be used and reconfirmed if the target MCU or revision changes.

As explained in [Operating Modes and Privilege Levels on page 38](#), when the CPU is executing code in Thread mode, the CONTROL register can be configured to use the Process Stack Pointer (PSP) or Main Stack Pointer (MSP). In Handler mode, the MSP is always used. Note that the CPU enters Thread mode and uses MSP when it comes out of reset. The software must take special care while setting up the system call interrupts because this is dependent on the CPU mode (Thread or Handler) of CM0+ and the stack pointer (MSP or PSP) used when the system call was triggered.

Case 1: If the software triggers system calls only when CM0+ is in Handler mode, then ensure that the software sets CM0+ IRQ1 with a higher priority than IRQ0. To do this, set the IRQ0 priority to '1'. By default, IRQ1 priority will be set to '0'.

Figure 37-1. System Call Interface Using IPC



Case 2: If the software triggers system calls with any other CPU states for CM0+ (for example, if Thread mode and PSP is used), then the software additionally needs to use another CM0+ interrupt (such as IRQ2), which acts as a manager for system calls. This approach in principle can also be used for any CPU state (including handler mode). This is a more generic approach to manage system calls under all CPU states.

The following steps are involved in the CM0+ IRQ setup with this approach.

1. Set up the system call manager function (Sys_Call_Manager) as IRQ handler for CM0+ IRQ2 in the user vector table.
2. Map the IPC Interrupt Structure 0 interrupt CM0+ IRQ2.
3. IRQ2 must be given the lowest priority in relation with IRQ0 and IRQ1. The same highest priority is given for both IRQ0 and IRQ1. For example, set IRQ2 priority to 1; by default, IRQ0 and IRQ1 priority will be 0.
4. IRQ2 handler triggers IRQ0 in software.
5. IRQ2 handler clears the pending bit of IRQ0.

Thus, the CM0+ vector table will have the entries for the first three interrupts as shown in [Table 37-1](#).

Table 37-1. CM0+ Vector Table

Interrupt Number	Handler
...	...
IRQ0	Contents of address (0x00000040)
IRQ1	Contents of address (0x00000044)
IRQ2	Sys_Call_Manager
...	...

Note that instead of directly assigning Sys_Call_Manager as CM0+ IRQ2 handler, it can be combined with multiple other system interrupts when using a dispatcher implementation. This is explained in the [Interrupts chapter on page 144](#).

A pseudo code for the interrupt configuration needed for system call for Case 2 is as follows.

```
/* IRQ2 handler function for IPC Interrupt structure 0 interrupt. This is the system call
manager function */
void Sys_Call_Manager()
{
    /* Trigger IRQ0 in Software by writing to ISPR register */
    CM0P_SCS_ISPR = 1;

    /* Read back the register to ensure that the write has happened */
    CM0P_SCS_ISPR;

    /* Clear the NVIC Pending bit of IRQ0. This is done as a fallback in case the system call
was suppressed (e.g., by disabled interrupts) */
    CM0P_SCS_ICPR = 1;

    /* Read back the register to ensure that the write has happened */
    CM0P_SCS_ICPR;
}

/* Application function for interrupt configurations */
void interrupt_configure()
{
    /* Enable CM0+ IRQ0, IRQ1 and IRQ2 */
    CM0P_SCS_ISER = 7;

    /* Set Priority 0 for IRQ0, IRQ1 and Priority 1 for IRQ2 */
    CM0P_SCS_IPR0 = 0x00400000;

    /* Connect IPC Interrupt Structure 0 Interrupt (System Interrupt 0) to
IRQ2. The interrupt triggers Sys_Call_Manager */
}
```

```

CPUSS_CM0_SYSTEM_INT_CTL0.CPU_INT_IDX = 2;

/* Clear the PRIMASK register to enable the interrupts. This could also
be done by the application at a later point in time */
__ASM("cpsie i");
}

```

All system call requests from the master can arrive at the same time; the requests are prioritized at CM0+ > CM7_0 > CM7_1 > DAP.

The TVII-B-H device IPC component implements two 32-bit data registers, but only one of these two registers is used to pass parameters to the system calls. This argument is either a pointer to SRAM or a formatted opcode or argument value that cannot be a valid SRAM address. The encoding used for DAP and the CM7_0, CM7_1, or CM0+ is slightly different.

- DAP: If (opcode + argument) is less than or equal to 31 bits, store them in the data field and set the LSb of the data field as '1'. Upon completion of the call, a return value is passed in the IPC data register. For calls that need more argument data, the data field is a pointer to a structure in SRAM (aligned on a word boundary) that has the opcode and the argument. Therefore, it is a pointer if and only if the LSb is 0.
- CM7_0, CM7_1, or CM0+: A pointer is always used to a structure in SRAM. Commands that are issued as a single word by DAP can still be issued by CM0+ or CM7_0/CM7_1, but use an SRAM structure instead.

The IRQ0 interrupt handler for system calls works as follows.

- If the ROM boot process code is not initialized in the protection state (PROTECTION is still at its default/reset value UNKNOWN), the IRQ0 calls have no effect and the handler returns.
- A jump table is used to point to the code in ROM or flash. This jump table is in ROM or flash (as configured in supervisory flash).

The IPC mechanism is used to return the result of the system call. Two factors need to be considered.

- The result is to be passed in SRAM: CM0+ writes the result in the SRAM location provided by the requester and releases the IPC structure. The requester knows that the result is ready from the RELEASE interrupt.
- The result is scalar (32 bits): CM0+ writes the result to the data field of the IPC structure and releases it. The requester can read the data when the IPC structure lock is released. The requester polls the IPC structure to know when it is released.

External programmers program the flash memory of TVII-B-H-- using the JTAG or SWD protocol by sending the commands to the DAP. The programming sequence for TVII-B-H with an external programmer is given in the *Traveo II MCU Programming Specifications*. Flash memory can also be programmed by the CM7_0/CM7_1/CM0+ CPU by accessing the IPC interface. This type of programming is typically used to update a portion of the flash memory as part of a bootstrap operation, or other application requirements, such as updating a lookup table stored in the flash memory. All write operations to flash memory, whether from the DAP or from the CPU, are done through the CM0+.

37.2 System Call Implementation

37.2.1 System Call via CM0+ or CM7_0 or CM7_1

System calls can be made from the CM0+ or CM7_0/CM7_1 at any point during code execution. CM0+ or CM7_0/CM7_1 should acquire the IPC_STRUCT reserved for them and provide arguments in either of the methods described earlier and notify IPC interrupt 0 to trigger a system call.

37.2.2 System Call via DAP

When the debug interface is acquired, then the boot ROM enters "busy-wait loop" and waits for commands issued by the DAP. For a detailed description on acquiring the debug interface see the *Traveo II MCU Programming Specifications*.

37.2.3 Exiting from a System Call

When the API operation is complete, CM0+ will release the IPC structure that initiated the system call. If an interrupt is required upon release, then the corresponding mask bit should be set in IPC_INTR_STRUCT.INTR_MASK.RELEASE[i].

37.3 SROM API Library

SROM has two categories of APIs:

- Flash management APIs – These APIs provide the ability to program, erase, and test the flash macro.
- System management APIs – These APIs provide the ability to perform system tasks such as blowing eFuse and checksum.

Table 37-2 shows a summary of the APIs.

Table 37-2. List of System Calls

System Call	Opcode	Description	Access Allowed		
			Normal ^a	Secure	Dead
BlankCheck	0x2A	Performs blank check on the addressed work flash	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
BlowFuseBit	0x01	Blows an eFuse bit	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1	
Calibrate	0x13	Applies trims from eFuse and validates Sflash and then load the trims	CM0+, CM7_0, CM7_1, DAP		CM0+, CM7_0, CM7_1, DAP
CheckFactoryHash	0x27	Generates the FACTORY_HASH according to TOC1 and compares with the FACTORY1_HASH fuses	CM0+, CM7_0, CM7_1, DAP		
CheckFMStatus	0x07	Returns the status of the flash operation	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
Checksum	0x0B	Calculates the checksum of a flash region	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
ComputeBasicHash	0x0D	Computes the hash value of a flash region	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
ConfigureFMIInterrupt	0x08	Configures the flash macro interrupt	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
EnterFlashMarginMode	0x20	Enters margin mode for main flash	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
EraseAll	0x0A	Erases all flash	CM0+, CM7_0, CM7_1, DAP		DAP
EraseResume	0x23	Resumes a suspended erase operation	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
EraseSector	0x14	Erases a flash sector	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
EraseSuspend	0x22	Suspends and ongoing erase operation	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
ExitFlashMarginMode	0x21	Exits margin mode for main flash	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
GenerateHash	0x1E	Returns the truncated SHA-256 of the flash boot programmed in Sflash	CM0+, CM7_0, CM7_1, DAP		
SwitchOverRegulators	0x11	Switches between REGHC and linear regulators	CM0+	CM0+	
ConfigureRegulator	0x15	Configures high-current regulator (REGHC) for devices that include REGHC, or PMIC for devices that use PMIC control without REGHC	CM0+	CM0+	
ProgramRow	0x06	Programs the addressed flash page	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
ProgramWorkFlash	0x30	Programs the addressed work flash page	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP

Table 37-2. List of System Calls

System Call	Opcode	Description	Access Allowed		
			Normal ^a	Secure	Dead
ReadFuseByte	0x03	Reads addressed eFuse byte	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	
ReadFuseByteMargin	0x2B	Reads addressed eFuse byte marginally	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	
ReadSWPU	0x2C	Reads the identified SWPU from SRAM	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	
ReadUniqueID	0x1F	Reads the unique ID of the die from flash	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
SetEnforcedApproval	0x2E	Sets the Enforced Approval bit in SRAM	CM0+	CM0+	CM0+
SiliconID	0x00	Returns Family ID, Revision ID, Silicon ID and protection state	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
SoftReset	0x1B	Provides system reset or CM7_0, CM7_1 only reset	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP
TransitiontoRMA	0x28	Converts parts from SECURE to RMA life-cycle stage	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	
TransitiontoSecure	0x2F	Converts parts to Secure life-cycle stage	CM0+, CM7_0, CM7_1, DAP		
DirectExecute	0x0F	Directly executes code located at a configurable address	DAP		
WriteRow	0x05	Programs Sflash	CM0+, CM7_0, CM7_1, DAP		CM0+, CM7_0, CM7_1, DAP
WriteSWPU	0x2D	Updates the identified SWPU in SRAM	CM0+, CM7_0, CM7_1, DAP	CM0+, CM7_0, CM7_1, DAP	
DebugPowerUpDown	0x12	Enables/disables CM7 debugging	CM0+, DAP	CM0+, DAP	CM0+, DAP
LoadRegulatorTrims	0x16	Sets proper trims to PWR_TRIM_HT_PWRSYS_CTL	CM0+	CM0+	

a. Refer to the [Chip Operational Modes](#) chapter on page 160.

37.4 System Calls

Table 37-2 lists all the system calls supported in TVII-B-H device along with the function description and availability in device protection modes. See the [Device Security chapter on page 158](#) for more information on the device protection settings. Note that some system calls cannot be called by the CM7_0, CM7_1, CM0+, or DAP as given in the table. Detailed information on each system call follows the table.

Note: System calls that require more than 32-bit arguments, such as Program Row and Erase Sector APIs, will first fetch the parameter address from IPC_DATA0 to derive further arguments and expect it to be a 32-bit aligned address in SRAM; if this is not followed, then the SROM API will trigger a HardFault.

37.4.1 BlankCheck

Performs blank check on the addressed work flash in blocking mode.

Table 37-3. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x2A	Blank Check opcode.
Bits [23:16]		Not used.
Bits [15:8]		Not used.
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:0]		Work flash address whose blank check needs to be performed. It should be provided in 32-bit system address format.
SRAM_SCRATCH_ADDR + 0x08		
Bits [31:16]		Not used.
Bits [15:0]	0: 1 word 1: 2 words ...	Number of words to be checked.

Table 37-4. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in background 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [27:24]		Not used.
Bits [23:8]		In case of fail, first failed word number
Bits [7:0]		In case of fail, error code (see SROM API status codes)

37.4.2 BlowFuseBit

This function blows the addressed eFuse bit. The read value of a blown eFuse bit is '1'. Parameters and result are described here.

Call the BlowFuseBit API with the ROM boot default clock configuration settings. One way to achieve this is to set SFLASH.TOC2_FLAGS.CLOCK_CONFIG = 0x3 and acquire before the user application modifies the clock settings.

The following are recommended clk_hf0 configurations:

F = 8MHz; FLL_ENABLE = 0 or

F = 25 MHz; FLL_ENABLE = 1 (FLL = 50 MHz) and CLK_ROOT_SELECT.ROOT_DIV = DIV_BY_2; or

F = 50 MHz; FLL_ENABLE = 1 (FLL = 100 MHz) and CLK_ROOT_SELECT.ROOT_DIV = DIV_BY_2;

Table 37-5. Arguments if IPC_DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x01	Blow fuse bit opcode.
Bits [23:16]	Byte Address	Refer to the Device Security chapter on page 158 for more details.
Bits [15:12]	Macro Address	
Bits [11]	Not used.	
Bits [10:8]	Bit Address	
Bits [7:1]	Not used.	
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-6. Arguments if IPC_DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x01	Blow fuse bit opcode.
Bits [23:16]	Byte Address	Refer to the Device Security chapter on page 158 for more details.
Bits [15:12]	Macro Address	
Bits [11]	Not used.	
Bits [10:8]	Bit Address	
Bits [7:0]	Not used.	

Note: Because the SRAM address is 32-bit aligned, the last two bits of the address are 0. Therefore, IPC_DATA[0] is 0.

Refer to [Customer eFuses on page 739](#) for details about Macro Address/Byte Address calculation.

Table 37-7. Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-8. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.3 Calibrate

Applies trims from eFuse and validates Sflash and then load the trims. This API can be invoked before blowing the NORMAL fuse to check if the nonvolatile memory configurations are correct.

If the hash fails then computed hash will be OR'd with STATUS_HASH_FAIL.

Table 37-9. Arguments if IPC_DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x13	Calibrate opcode.
Bits [23:16]	Not used.	
Bits [15:8]	1: load trims from eFuse Other: do not load trims from eFuse	Enable eFuse
Bits [7:1]	Not used.	
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-10. Parameters if CM0+/CM7_0/CM7_1 is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x13	Calibrate opcode.
Bits [23:16]	Not used.	
Bits [15:8]	1: load trims from eFuse Other: do not load trims from eFuse	Enable eFuse
Bit [0]	Not used.	

Table 37-11. Return if IPC_DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR 0xF1 = Computed Hash	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any) or computed hash	See 37.5 System Call Status for details.

Table 37-12. Return if CM0+/CM7_0/CM7_1 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR 0xF1 = Computed Hash	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any) or computed hash	See 37.5 System Call Status for details.

37.4.4 CheckFactoryHash

Generates the FACTORY_HASH according to TOC1 and compares with the FACTORY1_HASH fuses.

Table 37-13. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x27	Check Factory Hash opcode.
Bits [23:1]	Not used.	
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-14. Parameters if CM0+/CM7_0/CM7_1 is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x27	Check Factory Hash opcode.
Bits [23:0]	Not used.	

Table 37-15. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR (INVA-LID_FACTORY_HASH)	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-16. Return if CM0+/CM7_0/CM7_1 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR (INVA-LID_FACTORY_HASH)	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.5 CheckFMStatus

This API returns the status of the flash operation.

Note: The flash operation status can be retrieved by directly reading the FLASHC_STATUS register without the use of the CheckFMStatus API.

Table 37-17. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x07	Check FM Status opcode.
Bits [23:1]	Not used.	
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-18. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x07	Check FM Status opcode.
Bits [23:0]	Not used.	

Table 37-19. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [8:0]	Error code (0x1) or status	See 37.5 System Call Status for details.

Table 37-20. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code or status	See 37.5 System Call Status for details.

Table 37-21. Status

Bits	Name	Description
0	PGM_CODE	Indicates if active PGM operation to the code flash is taking place 0: not running 1: running
1	PGM_WORK	Indicates if active PGM operation to the work flash is taking place 0: not running 1: running
2	ERASE_CODE	Indicates if active Erase operation to the code flash is taking place 0: not running 1: running
3	ERASE_WORK	Indicates if active Erase operation to the work flash is taking place 0: not running 1: running
4	ERS_SUSPEND	Indicates if Erase operation (code/work) is currently being suspended 0: not suspended 1: suspended
5	BLANK_CHECK_WORK	Indicates if Blank Check mode is currently running on the work flash 0: not running 1: running
6	BLANK_CHCEK_PASS	Indicates the Blank check command result is PASS (Blank) 0: Not Blank 1: Blank (PASS)+G76
7	HANG	After embedded operation (pgm/erase) this flag will tell if it was successful or failed 0: PASS 1: FAIL
8	BUSY	Whenever the device is in embedded mode the RDY goes low. Should be the same as the c_interrupt pin (but inverted) 1: busy in embedded 0: rdy (high also in erase suspend)

37.4.6 Checksum

Checksum reads either the whole flash or a row of flash, and returns the sum of each byte read.

Bytes 1 and 2 of the parameter select whether the checksum is performed on the whole flash, or a row of flash. The row of Sflash or main or work flash is determined by the Row Id Lo and Row Id Hi parameters.

This API checks if the client has read access to the requested memory region by looking into DAP MPUs and SMPUs. If the client does not have read access, then STATUS_ROW_PROTECTED status is returned.

If the flash is configured in dual bank mode, then the appropriate bank needs to be provided when the whole flash option is selected. If bank 1 is selected in single bank mode, then API will return invalid argument status. Note that only one bank of Sflash is exposed.

Parameters and result are described here.

Table 37-22. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x0B	Checksum opcode
Bits [23:22]	0 - code 1 - work Other - supervisory	Flash region
Bit [21]	0 - page 1 - whole memory	Page or whole memory
Bits [20:8]		Row ID
Bit [7]	0 - Bank 0 1 - Bank 1	Bank (Only for dual bank device)
Bits [6:1]		Not used.
Bit [0]	1	Indicates that all the arguments are passed in the IPC_DATA0 register

Table 37-23. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x0B	Checksum opcode
Bits [23:22]	0 - code 1 - work Other - supervisory	Flash region
Bit [21]	0 - page 1 - whole memory	Page or whole memory
Bits [20:8]		Row ID
Bit [7]	0 - Bank 0 1 - Bank 1	Bank (Only for dual bank device)
Bits [6:1]		Not used.
Bit [0]	0	

Table 37-24. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA1 Register		
Bits [31:0]	Checksum	
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-25. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
IPC_DATA1 Register		
Bits [31:0]	Checksum	
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.7 ComputeBasicHash

This function generates the hash of the flash region provided using the formula:

$$H(n+1) = \{H(n) \times 2 + \text{Byte}\} \% 127; \text{ where } H(0) = 0$$

This function returns an invalid address status if called on an out-of-bound flash region.

This function checks if the client has read access to the requested memory region by looking into DAP MPUs and SMPUs. If the client does not have read access, then STATUS_ROW_PROTECTED status is returned.

The first byte of the parameter specifies if a CRC8SAE is computed based on the following polynomial

$$x^8 + x^4 + x^3 + x^2 + 1$$

Table 37-26. Parameters

Address	Value to be Written	Description
IPC_DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x0D	Compute Hash opcode.
Bits [23:16]		Not used.
Bits [15:8]	0x01 - CRC8SAE Other - Basic Hash	Hash type
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:0]		Start address (32-bit system address of the first byte of the data).
SRAM_SCRATCH_ADDR + 0x08		
Bits [31:0]	0 - 1 byte 1 - 2 bytes, 2 - 3 bytes and so on	Number of bytes.

Table 37-27. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Hash of the data	Hash of data if status is SUCCESS; otherwise, error code.

37.4.8 ConfigureFMInterrupt

Configures the flash macro interrupt.

The following functionalities are provided:

- Set interrupt mask
- Clear interrupt mask
- Clear interrupt

Table 37-28. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x08	Configure Flash Macro Interrupt opcode.
Bits [15:8]	0: Clear interrupt mask 1: Set interrupt mask Other: Clear interrupt	
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-29. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x08	Configure Flash Macro Interrupt opcode.
Bits [15:8]	0: Clear interrupt mask 1: Set interrupt mask Other: Clear interrupt	
Bits [7:0]	Not used.	

Table 37-30. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [8:0]	Error code (0x1) or status	See 37.5 System Call Status for details.

Table 37-31. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [8:0]	Error code (0x1) or status	See 37.5 System Call Status for details.

37.4.9 EnterFlashMarginMode

Enters margin mode for main flash. When flash enters margin mode, all other flash program/erase APIs will be blocked and any attempt to call them will return STATUS_EMB_ACTIVE.

Table 37-32. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x20	Enter Flash Margin Mode opcode.
Bits [17]	0: internal device defaults used from margin reads reference current 1: MARGIN_DCS_TRIM configuration is used during margin read	DCS_TRIM_EN
Bits [16:8]	Trim that serves as the reference current to the area between 4–8 μ A (around the 6 μ A normal static reference current)	DCS_TRIM
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:24]		Not used.
Bits [16]	0: ERS Margin is checked 1: PGM Margin is checked	PGM_ERS_B
Bits [15:8]	rdreg_c trim to be used in Margin mode if enabled by MARGIN_MODE_RDREG_CHNG_EN	RD_REG_TRIM
Bits [7:0]	0x01 - Enable rd_reg Other - Do not enable rd_reg	RD_REG_EN

Table 37-33. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]		In case of fail, error code (see SROM API status codes)

37.4.10 EraseAll

This function erases the whole flash macro specified. This API will erase only the code flash. The API returns fail status if user does not have write access to flash according to SMPU settings.

Note that when in dual bank mode, the API will always erase the alternate bank addressed from 0x12000000.

Table 37-34. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x0A	Erase All opcode.
Bits [23:1]		Not used.
Bit [0]	1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-35. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x0A	Erase All opcode.
Bits [23:0]		Not used.

Table 37-36. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-37. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.11 EraseResume

This function resumes a suspended erase operation.

Table 37-38. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x23	Erase Resume opcode.
Bits [23:16]	0x01 - Set FM interrupt mask. Other - FM interrupt mask not set.	Interrupt mask, only applicable when non-blocking.
Bits [15:8]	0x01 - API blocks CM0+ Other - non-blocking	Blocking mode
Bits [7:1]		Not used
Bit [0]	0x1	Indicates all arguments are passed in the IPC_DATA0 register.

Table 37-39. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x23	Erase Resume opcode.

Table 37-39. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
Bits [23:16]	0x01 - Set FM interrupt mask. Other - FM interrupt mask not set.	Interrupt mask, only applicable when non-blocking.
Bits [15:8]	0x01 - API blocks CM0+	
Other - non-blocking	Blocking mode	
Bits [7:0]		Not used

Table 37-40. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-41. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.12 EraseSector

This function starts erase operation on the specified sector. This function cannot be called on Sflash; the API will return STATUS_INVALID_FLASH_ADDR if invoked on Sflash.

EraseSector is allowed on the sector that is erase suspended. If EraseSector is called on a sector other than the suspended one, then the new sector will be erased and the suspended sector will be in an unknown state. EraseSector can be called on the suspended sector to restore to blank state.

Table 37-42. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x14	Erase Sector opcode.
Bits [23:16]	0x01 - Set FM interrupt mask. Other - FM interrupt mask not set.	Interrupt mask, only applicable when non-blocking.
Bits [15:8]	0x01 - API blocks CM0+ Other - non-blocking	Blocking mode
Bits [7:0]		Not used
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:0]		Flash address to be erased. Should be provided in 32-bit system address format. For example, to erase the second sector you need to provide the 32-bit system address of any of the bytes lying in the second sector.

Table 37-43. Return

Address	Return Value	Description
SRAM_SCRATCH Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.13 EraseSuspend

This function suspends an ongoing erase operation. User should not read from a sector that is suspended. The Program Row API function will return error if invoked on suspended sector.

Table 37-44. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x22	Erase Suspend opcode.
Bits [23:1]		Not used
Bit [0]	0x1	Indicates all arguments are passed in the IPC_DATA0 register.

Table 37-45. Parameters if CM0+/CM7_0/CM7_1 is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x22	Erase Suspend opcode.
Bits [23:0]		Not used

Table 37-46. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [27:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-47. Return if CM0+/CM7_0/CM7_1 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [27:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.14 ExitFlashMarginMode

Exits margin mode for main flash.

Table 37-48. Parameters if DAP is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x21	Exit Flash Margin Mode opcode.
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-49. Parameters if CM0+/CM7_0/CM7_1 is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x21	Exit Flash Margin Mode opcode.
Bits [23:0]	Not used.	

Table 37-50. Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-51. Return if CM0+/CM7_0/CM7_1 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.15 GenerateHash

This API returns the truncated SHA-256 of the flash boot programmed in Sflash and optionally includes public key and other objects as indicated in Table of Contents (TOC).

This function gets the flash boot size from TOC.

Typically, this function will be called to check if the HASH to be blown into eFuse matches with what ROM boot expects it to be.

Note: If the TOC1/TOC2 hash object start address is wrong or is an unaligned address (not 4-byte aligned), a hard-fault will be generated.

Table 37-52. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x1E	Generate hash opcode.
Bits [23:16]		Not used.
Bits [15:8]	0x1: returns FACTORY_HASH Other: returns hash of all objects according to TOC1 and 2	Factory
Bits [7:0]		Not used.

Table 37-53. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in background 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]		In case of fail, error code (see SROM API status codes)
SRAM_SCRATCH_ADDR + 0x4		
Bits [31:0]	HASH_WORD0	
SRAM_SCRATCH_ADDR + 0x8		
Bits [31:0]	HASH_WORD1	
SRAM_SCRATCH_ADDR + 0xC		
Bits [31:0]	HASH_WORD2	
SRAM_SCRATCH_ADDR + 0x10		
Bits [31:0]	HASH_WORD3	
SRAM_SCRATCH_ADDR + 0x14		
Bits [31:0]	HASH_ZEROS	

37.4.16 SwitchOverRegulators

This function is used to switch between the high-current regulator (REGHC or PMIC without REGHC) required to run CM7 and the linear regulator (LDO). It should be called to switch from LDO to REGHC before enabling CM7. The Configure Regulator system call should be called before using this function.

Note: If the API is called in the blocking mode it will handle setting of the proper regulator trims. However, if the API is called in the non-blocking mode, then the proper trims will not be set for transition to the external regulator. Therefore, the trims should be set when transition is complete.

For transition to LDO, the required trims will be set by the API before the transition is initialized.

Table 37-54. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x11	Opcode.
Bit [3]	Blocking: 0- Non-blocking call 1- Blocks CM0+ until the transition completes	
Bit [2]	Regulator: 0 - Switch over to REGHC (PMIC without REGHC) 1 - Switch over to linear regulator	
Bit [1]	Operating mode: 0 - External transistor 1 - External PMIC	For devices without REGHC, the Operating Mode is ignored.
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-55. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x11	Opcode.
Bits [23:16]	Blocking: 0- Non-blocking call 1- Blocks CM0+ until the transition completes	
Bits [15:8]	Select Regulator: 0 - Switch over to REGHC 1 - Switch over to LDO	
Bit [1]	Operating mode: 0 - External transistor 1 - External PMIC	For devices without REGHC, the operating mode is ignored.
Bit [0]	Not used	

Table 37-56. Return if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR 0xF1 = Computed Hash	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any) or computed hash	See 37.5 System Call Status for details.

Table 37-57. Return if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR 0xF1 = Computed Hash	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any) or computed hash	See 37.5 System Call Status for details.

37.4.17 ConfigureRegulator

This function is used to configure the high-current regulator (REGHC) for devices that include REGHC, or PMIC for devices that use PMIC control without REGHC. It should be called to configure the desired regulator only once before switching to the regulator using the Switch Over Regulators system call.

Table 37-58. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x15	Opcode.
Bits [15:13]	Radj Value	RADJ reset threshold value These bits are invalid in CYT4D and CYT3D series devices
Bits [12:8]	$v_{adj} = (1.020 \text{ V} + V_{adjTrim} \times 0.005 \text{ V})$	VADJ trim value (VadjTrim) used in the regulator output trim
Bit [7]	0 - Device generates VADJ when PMIC is enabled. 1 - Device does not generate VADJ, and it must not be part of the PMIC feedback loop	This bit configures REGHC_PMIC_VADJ_DIS
Bit [6]	0 - no Radj 1 - use Radj	Use Radj to generate a reset threshold for the PMIC This bit is invalid in CYT4D and CYT3D series devices
Bit [5]	0 - Internal Active Linear Regulator is disabled after PMIC is enabled. OCD is disabled 1 - Internal Active Linear Regulator is kept enabled. See the related datasheet for the minimum PMIC V_{CCD} input to prevent OCD	UseLinReg
Bit [4]	0 - Allow HC regulator to operate in Normal mode 1 - Allow HC regulator to operate in DeepSleep mode	This bit configures PMIC behavior in DeepSleep mode This bit is invalid in CYT4D and CYT3D series devices
Bit [3]	Reset Polarity: 0 - Logic low is used for enable 1 - Logic high is used for enable	The polarity used to trigger a reset action based on the PMIC status input of the regulator

Table 37-58. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
Bit [2]	Enable Polarity: 0 – Logic low is used for enable 1 – Logic high is used for enable	Polarity used to enable the regulator
Bit [1]	0 – External transistor 1 – External PMIC	Operating mode For devices without REGHC, the operating mode is ignored.
Bit [0]	0x1	Indicates that all arguments are passed in the IPC_DATA0 register
IPC_DATA1 Register		
Bit [8:0]	Wait Count	Wait count in steps of 4 μ s after PMIC status is OK. This is used by the hardware sequencer to allow additional settling time before disabling the internal regulator. Note that the ConfigureRegulator API supports Wait Count field [8:0]; it does not support Wait Count bit [9]. Therefore, if the user needs Wait Count more than 0x1FF, set the PWR_REGHC_CTL.REGHC_PMIC_STATUS_WAIT[29:20] register after the ConfigureRegulator API is successful.

Table 37-59. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR1	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
IPC_DATA1 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR2	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR1		
Bits [31:24]	0x15	Opcode.
Bits [15:13]	Radj Value	RADJ reset threshold value These bits are invalid in CYT4D and CYT3D series devices
Bits [12:8]	$v_{adj} = (1.020\text{ V} + V_{adjTrim} \times 0.005\text{ V})$	VADJ trim value ($V_{adjTrim}$) used in the regulator output trim.
Bit [7]	0 - Device generates VADJ when PMIC is enabled. 1 - Device does not generate VADJ, and it must not be part of the PMIC feedback loop	This bit configures REGHC_PMIC_VADJ_DIS
Bit [6]	0 - no Radj 1 - use Radj	Use Radj to generate a reset threshold for the PMIC This bit is invalid in CYT4D and CYT3D series devices
Bit [5]	0 - Internal Active Linear Regulator is disabled after PMIC is enabled. OCD is disabled 1 - Internal Active Linear Regulator is kept enabled. See the datasheet for the minimum PMIC V_{CCD} input to prevent OCD	UseLinReg
Bit [4]	0 - Allow HC regulator to operate in Normal mode 1 - Allow HC regulator to operate in DeepSleep mode	This bit configures PMIC behavior in DeepSleep mode This bit is invalid in CYT4D and CYT3D series devices

Table 37-59. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
Bit [3]	Reset Polarity: 0 – Logic low is used for enable 1 – Logic high is used for enable	The polarity used to trigger a reset action based on the PMIC status input of the regulator.
Bit [2]	Enable Polarity: 0 – Logic low is used for enable 1 – Logic high is used for enable	Polarity used to enable the regulator.
Bit [1]	0 – External transistor 1 – External PMIC	Operating mode. For devices without REGHC, the operating mode is ignored.
Bit [0]	Not used.	
SRAM_SCRATCH_ADDR2		
Bit [8:0]	Wait Count	Wait count in steps of 4 μ s after PMIC status is OK. This is used by the hardware sequencer to allow additional settling time before disabling the internal regulator. Note that the ConfigureRegulator API supports Wait Count field [8:0]; it does not support Wait Count bit [9]. Therefore, if the user needs Wait Count more than 0x1FF, set the PWR_REGHC_CTL.REGHC_PMIC_STATUS_WAIT[29:20] register after the ConfigureRegulator API is successful.

Table 37-60. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-61. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR1		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.18 ProgramRow

This function programs the addressed flash page (the flash can be code flash or work flash). The user needs to provide the data to be loaded and the flash address to be programmed. The flash page should be in the erased state before calling this function. Otherwise, it will return an error status. The function returns a fail status if the user does not have write access to flash according to SMPU/SWPU settings.

The FM interrupt mask option can be set to generate an interrupt from the flash macro when running with non-blocking option.

Note that the user should perform a dummy read from work flash after the program operation is complete if ProgramRow is invoked in non-blocking mode. Dummy read is required to make the logical bank of work flash ready for read operation after a

program or erase operation. This is not applicable if ProgramRow is invoked in blocking mode.

Table 37-62. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x06	Program Row opcode.
Bits [23:16]	0x01 - Skips the blank check step. Other - Perform blank check	Skip blank check
Bits [15:8]	0x01 - API blocks CM0+ Other - non-blocking	Blocking mode
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:24]	0x01 - Set FM interrupt mask. Other - FM interrupt mask not set.	Interrupt mask, only applicable when non-blocking.
Bits [23:16]		Not used.
Bits [15:8]		Not used.
Bits [7:0]	3 - 64 bits 5 - 256 bits 9 - 4096 bits	Data size for code flash. For work flash the data size is always 32 bits.
SRAM_SCRATCH_ADDR + 0x08		
Bits [31:0]		Flash address to be programmed. This should be provided in 32-bit system address format.
SRAM_SCRATCH_ADDR + 0x0C		
Bits [31:0]	SRAM_SCRATCH_DATA_ADDR	Address of SRAM where data to be programmed is stored
SRAM_SCRATCH_DATA_ADDR		
Bits [31:24]		Data byte 3 to be programmed in flash
Bits [23:16]		Data byte 2 to be programmed in flash
Bits [15:8]		Data byte 1 to be programmed in flash
Bits [7:0]		Data byte 0 to be programmed in flash
SRAM_SCRATCH_DATA_ADDR + (n-3)		
Bits [31:24]		Data byte n to be programmed in flash
Bits [23:16]		Data byte n-1 to be programmed in flash
Bits [15:8]		Data byte n-2 to be programmed in flash
Bits [7:0]		Data byte n-3 to be programmed in flash

Table 37-63. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in background 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [27:0]	Error code (if any)	See 37.5 System Call Status for details. In case of success: - 0x0 indicates successful completion of API (second phase) - 0x9 indicates successful completion of first phase, program command is ongoing in the background.

37.4.19 ProgramWorkFlash

This function programs the addressed work flash page. The function is not applicable for programming of any flash other than work flash and will return an error status when called on non-work flash. The user must provide the data to be loaded and the work flash address to be programmed. The flash page should be in the erased state before calling this function. Otherwise, it will return an error status. The function returns a fail status if the user does not have write access to flash according to SMPU/ SWPU settings.

Table 37-64. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits[31:24]	0x30	ProgramWorkFlash opcode
Bits[23:16]	0x01 - Skips the blank check step. Other - Perform blank check	Skip blank check
Bits[15:8]	0x01 - API blocks CM0+	Only blocking mode is supported
Bits[7:0]		Reserved.
SRAM_SCRATCH_ADDR + 0x04		
Bits[31:24]		Reserved.
Bits [23:16]		Reserved.
Bits [15:8]		Reserved.
Bits[7:0]	2 - 32 bits 3 - 64 bits 4 - 128 bits 5 - 256 bits 6 - 512 bits 7 - 1024 bits 8 - 2048 bits 9 - 4096 bits	Data size
SRAM_SCRATCH_ADDR + 0x08		
Bits [31:0]		Work flash address to be programmed. This should be provided in 32-bit system address format.
SRAM_SCRATCH_ADDR + 0x0C		
Bits [31:0]	SRAM_SCRATCH_DATA_ADDR	Address of SRAM where data to be programmed is stored
SRAM_SCRATCH_DATA_ADDR		
Bits [31:24]		Data byte 3 to be programmed in work flash
Bits [23:16]		Data byte 2 to be programmed in work flash
Bits [15:8]		Data byte 1 to be programmed in work flash
Bits [7:0]		Data byte 0 to be programmed in work flash
SRAM_SCRATCH_DATA_ADDR + (n-3)		
Bits [31:24]		Data byte n to be programmed in work flash
Bits [23:16]		Data byte n-1 to be programmed in work flash
Bits [15:8]		Data byte n-2 to be programmed in work flash
Bits [7:0]		Data byte n-3 to be programmed in work flash

Table 37-65. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in the background 0xF = ERROR	Status code (see 33.5 System Call Status for details).
Bits [27:0]	Error code (if any)	See 33.5 System Call Status for details. In case of success: 0x0 indicates successful completion of API (second phase) 0x9 indicates successful completion of first phase; program command is ongoing in the background.

37.4.20 ReadFuseByte

This function returns the value of an eFuse. The read value of a blown eFuse bit is '1' and that of a not blown eFuse bit is '0'. This API inherits the client protection context.

Table 37-66. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x03	Read Fuse Byte opcode.
Bits [23:8]	Value in the range [0,511]	eFuse address
Bits [7:0]	0x01	Indicates all arguments are passed in IPC_DATA0.

Table 37-67. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x03	Read Fuse Byte opcode.
Bits [23:8]	Value in the range [0,511]	eFuse address
Bits [7:0]		Not used.

Table 37-68. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	eFuse byte	Byte read from eFuse if status is success; otherwise, error code.

Table 37-69. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	eFuse byte	Byte read from eFuse if status is success; otherwise, error code.

37.4.21 ReadFuseByteMargin

API returns the eFuse contents of the addressed byte read marginally.

The read value of a blown eFuse bit is '1' and that of not blown is '0'.

This API inherits client's protection context.

Table 37-70. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x2B	Read Fuse Byte Margin opcode.
Bits [23:20]	0: Low resistance, -50% from nominal 1: Nominal resistance (default read condition) 2: High resistance (+50% from nominal) Other: Higher resistance (+100% from nominal)	Margin control
Bits [19:8]	Value in the range [0,511]	eFuse address
Bit [0]	0x01	Indicates all arguments are passed in IPC_DATA0.

Table 37-71. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x2B	Read Fuse Byte Margin opcode.
Bits [23:20]	0: Low resistance, -50% from nominal 1: Nominal resistance (default read condition) 2: High resistance (+50% from nominal) Other: Higher resistance (+100% from nominal)	Margin control
Bits [19:8]	Value in the range [0,511]	eFuse address
Bits [7:0]		Not used.

Table 37-72. Return if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	eFuse byte	Byte read from eFuse if status is success; otherwise, error code.

Table 37-73. Return if IPC_STRUCT.DATA[0] = 0

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	eFuse byte	Byte read from eFuse if status is success; otherwise, error code.

37.4.22 ReadSWPU

Reads the identified SWPU from SRAM. The PU ID is based on the storage of SWPU in Sflash. There is only one contiguous SWPU indexing in Sflash even though there are two physically separate storage in Sflash.

Table 37-74. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x2C	Read SWPU opcode.
Bits [23:16]	1: eFuse Write 2: eFuse Read Other: Flash Write	PU type
Bits [15:8]	Structure ID to be read. Indexed from 0	PU ID
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:0]	SRAM_DATA_ADDRESS	

Table 37-75. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in background 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]		In case of fail, error code (see SROM API status codes)
SRAM_DATA_ADDRESS		
Bits [31:0]	SL_OFFSET/SL_ADDRESS	
SRAM_DATA_ADDRESS + 0x4		
Bits [31:0]	SL_SIZE	
SRAM_DATA_ADDRESS + 0x8		
Bits [31:0]	SL_ATT	
SRAM_DATA_ADDRESS + 0xC		
Bits [31:0]	MS_ATT	

37.4.23 ReadUniqueID

Returns the unique ID of the die from Sflash.

Table 37-76. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x1F	Read Unique ID opcode.
Bits [23:0]		Not used.

Table 37-77. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in background 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error if any or DIE_ID0	In case of fail, error code (see SROM API status codes)
SRAM_SCRATCH_ADDR + 0x4		
Bits [31:0]	DIE_ID1	
SRAM_SCRATCH_ADDR + 0x8		
Bits [31:0]	DIE_ID2	

Note: ID includes production date of the device as well as other manufacturing information such as lot, wafer, and die serial numbers, which in combination ensures the uniqueness of the ID within the Traveo II family.

37.4.24 SetEnforcedApproval

Sets the EnforcedApproval bit in SRAM. EnforcedApproval bit is stored in PC1 private SRAM. If this bit is set then API checks for supervised marker.

Table 37-78. Parameters if DAP is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x2E	Set Enforced Approval opcode.
Bit [0]	0x01	Indicates all arguments are passed in IPC_DATA0.

Table 37-79. Parameters if CM0+/CM7_0/CM7_1 is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x2E	Set Enforced Approval opcode.
Bits [23:0]		Not used.

Table 37-80. Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]		Error code if any

Table 37-81. Return if CM0+/CM7_0/CM7_1 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]		Error code if any

37.4.25 SiliconID

This function returns a 12-bit family ID, 16-bit silicon ID, 8-bit revision ID, and the current protection state.

Note that only 32 bits are available to store the return value in the IPC structure. Therefore, the API takes a parameter ID type based on which it will return family ID and revision ID if the ID type is set to '0'. It will return silicon ID and protection state if the ID type is set to '1'.

Table 37-82. Silicon ID

Cypress IDs	Memory Location	Data
Family ID [7:0]	0xF0000FE0	Part Number [7:0]
Family ID [11:8]	0xF0000FE4	Part Number [3:0]
Major Revision	0xF0000FE8	Revision [7:4]
Minor Revision	0xF0000FEC	Rev and Minor Revision Field [7:4]
Silicon ID	Sflash	Silicon ID [15:0]
Protection state	MMIO	Protection [3:0]

Table 37-83. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x00	Silicon ID opcode
Bits [15:8]	0 - returns 0. Use the CPUSS_PRODUCT_ID register to get family ID and revision ID 1 - returns 16-bit silicon ID and protection state 2 - returns SROM firmware version Others - returns invalid argument status	ID type
Bits [7:1]		Not used.
Bit [0]	0x1	Indicates that all the arguments are passed in IPC_DATA0.

Table 37-84. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.

Table 37-84. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
SRAM_SCRATCH		
Bits [31:24]	0x00	Silicon ID opcode
Bits [15:8]	0 - returns 12-bit family ID and revision ID 1 - returns 16-bit silicon ID and protection state 2 - returns SROM firmware version Others - returns invalid argument status	ID type
Bits [7:0]		Not used.

Return if IPC_STRUCT.DATA[0] = 1

Table 37-85. If ID Type is 0

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:20]	Major Revision ID	See the <i>Traveo II MCU Programming Specifications</i> for these values.
Bits [19:16]	Minor Revision ID	
Bits [15:8]	Family ID Byte High	See the device datasheet for silicon ID values for different part numbers.
Bits [7:0]	Family ID Byte Low	

Table 37-86. If ID Type is 1

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [27:24]		Not used.
Bits [23:20]	0: VIRGIN 1: NORMAL 2: SEC_W_DBG 3: SECURE 4: RMA 5: SORT 6: PROVISIONED 7: NORMAL_PROVISIONED 9: CORRUPTED	Note that devices are in the NORMAL_PROVISIONED stage when shipped. The VIRGIN, NORMAL, SORT, and PROVISIONED life-cycle stages are not applicable for final samples.
Bits [19:16]	0: UNKNOWN 1: VIRGIN 2: NORMAL 3: SECURE 4: DEAD	Protection state
Bits [15:8]	Silicon ID Byte High	See the device datasheet for silicon ID values for different part numbers.
Bits [7:0]	Silicon ID Byte Low	

Table 37-87. If ID Type is 2

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [27:24]	Flash boot major version	
Bits [23:16]	Flash boot minor version	
Bits [15:8]	SROM firmware major version	
Bits [7:0]	SROM firmware minor version	

Return if IPC_STRUCT.DATA[0] = 0

Same values as for DAP but located in the SRAM_SCRATCH location.

37.4.26 SoftReset

Resets the system by setting CM0+ AIRCR system reset bit. This will result in a system-wide reset except for debug logic. This API can also be used to selective reset just CM7_0/CM7_1 core based on 'type' parameter. CM7_0/CM7_1 should be in DeepSleep mode when selectively resetting CM7_0/CM7_1. Soft Reset API called with Type parameter set to 1 will result in CM7_0/CM7_1 transition to Enabled state. Note that this API will return an error status if CM7_0/CM7_1 core reset is requested when CM7_0/CM7_1 is in active mode.

Table 37-88. Parameters if DAP is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x01B	Soft Reset opcode.
Bits [23:8]	Value in the range [0,511]	eFuse address
Bits [7:1]	0: System Reset 1: Only CM7_0/CM7_1 resets	Type
Bit [0]	0x01	Indicates all arguments are passed in IPC_DATA0.

Table 37-89. Parameters if CM0+/CM7_0/CM7_1 is Master

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x01B	Soft Reset opcode.
Bits [23:8]	Value in the range [0,511]	eFuse address
Bits [7:1]	0: System Reset 1: Only CM7_0/CM7_1 resets	Type
Bit [0]		Not used

Table 37-90. Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

Table 37-91. Return if CM0+/CM7_0/CM7_1 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code (if any)	See 37.5 System Call Status for details.

37.4.27 TransitiontoRMA

Converts parts from SECURE or SECURE WITH DEBUG to the RMA life-cycle stage. This API returns the 0xF00000A9 failure code if any active embedded flash operations are going on.

Table 37-92. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x28	Transition to RMA opcode.
Bits [23:0]		Not used.
SRAM_SCRATCH_ADDR + 0x4		
Bits [31:0]		Object size in bytes (including itself). It should always be 20 bytes.
SRAM_SCRATCH_ADDR + 0x8		
Bits [31:0]	0x120028F0	Command ID
SRAM_SCRATCH_ADDR + 0xC		
Bits [31:0]		Unique ID word 0
SRAM_SCRATCH_ADDR + 0x10		
Bits [31:0]		Unique ID word 1
SRAM_SCRATCH_ADDR + 0x14		
Bits [31:0]		Unique ID word 2 (3 bytes)
SRAM_SCRATCH_ADDR + 0x18		
Bits [31:0]		SRAM address where signature is stored (4 bytes)

Table 37-93. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error if any	In case of fail, error code (see SROM API status codes)

37.4.28 TransitiontoSecure

Validates the FACTORY_HASH and programs SECURE_HASH, secure access restrictions and dead access restrictions into eFuse.

Programs secure or secure with debug fuse to transition to SECURE or SECURE with DEBUG life-cycle stage.

Only allowed in NORMAL_PROVISIONED life-cycle stage

Table 37-94. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x2F	Transition to Secure opcode.
Bits [15:8]	1: Blow D fuse Other: Blow S fuse	Debug
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x4		
Bits [31:0]	bit[1:0] AP_CTL_CM0_DISABLE bit[3:2] AP_CTL_CM7_0/CM7_1_DISABLE bit[5:4] AP_CTL_SYS_DISABLE bit[6] SYS_AP_MPU_ENABLE bit[7] DIRECT_EXECUTE_DISABLE bit[10:8] FLASH_ALLOWED bit[13:11] SRAM_ALLOWED bit[15:14] WORK_FLASH_ALLOWED bit[17:16] SFLASH_ALLOWED bit[19:18] MMIO_ALLOWED	SECURE_ACCESS_RESTRICT
SRAM_SCRATCH_ADDR + 0x8		
Bits [31:0]		DEAD_ACCESS_RESTRICT

Table 37-95. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error if any	In case of fail, error code (see SROM API status codes)

37.4.29 DirectExecute

Directly executes code located at a configurable address. API is allowed in VIRGIN state. In NORMAL state, API is allowed only if the corresponding DIRECT_EXECUTE_DISABLE bit (in Sflash/eFuse) is '0'.

Table 37-96. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x0F	Direct Execute opcode.
Bits [23:2]	Address[21:0]	
Bit [1]	0: SRAM 1: Flash	Memory
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 register.

Table 37-97. Parameters if IPC0_DATA0[0] is 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x0F	Direct Execute opcode.
Bits [23:2]	Not used.	
Bits [1:0]	0: (void, void) 1: (void, long32) 2: (long32, void) 3: (long32, long32)	FuncType Value: (return,arg)
SRAM_SCRATCH_ADDR + 4		
Bits [31:0]		Argument
SRAM_SCRATCH_ADDR + 8		
Bits [31:0]		Address (32-bit system address of the code to execute)
SRAM_SCRATCH_ADDR + 12		
Bits [31:0]		Return
SRAM_SCRATCH_ADDR + 16		
Bits [31:0]		FM API status (this field is primarily used by s40flash functions to return status)

Return when arguments are passed only in IPC_DATA.

Table 37-98. On Successful Execution

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	0xF = ERROR	Does not return any status on successful execution. The function that is being executed should return a meaningful status.

Table 37-99. On Error

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:28]	0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code	See 37.5 System Call Status for details.

Table 37-100. Return when Arguments are Passed in SRAM_SCRATCH

Address	Return Value	Description
IPC_DATA0 Register		
Bits [31:0]	0xA0000000	Success status on completion of execution

37.4.30 WriteRow

This API is used to program flash. The user needs to provide data to be loaded and flash address to be programmed.

This API can be called only on Sflash.

The API is allowed only in single bank mode. When called in dual bank mode will return STATUS_INVALID_BANK_MODE.

All operations performed are blocking CM0+ & IPC used to invoke the call. This API returns an error status when called during an active embedded operation.

The API returns an invalid address error status if called on wounded flash.

The API returns fail status if the user does not have write access to flash according to SMPU settings.

This API can also be called in 'blocking' mode by setting blocking parameter as 1, in which case the API will return only after all flash operation completes. The API will be polling for each of the timer to expire instead of configuring the flash interrupt and splitting up in phases.

This API does not operate on Sflash in protection states other than VIRGIN and NORMAL.

This API can be used to program all of Sflash rows only in VIRGIN state.

This API can be used to program user Sflash rows (row 4 to 7), NORMAL Access Restriction row (row13) (refer to [Table 37-103](#) for the encoding scheme details), public key rows (row 50 to 55), and the TOC2 row (row 62) in NORMAL state. When used to program the allowed Sflash rows the API copies the flash high-voltage parameters into a local array of 512 bytes increasing the stack size accordingly. Sflash programming is always CM0+ blocking. For Traveo II,

the application protection settings (row 59) are also considered as user row and can be updated using WriteRow API.

To define flash (program/erase) access restrictions, SWPU objects need to be configured in row 59 of Sflash. This specific row in Sflash is updated using the Write Row API. See the [BootROM chapter on page 133](#) for more details.

When NORMAL access restrictions are requested to be updated in NORMAL state and if new restrictions are wider than the existing ones, the API will return the STATUS_INVALID_ACCESS_RESTRICTION status.

If WriteRow is used to program the NORMAL Access Restriction row (row13) of Sflash, first disable CM0+ cache before call to WriteRow. This can be done by writing '0' to the FLASHC_CM0_CA_CTL0.CA_EN bit. After the API is executed successfully, CM0+ cache can be again enabled by writing '1' to the FLASHC_CM0_CA_CTL0.CA_EN bit.

It does not support RWW feature. Bank 1 sector is used as backup sector as entire Bank 0 sector is erased during a WriteRow operation. The user needs to provide corresponding Bank 0 address when invoking the API. For example, if the device is in single bank mode or mapping A dual bank mode, then 0x17000000 logical base address is considered. If device is in mapping B of dual bank mode, then 0x178000000 logical base address is considered. Any request on other bank returns invalid address status. All operations performed are blocking CM0+ and IPC used to invoke the call. This API returns an error status when called during an active embedded operation.

Note: The "Inject Public Key", "Write Normal Access Restriction", and "Write TOC2" APIs have been removed from the system calls. The user must use the "Write Row" API to update normal access restriction, public key, and TOC2 to the Sflash.

Table 37-101. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x05	Write Row opcode.
Bits [23:16]		Not used.
Bits [15:8]	0x01 - API blocks CM0+ Other - non-blocking	Blocking mode
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:0]		Not used.
SRAM_SCRATCH_ADDR + 0x08		
Bits [31:0]		Flash address to be programmed. This should be provided in 32-bit system address format.

Table 37-101. Parameters

Address	Value to be Written	Description
SRAM_SCRATCH_ADDR + 0x0C		
Bits [31:0]	SRAM_SCRATCH_DATA_ADDR	Address of SRAM where data to be programmed is stored
SRAM_SCRATCH_DATA_ADDR		
Bits [31:24]		Data byte 3 to be programmed in flash
Bits [23:16]		Data byte 2 to be programmed in flash
Bits [15:8]		Data byte 1 to be programmed in flash
Bits [7:0]		Data byte 0 to be programmed in flash
SRAM_SCRATCH_DATA_ADDR + (n-3)		
Bits [31:24]		Data byte n to be programmed in flash
Bits [23:16]		Data byte n-1 to be programmed in flash
Bits [15:8]		Data byte n-2 to be programmed in flash
Bits [7:0]		Data byte n-3 to be programmed in flash

Table 37-102. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in background 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [27:0]	Error code (if any)	See 37.5 System Call Status for details. In case of success: 0x0 indicates successful completion of all phases 0x9 indicates successful completion of first phase, program command is ongoing in the background.

Table 37-103. Access Restrictions Encoding

Name	Description
bit[1:0] AP_CTL_M0_DISABLE	00 – Enable M0-AP 01 – Disable M0-AP 1x – Permanently Disable M0-AP
bit[3:2] AP_CTL_CM7_0/CM7_1_DISABLE	00 – Enable CM7_0/CM7_1-AP 01 – Disable CM7_0/CM7_1-AP 1x – Permanently Disable CM7_0/CM7_1 AP
bit[5:4] AP_CTL_SYS_DISABLE	00 – Enable SYS-AP 01 – Disable SYS-AP 1x – Permanently Disable SYS AP
bit[6] SYS_AP_MPU_ENABLE	Indicates that the MPU on the system debug port must be programmed and locked according to the settings in the next four fields.
bit[7] DIRECT_EXECUTE_DISABLE	Disables DirectExecute system call functionality

Table 37-103. Access Restrictions Encoding

Name	Description
bit[10:8] FLASH_ALLOWED	This field indicates what portion of the flash main region is accessible through the system debug port. Only a portion of flash starting at the bottom of the area is exposed. Encoding is as follows: 0: Entire region 1: Seven-eighth 2: Three-fourth 3: One-half 4: One-quarter 5: One-eighth 6: One-sixteenth 7: Nothing
bit[13:11] SRAM_ALLOWED	This field indicates what portion of the SRAM region is accessible through the system debug port. Only a portion of SRAM starting at the bottom of the area is exposed. Encoding is the same as FLASH_ALLOWED.
bit[15:14] WORK_FLASH_ALLOWED	This field indicates what portion of work flash is accessible through the system access port. Only a portion of work flash starting at the bottom of the area is exposed. Encoding is as follows: 0: Entire region 1: One-half 2: One-quarter 3: Nothing
bit[17:16] SFLASH_ALLOWED	This field indicates what portion of the flash supervisory region is accessible through the system debug port. Only a portion of Sflash starting at the bottom of the area is exposed. Encoding is as follows: 0: Entire region 1: One-half 2: One-quarter 3: Nothing
bit[19:18] MMIO_ALLOWED	This field indicates what portion of the MMIO region is accessible through the system debug port. Encoding is as follows: 0: All MMIO registers 1: Only IPC MMIO registers accessible (system calls) 2,3: No MMIO access

37.4.31 WriteSWPU

Updates the identified SWPU in SRAM if client has appropriate access. The PU ID is based on the storage of SWPU in Sflash. Only one contiguous SWPU indexing in Sflash even though there are two physically separate storage in Sflash.

The MS_ATT field of selected PU defines who can modify the specific PU structure.

The update is allowed only if the PC that is requesting the update is in the PC_MASK of MS_ATT. The update only modifies the fields SL_ATT, MS_ATT, and SL_SIZE.ENABLE. For a successful update, the other fields SL_ADDR and

SL_SIZE.REGION_SIZE should match what is stored in that entry. The safe way to update is to first read the entry, modify, and write it back.

Table 37-104. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x2D	Write SWPU opcode.
Bits [23:20]	0: Update SWPU 1: Enable SWPU Other: Disable SWPU	Control
Bits [19:16]	1: eFuse Write 2: eFuse Read Other: Flash Write	PU type
Bits [15:8]	Structure ID to be read. Indexed from 0	PU ID
Bits [7:0]		Not used.
SRAM_SCRATCH_ADDR + 0x04		
Bits [31:0]	SRAM_DATA_ADDRESS	

Table 37-105. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS/Program command ongoing in background 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]		In case of fail, error code (see SROM API status codes)
SRAM_DATA_ADDRESS		
Bits [31:0]	SL_OFFSET/SL_ADDRESS	Read only
SRAM_DATA_ADDRESS + 0x4		
Bits [31:0]	SL_SIZE	Read only
SRAM_DATA_ADDRESS + 0x8		
Bits [31:0]	SL_ATT	Used only when control is 0
SRAM_DATA_ADDRESS + 0xC		
Bits [31:0]	MS_ATT	Used only when control is 0

37.4.32 DebugPowerUpDown

The DebugPowerUpDown function is used for handling the power transitions of CM7_0/1 power domains to properly connect/disconnect debug probe to/from the device. The system call does not switch off the CM7_0_PWR_CTL. The function first waits until the RegHC power-up is complete. Then, it sets the PWR_MODE for CM7 power domain to ENABLED. When the CM7 power domain is ON (PWR_DONE = 1), the function restores the remembered PWR_MODE. This step is realized for CM7_0 and CM7_1 cores. The function temporarily enables CM7 power because the CM7 power FSM runs on the CM7 clock, which may be OFF.

Table 37-106. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x12	Opcode
Bit [1]	0: Power down 1: Power up	
Bit [0]	0x1	Indicates that all the arguments are passed in the IPC_DATA0 Register.

Table 37-107. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x12	Opcode
Bits [23:2]	Not used.	
Bit [1]	0: Power down 1: Power up	
Bit [0]	Not used.	

Table 37-108. Return if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code	See 37.5 System Call Status for details.

Table 37-109. Return if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code	See 37.5 System Call Status for details.

37.4.33 LoadRegulatorTrims

This API is used to adapt the output voltage for internal regulators during handover. This API must be called every time a load transition requires switching between external and internal regulators, except when using the SwitchOverRegulators API with a blocking call.

Table 37-110. Arguments if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:24]	0x16	Opcode
Bits [3:2]	0 – Force trim setting. The syscall will ignore regulator configuration and will set requested trims 1 – Deep Sleep Entry use case 2. Deep Sleep Exit use case 3. Reset Recovery use case	Use case
Bit [1]	0 - Internal regulator (LDO) 1 - REGHC	Operating mode
Bit [0]	0x1	Indicates that all the arguments are passed in IPC_DATA0 Register.

Table 37-111. Arguments if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x16	
Bits [3:2]	0 – Force trim setting. The syscall will ignore regulator configuration and will set requested trims 1 – Deep Sleep Entry use case 2. Deep Sleep Exit use case 3. Reset Recovery use case	Use case
Bit [1]	0 - Internal regulator (LDO) 1 - REGHC	Operating mode
Bit [0]	Not used.	

Table 37-112. Return if IPC_STRUCT.DATA[0] = 1

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code	See 37.5 System Call Status for details.

Table 37-113. Return if IPC_STRUCT.DATA[0] = 0

Address	Value to be Written	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details).
Bits [23:0]	Error code	See 37.5 System Call Status for details.

37.4.34 OpenRMA

This API enables full access to the device in the RMA life-cycle stage upon successful execution. The API returns the 0xF00000A9 failure code if there are any active embedded Flash operations. Users can trigger this API with DAP as Master after transitioning the device to RMA. In the RMA life-cycle stage, before successful OpenRMA execution, DAP will only have access via SYSTEM AP to IPC MMIOs and one-sixteenth of SRAM0. Only OpenRMA system call is allowed before successful OpenRMA execution.

Table 37-114. Parameters

Address	Value to be Written	Description
IPC_DATA0 Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH_ADDR		
Bits [31:24]	0x29	OpenRMA opcode
Bits [23:0]		Not used
SRAM_SCRATCH_ADDR + 0x4		
Bits [31:0]		Object size in bytes including itself. It should always be 20 bytes.
SRAM_SCRATCH_ADDR + 0x8		
Bits [31:0]	0x120029F0	Command ID
SRAM_SCRATCH_ADDR + 0xC		
Bits [31:0]		Unique ID word 0
SRAM_SCRATCH_ADDR + 0x10		
Bits [31:0]		Unique ID word 1
SRAM_SCRATCH_ADDR + 0x14		
Bits [31:0]		Unique ID word 2 (3 bytes)
SRAM_SCRATCH_ADDR + 0x18		
Bits [31:0]		SRAM address where signature is stored (4 bytes)

Table 37-115. Return

Address	Return Value	Description
SRAM_SCRATCH_ADDR		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see 37.5 System Call Status for details)
Bits [23:0]	Error if any	In case of fail, error code (see SROM API status codes)

37.5 System Call Status

At the end of every system call, a status code is written over the arguments in the IPC data structure or the SRAM address pointed by the IPC location. A success status is 0xAXXXXXXX, where X indicates don't care values or return data for system calls that return a value. A failure status is indicated by 0xF00000XX, where XX is the failure code.

If any address of SRAM_SCRATCH is protected, a failure status is indicated by 0xF00000F1.

Table 37-116. System Call Status

Status Code	Description
0xAXXXXXXX	Success – The X denotes a don't care value, which has a value of '0' returned by the SROM.
0xA0000009	Command in progress.
0xF0000001	Invalid protection state – This API is not available in current protection state.
0xF0000002	Invalid eFuse address.
0xF0000004	Wrong or out-of-bound flash address.
0xF0000005	Page is write protected.
0xF0000006	Client did not use its reserved IPC structure for invoking system call.
0xF0000008	Returned by all APIs when client does not have access to the region it is using to pass arguments.
0xF0000009	Command in progress. The code begins with "F" from fail. To be replaced by the next code in the future.
0xF000000A	Checksum of flash resulted in non-zero.
0xF000000B	The opcode is not a valid API opcode.
0xF000000E	Invalid address range.
0xF000000F	Invalid arguments passed to the API.
0xF0000010	Boot flash authentication failed
0xF0000011	Indicates that TEST_KEY_DFT_EN was set during boot up.
0xF0000012	Indicates that TST_KEY_SAFE_MODE was set during boot up.
0xF0000013	Invalid arguments location.
0xF0000015	Invalid trims length.
0xF0000016	Invalid HASH object.
0xF0000017	Number of zeros in the HASH computed by ROM boot and number of zeros stored in eFuse do not match.
0xF0000018	Invalid table of contents 1's CRC.
0xF000001A	Returned during secure boot if Sflash bank 1 authentication check fails
0xF0000020	Invalid table of contents 2's CRC.
0xF0000022	Returned when flash embedded operations are invoked during margin mode operation.
0xF0000080	Flash trim hunt failed for SONOS. For eCT magic number was not found in TOC1
0xF0000091	Program is called on a sector that is suspended from erase.
0xF0000092	EraseResume is called when no sector is suspended from erase.
0xF0000095	The requested system call is not approved by TEE.
0xF00000A0	FUR download fails with POR_NATIVE = 1.
0xF00000A1	FUR download fails due to ECC error.
0xF00000A2	IRAM download fails due to ECC error.
0xF00000A3	8051 software download fails due to ECC error.
0xF00000A4	ProgramRow is invoked on non-erased cells or blank check fails.
0xF00000A5	EraseSuspend when called with no ongoing erase operation.
0xF00000A6	ProgramRow when active erase operation is going on.
0xF00000A7	Embedded operation fails.
0xF00000A8	Invalid program width option is provided.

Table 37-116. System Call Status

Status Code	Description
0xF00000A9	WriteRow when invoked during an active embedded operation.
0xF00000AA	Returned by FLASH program/erase APIs when writes are disabled in safety register.
0xF00000AB	Returned by WriteRow when invoked in dual bank mode.
0xF00000B0	Invalid number of entries are passed.
0xF00000B1	Returned when WriteNormalAccessRestrict is called to restrict less.
0xF00000B2	Returned when WriteRow is called on invalid Sflash rows in NORMAL state.
0xF00000B3	Invalid unique ID is passed during RMA.
0xF00000B4	Invalid signature is found during RMA.
0xF00000B5	Invalid FACTORY_HASH.
0xF00000B8	Returned when more than 15 HASH objects are indicated in TOC1.
0xF00000B9	Returned when more than 15 HASH objects are indicated in TOC2.
0xF00000BA	Returned by TransitionRMA and OpenRMA when public key structure is invalid.
0xF00000BB	Backup sector is in erased state.
0xF00000BC	Returned during boot when SWPU in Sflash is more than expected.
0xF00000BD	Returned during boot when SWPU in Sflash is more than expected.
0xF00000BE	Returned during boot when SWPU in Sflash is more than expected.
0xF00000BF	Returned during boot when SWPU in Sflash is more than expected.
0xF00000C0	Returned during boot when SWPU in Sflash is more than expected.
0xF00000C1	Returned during boot when SWPU in Sflash is more than expected.
0xF00000C2	Returned by Read or WriteSWPU API when invalid ID is passed.
0xF00000C3	Returned by WriteSWPU API when client does not have access to update SWPU.
0xF00000C4	Returned by WriteSWPU API when client does not provide matching SL_ADDR and SL_SIZE.
0xF00000C5	Returned by ReadSWPU API if ECC error occurred during SRAM read operation.
0xF00000C6	Returned by Read and WriteSWPU API if the ID'd PU was rejected during boot due to overlap or out-of-order region.
0xF00000C7	Returned by Read and WriteSWPU APIs if there was a pending ECC error before performing SWPU operations.
0xF00000C8	Returned during boot if valid life-cycle fuse combinations are not read from eFuse.
0xF00000CB	Returned by BlowFuseBit API when read value from programmed fuse is 0.
0xF00000CF	User has provided arguments in protected region.
0xF00000D0	Address pointer fetched from TOC/patched syscall table is not in Sflash.
0xF00000D1	The bootrow is not zero in VIRGIN.
0xF00000D2	SRAM BIHR repair operation fails.
0xF00000D3	SRAM repair fuse redundancy check fails.
0xF00000D5	Returned when Sflash markers are corrupted during boot.
0xF00000D6	Returned by WriteRow when marker overflows by 2^32 times.
0xF00000D7	Returned by SoftReset API when CM7 reset is requested with CM7 not being in DeepSleep mode.
0xF00000D8	Invalid life cycle
0xF00000E0	REGHC is configured for Manual mode.
0xF00000E1	REGHC is currently in transition.
0xF00000E2	REGHC is already enabled.
0xF00000E3	Regulator is not configured with ConfigureRegulator().
0xF00000E4	Returned by SwitchOverRegulators() when the syscall is called with a different OpMode parameter than ConfigureRegulator().
0xF00000F0	HardFault occurs during bootup.

Table 37-116. System Call Status

Status Code	Description
0xF00000F1	HardFault occurs in context of system calls.
0xF4000000	Invalid programmable PPU access.
0xF5000000	Invalid fixed PPU access.
0xF6000019	Returned when bootrow fuse and MMIO do not match.
0xF6000029	Key in the bootrow mismatch.
0xF6000039	Returned when trim and trim inverse in bootrow are not equal.
0xF6000049	Returned when life-cycle fuses fail its redundancy check.
0xF6000059	Returned when invalid life-cycle fuse combinations are blown.
0xF1000000	Hash on Sflash trims failed. The computed hash is OR'd with this status.
0xF2000000	CRC8 of the eFuse group failed. The computed CRC is OR'd with this status.
0xF3000000	Returned during boot in IPC_STRUCT0.DATA1 if fault structure 0 valid bit is set. The LSBs will hold the fault ID information.

37.6 eFuse Memory

The eFuse memory consists of a set of eFuse bits. When an eFuse bit is programmed, or “blown”, its value can never be changed. Some of the eFuse bits are used to store various unchanging device parameters, including critical device factory trim settings, device life-cycle stages, DAP security settings, and encryption keys. Other eFuse bits are available for customer use.

37.6.1 Features

eFuses have the following features:

- A total of 1024 eFuse bits.
- The eFuse bits are programmed one at a time, in a manufacturing environment. The eFuse bits cannot be programmed in the field.
- Multiple eFuses can be read at the bit or byte level through an SROM call. An unblown eFuse reads as logic 0 and a blown eFuse reads as logic 1. There are no hardware connections from eFuse bits to elsewhere in the device.
- SROM system calls are available to program and read eFuses.

37.6.2 Customer eFuses

eFuses have bits available for customer use. They can be programmed in the NORMAL life cycle stage via CM0+ and CM7/ DAP, and in the SECURE protection state via CM0+ and CM7.

Offset	Width	Name
0x068	32	Customer Data

To program customer data, the Blow Fuse Bit system call must be called; the logic for calculation is:

macro Address = AddressOffset% EFUSE_NR

Byte Address = AddressOffset/EFUSE_NR

Where EFUSE_NR = number of eFuse macros (that is, number of columns) available for a product.

38. Flash Boot



Flash boot is the firmware that resides in SFlash, runs on the security processor (Cortex-M0+), and is executed after ROM boot has completed the basic hardware configuration and trim.

The purposes of flash boot are as follows:

- Initial configuration for a hardware subset
- Security configuration that must be done at programming context (PC) = 0
- Initialization of a debugger pin and the debug access port (DAP) subsystem
- Authentication for secure application
- Launching the application in a boot chain

Flash boot performs the following tasks:

- Configures the hardware that is not part of ROM boot
- Validates TOC2
- Sets up the CM0+ and peripheral clocks based on TOC2_FLAGS
- Enables system calls
- Configures SWD and JTAG pins and enables DAP
- Configures and enables a listen window for DAP
- Validates user applications structure
- Validates an RSA public key structure
- Authenticates secure applications by verifying their digital signature
- Sets a PC value – either PC = 2 or PC = 4
- Launches a bootloader for end-of-line programming, controlled through TOC2_FLAGS
- Launches a user application if there are no errors
- Enters NORMAL_DEAD or DEAD protection mode if an error occurs

38.1 Features

- Secure boot support
 - Digital signature verification by RSASSA-PKCS1-v1.5 with SHA-256 and RSA 2048 bit
 - Public key in SFlash for RSA 2048
 - Control enabling DAP by access restrictions (AR)
- User configuration through TOC2
 - The next launched application's address and format
 - A listen window to facilitate debugging
 - Boot time and power consumption
 - Authentication options for secure applications
- Embedded CAN and LIN bootloader to replace SWD or JTAG for factory programming
 - CAN at 100 or 500 kbps
 - LIN at 20000 or 115200 bps

38.2 Using Flash Boot

38.2.1 Flash Boot Shared Functions

Flash boot contains a few functions that may be executed from user code. [Table 38-1](#) provides memory locations for the function pointers and a short description for each function.

Table 38-1. Flash Boot Functions

Memory Location	Function Name	Comment
0x1700_2040	Cy_FB_VerifyApplication	Validates the application signature with RSASSA-PKCS1-v1.5 (2048 bit)
0x1700_2044	Cy_FB_IsValidKey	Validates the public key

38.2.1.1 Cy_FB_VerifyApplication

■ Function Description

Used in flash boot for authentication of the next application.

Can be used by the other code to authenticate RSASSA-PKCS1-v1.5 for any data; it need not be a signed application image.

■ Parameters

uint32_t address: The start address of the data image to be authenticated.

uint32_t length: The length of the data image.

uint32_t signature: The start address of the signature for the data image.

cy_stc_crypto_rsa_pub_key_t* publicKey: The pointer to a public key structure.

■ Return Value

uint32_t

0 - digital signature is invalid

1 - digital signature is valid

2. Check if the address of the Object Size member of a public key object is a valid location in the internal memory
3. Check if the Object Size value is within the allowed range [MIN, MAX]. MIN and MAX depend on the signature scheme and are implementation details
4. Check if the address of the last word in the public key object points to a valid location within the internal memory
5. Check if the Signature Scheme value of the public key object is valid. For Signature Scheme 0: RSASSA-PKCS1-v1.5 with RSA-2048 and SHA2-256
6. Check if the RSA public key exponent size is less than or equal to 32×8 bits. Check if the RSA public key module size is less than or equal to 256×8 bits
7. Check if the values of RSA public key module and exponent members of the public key structure are inside the memory region for the public key object
8. Validate the RSA optional coefficients (barretCoefPtr, inverseModuloPtr, rBarPtr). Their values should either be zero or the addresses inside the memory range of the public key object

38.2.1.2 Cy_FB_IsValidKey

■ Function Description

Checks whether the public key structure is valid.

Note: It may be used only for a public key that is referenced in TOC2.

■ Parameters

uint32_t address: The address of TOC2.

cy_stc_crypto_rsa_pub_key_t* publicKey: The pointer to public key structure.

■ Return Value

uint32_t:

0 - public key is invalid

1 - public key is valid

If any of the following steps results in false, the function returns 0, otherwise it returns 1:

1. Check if the address of a public key in TOC2 points to a valid location in the internal memory

38.2.2 Using a Bootloader

38.2.2.1 Bootloader Host Requirements

■ Bootloader Packet Structure

Figure 38-1 shows the structure of communication packets sent from the host to the MCU.

Figure 38-1. Bootloader Command Packet Structure





Figure 38-2 shows the structure of response packets sent from the MCU to the bootloader host.

Figure 38-2. Bootloader Response Packet Structure



All multi-byte fields are little endian.

Bootloader packet length is limited to four CAN messages, each with 8 bytes of data.

Bootloader packet length is limited to four LIN messages, each with up to 8 bytes of data.

Each CAN or LIN message may contain up to 8 bytes of user data, which hold bootloader command data.

Bootloader Commands

■ Enter Bootloader

Begins a bootload operation. All the other commands except Exit Bootloader are ignored until this command is received. Responds with device information and Bootloader SDK version.

Input

- Command Byte: 0x38
- Data Bytes:
4 bytes: Product ID. Internal bootloader requires Product ID = 0x01020304

Output

- Status/Error Codes:
Success
Error Command
Error Data, used for product ID mismatch
Error Length
Error Checksum
- Data Bytes:
4 bytes: Device JTAG ID
1 byte: Device revision
3 bytes: Bootloader SDK version

■ Sync Bootloader

Resets the bootloader to a known state, making it ready to accept a new command. Any buffered data is discarded. This command is needed only if the bootloader and the host are out of sync with each other.

Input

- Command Byte: 0x35
- Data Bytes: N/A

Output

N/A

This command is not acknowledged

■ Exit Bootloader

Exits from the bootloader and ends the bootload operation

After this command is received, the internal bootloader stops reading a bootloading communication, verifies the bootloadable application image, and launches it if it is valid.

Input

- Command Byte: 0x3B
- Data Bytes: N/A

Output

N/A

This command is not acknowledged

■ Send Data

Transfers a block of data to the bootloader. This data is buffered in anticipation of a Program Data command. If a sequence of multiple Send Data commands are sent, the data is appended to the previous block. This command is used to break up large data transfers into

smaller pieces, to prevent channel starvation in some communication protocols.

Input

- ❑ Command Byte: 0x37
- ❑ Data Bytes:
 - n bytes: Data to be appended to the buffer
 - CAN allows up to 25 bytes of data
 - LIN allows up to 21 bytes of data

Output

- ❑ Status/Error Codes:
 - Success
 - Error Command
 - Error Data
 - Error Length
 - Error Checksum
- ❑ Data Bytes: N/A
- Send Data Without Response

Same as the Send Data command, except that no response is generated by the bootloader. This reduces bootloading time.

Input

- ❑ Command Byte: 0x47
- ❑ Data Bytes:
 - n bytes: Data to be appended to the buffer

Output

N/A

This command is not acknowledged.

- Program Data

Writes data to one row of device internal flash or page of external nonvolatile memory (NVM). May follow a series of Send Data or Send Data Without Response commands.

Input

- ❑ Command Byte: 0x49
- ❑ Data Bytes:
 - 4 bytes: Address. Must be within the correct memory address space, and 32-bit aligned
 - 4 bytes: CRC-32C of the entire data buffer to be written

Note: The buffer includes data that is already appended to it with Send Data or Send Data without Response commands that precede Program Data.

 - n bytes: Data to write into the flash row or external NVM page.

Output

- ❑ Status/Error Codes:
 - Success
 - Error Command
 - Error Data
 - Error Length
 - Error Checksum

Error Flash Row
Error Flash Row Access

- ❑ Data Bytes: N/A

■ Verify Application

Reports whether the checksum for the bootloadable application image is valid.

Input

- ❑ Command Byte: 0x31
- ❑ Data Bytes:
 - 1 byte: App ID of the application to be verified. Must be the same value as in Set Application Metadata command.

Output

- ❑ Status/Error Codes:
 - Success - returned when either the application is valid
 - Error Command
 - Error Data
 - Error Length
 - Error Checksum
 - Error Flash Row Access
- ❑ Data Bytes:
 - 1 byte: 1/0 for application is valid or not valid

■ Set Application Metadata

This command is used to set a given application's metadata.

It must be the second bootloader command, which the bootloader host delivers to the MCU, the first one is Enter Bootloader.

Input

- ❑ Command Byte: 0x4C
- ❑ Data Bytes:
 - 1 byte: App ID
 - 4 bytes: Bootloadable Application start address
 - 4 bytes: Bootloadable Application size in bytes

Output

- ❑ Status/Error Codes:
 - Success
 - Error Command
 - Error Length
 - Error Data
 - Error Checksum
 - Error Flash Row Access
- ❑ Data Bytes: N/A

Data Constraints

App ID may have the following values:

Table 38-2. Data Constraints

App ID Value	Description
0	For either LIN at 20 kbps or CAN.
1	For LIN at 115.2 kbps with a Fast Mode. See Switching between Normal and Fast modes .
2	For LIN at 115.2 kbps without a Fast Mode.

Bootloadable application start address must be within a valid RAM memory length - [RAM_START + 512, RAM_END – 4096].

Bootloadable application length must be a value for which the bootloadable application image fits into a RAM address range [RAM_START + 512, RAM_END – 4096].

38.2.2.2 Using CAN or LIN

Bootloader polls for Enter Bootloader command on CAN and LIN pins as follows:

1. Bootloader polls for CAN messages at 100 kbps; if no valid CAN message with Enter Bootloader command is received during 10 ms, it goes to (2). If a valid command is received, bootloader continues using CAN at 100 kbps for the next bootloader commands.

2. Bootloader polls CAN at 500 kbps for a duration of 10 ms. If no valid Enter Bootloader command is received it goes to (3).
3. Bootloader polls LIN at 20 kbps for a duration of 150 ms. If no valid Enter Bootloader command is received it goes to (1).
 - a. If a valid command is received and the next bootloader command is Set Application Metadata, and Set Application Metadata bootloader command has App ID = 1, then bootloader sends an OK response to the bootloading host. It then reconfigures LIN to 115200 bps and waits for the next bootloader command to use this new baud rate.
4. If bootloading has started on CAN or LIN, but later communication has stopped, the bootloader uses a timer, which detects that there was no bootloader communication for two seconds and resets the communication configuration. It then goes to (1), (2), or (3) depending on the communication channel for which the bootloading has failed.
5. If there are no valid Enter Bootloader command during 300 seconds, the bootloader stops. The device goes into Sleep power mode.

The following figures show a few examples of CAN and LIN bootloading communication.

Figure 38-3. Polling CAN and LIN with No Bootloader Commands




Figure 38-4. Successful Bootloading on CAN 500 kbps




Figure 38-5. An Example of a Failed Bootloading




Figure 38-6. Bootloading on LIN at 20 kbps




Figure 38-7. Bootloading on LIN at 115200 bps





Figure 38-8 shows a simplified logic to switch between CAN and LIN bootloading interfaces; a detailed logic is provided in the numbered list at the beginning of this section.

Figure 38-8. Bootloading Switching between CAN and LIN



(5) A CAN polling timeout is 10 ms at 100 kbps and 10 ms at 500 kbps.

(9) A LIN polling timeout is 150 ms at 20 kbps.

(11) An overall bootloading timeout is 300 seconds from the end of the last successful received bootloading command, or from the start of the bootloading if no commands have been received.

(12) A communication timed out flag is set if no valid bootloading command has been received for 2 seconds.

(13) LIN by default is configured at 20 kbps. But if the Send App Metadata bootloading command is received with AppID=1, then LIN is reconfigured to 115200 kbps. See [Switching between Normal and Fast modes](#).

38.2.2.3 CAN Configuration

See device datasheet for CAN configuration details.

■ CAN driver limitations

The CAN specification states the clock accuracy should be at most $\pm 0.5\%$ at 500 kbps. Internal generator (IMO) for Traveo II devices does not meet this accuracy, because it is specified to have frequency error up to $\pm 1.0\%$.

However, the CAN block may use SJW (ReSynchronisation Jump Width) to adjust CAN clock to the baud rate, which allows $\pm 1.0\%$ clock tolerance.


It is recommended to use a single point-to-point connection and have the wire length within the allowed range for CAN 500 kbps to have a stable communication at $\pm 1.0\%$ clock frequency tolerance.

38.2.2.4 LIN Configuration

See device datasheet for LIN configuration details

- **Switching between Normal and Fast modes**
 Some manufacturers of LIN transceivers allow “Fast mode” or “Flash mode”, which is used mainly for bootloading. Fast mode allows LIN communication speed to be increased to 115200 bps. A special sequence of signals on EN and TX pins of the LIN transceiver switches it to the Fast mode.

Figure 38-9. Switching to LIN Fast Mode



The T1, T2, and T3 limits may vary from manufacture to manufacturer. Flash boot uses average values:

$$T1 = T2 = T3 = 12 \mu s$$

Switching from the Fast mode to the Normal mode is done by applying the same sequence on EN and TX pins.

38.3 Flash Boot Internals

38.3.1 Definitions

- **Firmware Image**
 A specific format for a firmware module stored in the internal memory. See [Application Formats on page 754](#) for more details.
- **ROM Boot**
 ROM code stored at the device address range that starts at address 0x0000_0000. The first code is executed when the device is powered on. This is the first phase of the boot process.
- **Flash Boot**
 A firmware image stored in SFlash that provides code for the second phase of the boot process.
- **Secure Boot**
 Secure boot is the term used to include the entire secure chain of trust boot process. It includes ROM boot, flash boot, and optionally secure image.
- **DAP**
 Debug Access Port
- **SFlash**
 Supervisory flash is a dedicated flash region used by Cypress to store manufacturing information, hardware trim and wounding information, special user sections, TOC, and code for the second phase of the boot process and flash boot.
- **TOC, TOC1, and TOC2**
 Table of Contents. This table is broken up into two parts. The first part (TOC1) includes addresses of items frozen in the factory, such as items that are included in the FACTORY_HASH calculation and cannot be changed by the user. The second part (TOC2) includes addresses of the user application, public key, and other user configurable items that are used by secure boot. Entities from TOC1 and TOC2 are used to calculate SECURE_HASH.
- **Secure Application**
 An application in Cypress Secure Application Format (CySAF). This application contains digital signature and may be authenticated using RSASSA-PKCS1-v1.5 with RSA 2048 and SHA-256.
- **FACTORY_HASH**
 128 most significant bits of the SHA-256 hash value computed to authenticate objects frozen in the factory.
- **Private Key**
 A private key for RSA 2048 to sign the digital signature of the secure application.
- **Public Key**
 A public key for RSA 2048 to verify the digital signature of the secure application.
- **RSA**
 An asymmetric crypto algorithm by Rivest-Shamir-Adleman.
- **RSASSA-PKCS1-v1.5**
 A digital data authentication algorithm based on RSA and hashing functions.
- **SECURE_HASH**
 128 most significant bits of the SHA-256 hash value used to authenticate flash boot and public key in SECURE and SECURE_WITH_DEBUG life-cycle stages. When creating SECURE_HASH, factory frozen objects are authenticated using FACTORY_HASH. This makes sure that the flash boot authenticated by SECURE_HASH later is the same as the one created in the factory. SECURE_HASH is computed just before transition to SECURE, so public key needs to be known only then; the OEM provides the public key.
- **SHA2, SHA3, SHA-256, and SHAKE-128**
 SHA-based secure hash functions.

38.3.2 SFlash Address Mapping

The entire flash boot is located in SFlash. It starts at address 0x17001C00 and ends at 0x170063FF. You cannot overwrite or change the flash boot. The flash boot version can be read from address 0x17002018, which has an unsigned 32-bit integer value.

The area from 0x17000800 to 0x17000F00 can be used for user applications, and storing keys and other information. In this area, it is possible to write only in the Normal mode.

A public key is located at address 0x17006400. The maximum length is 3072 kbytes.

TOC2 is located near the end of SFlash at 0x17007C00. Both public key and TOC2 are available for write in Normal mode only.

Figure 38-10. SFlash Address Mapping




Figure 38-11 shows the flash boot program flow. The entry point to flash boot must be at a fixed offset inside the SFlash block. ROM boot code will transfer control to flash boot after its tasks are completed and SFlash is validated. Each section of the flow chart is labeled with an index number (n), which is used for reference in the next sections.

```

graph TD
    subgraph "Initial Boot Sequence"
        B1((1) From ROM boot) --> B2[Set-up SP]
        B2 --> B3[Initialization]
        B3 --> B5[Validate TOC2]
        B5 --> B6{Is TOC2 valid?}
        B6 -- No --> B26a((26) Branch DEAD)
        B6 -- Yes --> B8{Do Bootloading?}
        B8 -- Yes --> B28a((28) Branch Bootloader)
        B8 -- No --> B9[Get App #0 ResetHandler]
        B9 --> B10{Is ResetHandler valid?}
        B10 -- No --> B26b((26) Branch DEAD)
        B10 -- Yes --> B11{Authenticate App?}
        B11 -- No --> B26b
        B11 -- Yes --> B12{Is Public Key valid?}
        B12 -- Yes --> B26c((26) Branch DEAD)
        B12 -- No --> B13{Is Digital Signature valid?}
        B13 -- No --> B26b
        B13 -- Yes --> B14[Enable System Calls]
        B14 --> B15[Set-up DAP from AR]
        B15 --> B16[Set PC=2 or 4]
    end


    subgraph "Interrupt and System calls"
        B29[Interrupt and System calls] --> B9b[Get App #1 ResetHandler]
        B9b --> B10b{Is ResetHandler valid?}
        B10b -- No --> B26d((26) Branch DEAD)
        B10b -- Yes --> B11b{Authenticate App?}
        B11b -- No --> B26b
        B11b -- Yes --> B12b{Is Public Key valid?}
        B12b -- Yes --> B26e((26) Branch DEAD)
        B12b -- No --> B13b{Is Digital Signature valid?}
        B13b -- No --> B26f((26) Branch DEAD)
        B13b -- Yes --> B14b[Enable System Calls]
        B14b --> B15b[Set-up DAP from AR]
        B15b --> B16b[Set PC=2 or 4]
    end

    subgraph "Branch DEAD Sequence"
        B30[Branch DEAD] --> B31[Set Error Code]
        B31 --> B32{PROTECTION = VIRGIN?}
        B32 -- Yes --> B2a[Set-up SP]
        B2a --> B25a((25) Idle Loop)
        B32 -- No --> B33{LifeCycle = SECURE?}
        B33 -- Yes --> B33a[PROTECTION=DEAD]
        B33a --> B34a[Deploy SECURE_DEAD Access Restrictions]
        B34a --> B34b[Enable System Calls]
        B34b --> B35a[Set-up DAP from AR]
        B35a --> B35b[Apply System Protection]
        B35b --> B36a[Set PC=4]
        B33 -- No --> B33b[Do not change PROTECTION]
        B33b --> B34b[Deploy NORMAL_DEAD Access Restrictions]
        B34b --> B34c[Enable System Calls]
        B34c --> B35b[Set-up DAP from AR]
        B35b --> B35c[Apply System Protection]
        B35c --> B36b[Set PC=2 or 4]
        B36a --> B17{Is DAP enabled?}
        B36b --> B17
        B17 -- No --> B18[Configure SWJ]
        B17 -- Yes --> B18
        B18 --> B2b[Set-up SP]
        B2b --> B25b((25) Idle loop)
    end

    subgraph "Single-core and CM0+ Application Launch"
        B22{Is Single-core?} -- No --> B24[Launch CM0+ application]
        B22 -- Yes --> B23[Launch CM7 application]
    end

    B16 --> B17
    B16b --> B17
    B17 --> B18
    B18 --> B19{Wake-up from Hibernate?}
    B19 -- Yes --> B22
    B19 -- No --> B20[Listen window default is 20 ms]
    B20 --> B21{Test Mode?}
    B21 -- Yes --> B2b
    B21 -- No --> B22

```



38.3.3.1 Entry from ROM Boot (1)

At this stage ROM boot has finished its tasks and transfers the execution for CM0+ code to flash boot.

38.3.3.2 Set-up SP (2)

The same flash boot image is programmed in all Traveo II devices. Within the device family different devices have different sizes of SRAM. The SP register value for flash boot must be at the top of user SRAM. Thus, it is impossible to know the SP value at build time.

At the start of flash boot, the SP register value is 0. Flash boot calculates and sets the value of SP register at runtime.

38.3.3.3 Initialization (3)

This function executes a hardware-specific initialization code.

All Traveo II devices support SRAM ECC, and flash boot initializes the SRAM it uses for stack.

Flash boot enables S&H mode for SRSS.PWL_CTL2.BGREF_LPMODE.

38.3.3.4 Validate TOC2 (5)

The current procedure to validate TOC2:

- SFLASH_TOC2_OBJECT_SIZE <= 512 for Traveo II
- SFLASH_TOC2_OBJECT_SIZE >= 8
- SFLASH_TOC2_MAGIC_NUMBER == 0x01211220

If all the conditions above are true TOC2 state is VALID. TOC2 structure described in [TOC2 Structure on page 758](#).

38.3.3.5 Is TOC2 Valid (6)

TOC2 may be in three states:

- VALID, TOC2 structure, and CRC are valid
- ERASED, the first two 32-bit words at the start of TOC2 are equal to the SFlash erase value and protection mode is either VIRGIN or NORMAL.
For eCT SFlash the erased value is 0xFFFF_FFFF.
- CORRUPTED, when both ERASE nor VALID condition are false

If TOC2 state is ERASED then flash boot uses the default values for all the TOC2 elements instead of reading them from TOC2. A list of TOC2 elements for which the default values are used:

- SFLASH_TOC2_FIRST_USER_APP_ADDR is 0x1000_0000 (the start of flash)
- SFLASH_TOC2_FIRST_USER_APP_FORMAT is 0 (Basic Application Format)
- SFLASH_TOC2_FLAGS
 - 0x0000_0242 for Traveo II devices

The other TOC2 entries are not used when TOC2 state is ERASED.

38.3.3.6 Bootloading (8)

Bootloading triggers in VIRGIN and NORMAL protection modes if the following conditions are met:

- CPUSS.PROTECTION != SECURE
- TOC2 state is either ERASED or VALID and SFLASH.TOC2_FLAGS bit FB_BOOTLOADER_DISABLE is zero.
- The first two 32-bit words at 0x1000_0000 are equal to 0xFFFF_FFFF.

38.3.3.7 Get App #{0, 1} Reset Handler (9)

This step may be executed when TOC2 state is either VALID or ERASED (see [38.3.3.6 Bootloading \(8\)](#)).

If TOC2 state is ERASED and CPUSS.PROTECTION = NORMAL then:

1. Application start address is 0x1000_0000
2. Application format is CyBAF
3. Second application is ignored; thus, if a validation of the first application leads to an error, the second application is not validated and DEAD branch is executed.

Otherwise, TOC2 state is VALID and application parameters are calculated as shown in this section.

Flash boot reads application start address from TOC2 entries:

- SFLASH_TOC2_FIRST_USER_APP_ADDR for App#0
- SFLASH_TOC2_SECOND_USER_APP_ADDR for App#1

The application format is stored in TOC2 entries:

- SFLASH_TOC2_FIRST_USER_APP_FORMAT for App #0
- SFLASH_TOCR_SECOND_USER_APP_FORMAT for App #1

The reset handler address inside the application depends on the application format. See [Application Formats on page 754](#).

38.3.3.8 Valid Reset Handler (10)

Flash boot checks whether the address of the reset handler for the user application is inside a valid range.

The valid range is the following: SRAM, SFlash, code flash, and work flash.

Note: This check is made to prevent the HardFault that would otherwise occur if the reset handler for the user application points to an invalid memory location.

38.3.3.9 App Authentication (11)

Flash boot optionally authenticates a digital signature for the application image based on the value of TOC2_FLAGS bits APP_AUTH_DISABLE.

38.3.3.10 Is Public Key Valid (12)

The public key structure is filled by the user. It must be validated to ensure the correctness of the entries before being used.

The validation is done using Cy_FB_IsValidKey() function described in [Cy_FB_IsValidKey on page 741](#).

38.3.3.11 Valid Digital Signature (13)

The application to be launched by flash boot may be authenticated using a digital signature verification with RSASSA-PKCS1-v1_5 (RSA-2048, SHA2-256) see details in RFC3447.

The public key used for this operation may be placed in the User Public Key area of SFlash or in another internal flash; its pointer should be updated in TOC2_SIGNATURE_VERIF_KEY (See [Table 38-7](#)). The format of the key is shown in [RSA Public Key Format on page 757](#).

The application to be authenticated should be in the CySAF or Customer Formats. Applications in CyBAF may not be authenticated. The application formats are described in [Application Formats on page 754](#).

38.3.3.12 Enable System Calls (14)

The system calls are enabled in this stage. System calls are functions such as writing to code flash. The function calls are performed via the IPC communication to the CM0+ NMI interrupt. The SROM function EnableSystemCalls() is called to enable these system calls.

38.3.3.13 Set up DAP from AR (15)

Enable or disable DAP based on the current AR.

ROM boot function GetAccessRestrictStruct() is called to determine which APs (access points, one of CM0+ AP, CM7_0/CM7_1, and TC) are enabled. Based on this information the proper values are written to the CPUSS_AP_CTL register.

Note: For a VIRGIN protection mode the DAP setup is performed in the ROM boot; thus Flash boot skips this step.

38.3.3.14 Set PC (16)

ROM boot and Flash boot are being executed in PC = 0, system calls are executed in PC = 1, all other code must be executed in PC = 2 or higher. Flash boot is responsible for setting PC = 2 before launching a user application or jumping into the idle loop.

Flash boot sets PC = 4 in SECURE_DEAD branch (when the life-cycle stage is SECURE and protection mode is DEAD).

38.3.3.15 Is DAP Enabled (17)

For DEAD and Bootloader branches the result of step (17) is TRUE if CPUSS.AP_CTL enables DAP.

For a common branch (which ends by launching a user application), there is an additional check to determine if DAP is enabled. This check reads the TOC2_FLAG bit SWJ_PIN_CTL. If TOC2_FLAGS.SWJ_PINS_CTL is set and CPUSS.AP_CTL enables DAP then the result of step (17) is TRUE.

38.3.3.16 Configure SWJ (18)

Flash boot uses the ConfigureSWJ() ROM boot function if it is implemented for the device family; otherwise, this function is implemented in the Flash boot code base.

This function configures the GPIO pins to work in SWJ mode. All JTAG pins must be configured, except the TRST pin for the devices that have a problem.

38.3.3.17 Wake-up from Hibernate (19)

If the reason for a reset was "wake from Hibernate", skip the wait window and test mode check.

38.3.3.18 Listen Window (20)

The CPU delays execution for a period of time to allow the debug hardware to acquire the CM0+. The default is 20 ms, but other delay options may be set. If the Listen window is not required, the user may set the Listen window timeout to 0 ms.

This delay allows the debug hardware to acquire the debug interface before any user code is executed; it helps recovering the device in which a user code re-purposes SWJ pins.

38.3.3.19 Test Mode (21)

After the listen window delay, the firmware checks if the SRSS_TST_MODE register has either TEST_MODE or TEST_KEY_DFT_EN bit set. If either bit is set, execution is transferred to an endless loop in SROM. This is done by calling the ROM boot function BusyWaitLoop().

Some programmers use Listen window and set a Test mode bit to perform either programming or debugging tasks.

38.3.3.20 Is Single-Core (22)

Detects if a MCU is a single core. A single-core MCU does not allow a user code to be executed on CM0+.

Flash boot determines if MCU is a single-core by reading SFLASH_SINGLE_CORE_WOUND.

38.3.3.21 Launch CM0+ Application (24)

The procedure to launch a user application is as follows:

1. Set CPUSS_CM7_VECTOR_TABLE_BASE to 0xFFFF_0000 **Note:** For CM7 cores, the CPUSS_CM7_0/CM7_1_VECTOR_TABLE_BASE register should not be touched by Flash boot.
2. Set CPUSS_CM0_VECTOR_TABLE_BASE to the start of the user application interrupts vector table
3. Perform a CM0+ core reset
4. After a core-reset is performed a ROM boot is launched (on CM0+)
5. ROM boot checks if CPUSS_PROTECTION != 0, which means ROM boot is launched on CM0+ after a core-reset
6. If (5) is true, ROM boot sets SP and PC register values from the user interrupt vector table. The address of a user application interrupt vector table is stored at step (1) to CPUSS_CM0_VECTOR_TABLE_BASE
7. When ROM boot sets PC register value with the user reset handler address, user code starts executing

38.3.3.22 Idle Loop (25)

The SP register value is saved to R8 before entering Idle loop. Then SP is updated per step (2) in the [Set-up SP \(2\) on page 750](#).

Then CM0+ core is placed into a sleep power mode by calling ROM boot function BusyWaitLoop().

Note: Any interrupt (IPC system call or another interrupt source) may wake the device; therefore, the AR and other security settings should be properly configured by the user for each life-cycle stage.

38.3.3.23 Branch DEAD (8)

Flash boot goes into a DEAD branch if it detects any error. The list of the required errors is provided in [Set Error Code \(30\) on page 753](#).

38.3.3.24 Branch Bootloader (28)

If the bootloader feature is enabled for the device family and the bootloader launch condition is triggered, then flash boot launches a bootloader by going into this branch. The implementation may implement this branch either as a function call or as launching an application.

38.3.3.25 Interrupts and System Calls (29)

Flash boot should support patching the system calls using the system call patch table. Flash boot may patch a HardFault handler and re-configure the CM0+ interrupts during the startup.

38.3.3.26 Set Error Code (30)

Flash boot sets an error code into IPC_STRUCT[2].DATA0 register.

Table 38-3. Error Code

Error Name	Value	Description
CY_FB_STATUS_SUCCESS	0xA100_0100	Success status value.
CY_FB_STATUS_BUSY_WAIT_LOOP	0xA100_0101	Debugger probe acquired the device in Test Mode. The flash boot to entered a busy wait loop.
CY_FB_STATUS_BOOTLOADING	0xA100_0101	Bootloading in progress
CY_FB_STATUS_BTLD_OK	0xA100_0102	Bootloading finished successfully
CY_FB_ERROR_INVALID_APP_SIGN	0xF100_0100	App signature validation failed for the device families where flash boot launches only one application from TOC2. Either app structure or a digital signature is invalid for the device families for which Flash boot may launch either of two apps in TOC2.
CY_FB_ERROR_INVALID_TOC	0xF100_0101	Empty or Invalid TOC
CY_FB_ERROR_INVALID_KEY	0xF100_0102	Invalid Public Key
CY_FB_ERROR_UNREACHABLE	0xF100_0103	Unreachable Code
CY_FB_ERROR_TOC_DATA_CLOCK	0xF100_0104	TOC contains invalid CM0+ clock attribute.
CY_FB_ERROR_TOC_DATA_DELAY	0xF100_0105	TOC contains invalid listen window delay
CY_FB_ERROR_FLL_CONFIG	0xF100_0106	FLL configuration failed
CY_FB_ERROR_INVALID_APP0_DATA	0xF100_0107	App structure is invalid, for the device families where flash boot may launch only one app from TOC2.
CY_FB_ERROR_CRYPT0	0xF100_0108	Error in Crypto operation
CY_FB_ERROR_INVALID_PARAM	0xF100_0109	Invalid parameter value.
CY_FB_ERROR_UNEXPECTED_INTERRUPT	0xF100_010B	Any unexpected interrupt had happened in the Flash boot
CY_FB_ERROR_BOOTLOADER	0xF100_0140	Any bootloader error
CY_FB_ERROR_BOOT_LIN_INIT	0xF100_0141	Bootloader error, LIN initialization failed
CY_FB_ERROR_BOOT_LIN_SET_CMD	0xF100_0142	Bootloader error, LinSetCmd() failed
CY_FB_ERROR_BOOT_CAN_INIT	0xF100_0143	Bootloader error, CAN initialization failure
CY_FB_ERROR_BOOT_SECURE	0xF100_0144	Bootloader launched while CPUSS.PROTECTION=SECURE

38.3.3.27 *PROTECTION = VIRGIN (31)*

CPUSS_PROTECTION MMIO register value is being compared to the wished protection mode.

38.3.3.28 *LifeCycle = SECURE (32)*

A life-cycle stage is stored in eFuse. The life-cycle stage is not the same as protection mode. In this case, SECURE_WITH_DEBUG life-cycle stage is not equal to SECURE life-cycle stage, but for both, the protection mode equals to SECURE.

38.3.3.29 *PROTECTION = DEAD (33)*

Flash boot sets CPUSS_PROTECTION to DEAD in the DEAD branch only for SECURE life-cycle stage. For SECURE_WITH_DEBUG, NORMAL, and other life-cycle stages flash boot keeps the existing protection mode.

38.3.3.30 *Deploy AR (34)*

Flash boot deploys the AR that applies to the new protection mode.

NORMAL_DEAD AR are applied when entering DEAD branch from NORMAL, NORMAL_PROVISIONED, or SECURE_WITH_DEBUG life-cycle stages.

SECURE_DEAD AR are applied in the case of entering DEAD branch from SECURE life-cycle stage.

Assess restrictions are applied by calling ROM boot function RestrictAccess().

38.3.3.31 *Apply System Protection (35)*

System protection settings are applied. Usually they are applied by ROM boot code before entering flash boot. In the case of DEAD branch, the system protection settings may be changed and flash boot needs to call the ROM boot function ApplyProtectionSettings() to reconfigure them.

38.3.3.32 *Disable WDT (40)*

Bootloader may run for a significantly longer time then WDT timeout. Therefore, WDT must be either periodically reset or disabled at the start of bootloader.

38.3.3.33 *Set VECTOR_TABLE_BASE (41)*

CPUSS.CM0_VECTOR_TABLE_BASE and CPUSS.CM7_0/CM7_1_VECTOR_TABLE_BASE registers must be set to 0xFFFF_0000 value to show the debugger that no user application is running.

38.3.3.34 *Launch Bootloader (42)*

A bootloader firmware is launched. This firmware is a part of flash boot and launching it may be as simple as calling a function.

38.3.3.35 *CM0+ core reset (50)*

Flash boot resets CM0+ core by writing to the CPUSS_CM0_CTL register.

After the write, CM0+ must enter a sleep mode and wait until the core is reset.

38.3.3.36 *Launch a User App by ROM Boot (51)*

Flash boot performs the following to switch to the user application on CM0+:

1. Flash boot sets CPUSS.CM0_VECTOR_TABLE_BASE value with an address of the user CM0+ interrupt vector table.
2. Flash boot performs CM0+ core reset.
3. ROM boot starts up, tests if CPUSS.PROTECTION <> 0 which means ROM boot is launched from a software reset.
4. If (3) succeeds, ROM boot sets SP and PC register values from the users interrupt vector table, which is read out from CPUSS.CM0_VECTOR_TABLE_BASE.
5. When ROM boot sets the PC register value with the user's reset handler address, a user code starts executing.

38.3.4 Data Structures

38.3.4.1 *Application Formats*


Basic Application Format (CyBAF)

This is the most basic format and requires the least complicated setup and support. TOC2 can be left with default values, or in the ERASED state.

Note: CyBAF can be used only in VIRGIN and NORMAL protection modes. SECURE protection mode requires the format of the application, which is to be launched by Flash boot, to be CySAF.

Code flash and SRAM are divided into two parts, one for CM0+, the other for CM7_0/CM7_1 application. The CM0+ vector table usually starts at the beginning of code flash with the application code and the data following it, as shown in [Figure 38-13](#).

Figure 38-12. Basic Application Format Structure



The CM7_0/CM7_1 vector table is located at the lowest address location in the CM7_0/CM7_1 part of the code flash. The fraction of code flash and SRAM allocated for CM0+ and CM7_0/CM7_1 are controlled by the user.

The code does not have headers or footers and no predefined location for a digital signature for application validation. The user must validate the code if needed and have the CM0+ start up the CM7_0/CM7_1 CPU when ready.

If the application start address needs to be different from the start of code flash, update the TOC2 member SFLASH_TOC2_FIRST_USER_APP_ADDR with the new application start address.

Note: The Arm specification requires CM0+ interrupt vector table to be 256-byte aligned and CM7_0/CM7_1 interrupt vector table to be 1024-byte aligned.


Cypress Secure Application Format (CySAF)

This format is used for secure systems where the application code is authenticated using RSASSA-PKCS1-v1.5. Flash boot launches the application in the CySAF if:

1. SFLASH_TOC2_FIRST_USER_APP_ADDR points to a valid memory address
2. SFLASH_TOC2_FIRST_USER_APP_FORMAT value is 1, which means the app format is CySAF
3. The same as (1) and (2) but with SFLASH_TOC2_SECOND_USER_APP_ADDR, and SFLASH_TOC2_SECOND_USER_APP_FORMAT

The structure of CySAF is shown on [Figure 38-13](#).

Figure 38-13. CySAF Structure



CySAF consists of the following entries:

- **Application Header:** The header for the flash image is shown in [Table 38-4](#).

Table 38-4. Application Header

Offset	Size	Item	Description
0x0	4 bytes	Object Size	The size in bytes of an area for a digital signature authentication
0x4	4 bytes	Application ID/Version	Identifies the type of the app image. Flash boot does not use this value, however, the Cypress applications have the following values: Bit 31 - 16: Application Version Bit 15 - 0: Application ID 0x8001 - Flash boot 0x8002 - Security Image
0x8	4 bytes	Attribute	Reserved for future use
0xC	4 bytes	Number of Cores(N)	Number of MCU cores used by the application. Flash boot does not use this value.
0x10 + (4*i)	4 bytes	Core[i] VT offset	Offset to the interrupts vector table for Core[i]. An absolute address for Core[i] interrupt vector table is calculated as: Application start address + 0x10 + (4*i) + value of Core(i) VT offset. Flash boot does not use this information for i>0 and always launches the reset handler for Core[0].
0x10 + (4*N) + (4*i)	4 bytes	Core[i] CPU ID and Core Index	Customer assigned CPU ID and Core index. Bit 31 - 20: CPU ID. This is the part number value from the CPUID [15:4] register in an Arm device. Bit 7 - 0: Core Index The core index is used to distinguish between multiple cores within the system. The TVII-B-H system consists of a CM0+ and two CM7 cores. The CM0+ core is identified by a CPUID of '0xC60' and a core index of '0'; the first CM7 (CM7_0) core is identified by a CPUID of '0xC24' and a core index of '0' and the second CM7 (CM7_1) core has a CPUID of '0xC24' and a core index of '1'. Flash boot does not use this information.

■ MCU Code Segment

Each flash image in CySAF may contain one or more MCU code segments. At least one MCU code segment is required for the main MCU to be launched, this is usually CM0+. Flash boot requires the application address in SFLASH_TOC2_FIRST_USER_APP_ADDR and SFLASH_TOC2_SECOND_USER_APP_ADDR to have the first MCU code segment for CM0+ application.

Table 38-5. MCU Code Segment

Absolute address	Item	Description
App start address + 0x10 + (4*i) + Core(i) VT offset	Interrupts Vector Table for Core[i]	An offset to an Interrupts Vector Table for Core[i]
-	Core[i] Code and data	Code and Data for the Core(i) code segment

■ Application Footer


The footer of the application in CySAF contains the signature for authentication. Flash boot authenticates the application; it launches using RSASSA-PKCS1-v1.5 with RSA 2048 bit and SHA-256.

38.3.4.2 RSA Public Key Format

The RSA public key may be stored anywhere in the internal flash. For convenience, SFlash contains a region allocated for user data where RSA public key is assumed to be placed by default.

Figure 38-14 shows the structure of the RSA public key object used for signature authentication. The “Signature Scheme” (specified in TOC) defines the structure of the key. Figure 38-14 shows the key structure for RSASSA0PKCS1-v1_5.

Figure 38-14. Public Key Format



* Modulus, Exponent, K1, K2, and K3 must be 32-bit aligned, the data is little endian.

Table 38-6. Public Key Format

Public Key Object Member Name	Description
Object Size	A size in bytes used in SECURE_HASH calculation for a Public Key data protection.
Signature Scheme	A signature scheme. 0 - RSASSA-PKCS1-v1.5 with RSA-2048 and SHA-256 Other values are reserved.
Pointer to Modulus	A pointer to an RSA public key modulus data.
Length in bits of Modulus	A length in bits of an RSA public key modulus.
Pointer to Exponent	A pointer to an RSA public key exponent data.
Length in bits of Exponent	A length in bits of an RSA public key exponent data.
Pointer to K1	A pointer to an optional RSA public key coefficient, named Barrett coefficient. Or NULL if not present.
Pointer to K2	A pointer to an optional RSA public key coefficient, named inverse modulus. Or NULL if not present.
Pointer to K3	A pointer to an optional RSA public key coefficient, named rBarr coefficient. Or NULL if not present.

Note: Pointers to K1, K2, and K3 coefficients are optional. When they are set to NULL, flash boot will calculate the value for these coefficients at run time. Providing pre-calculated values for K1, K2, and K3 in the RSA public key object speeds up an RSA calculation up to three times.

The public key structure format and the public key object data is designed to be compatible with SDL functions for RSA operations (struct name `cy_stc_crypto_rsa_pub_key_t`).

38.3.4.3 TOC2 Structure

TOC2 is a structure stored in SFlash, which is used to configure flash boot and ROM boot firmware. It contains reference to the SMIF configuration structure used by the programming tools.

Table 38-7. TOC2 Structure

Offset	Name	Purpose
0x00	TOC2_OBJECT_SIZE	Object size in bytes starting from offset 0x00 until the last entry in TOC2.
0x04	TOC2_MAGIC_NUMBER	Magic number (0x01211220)
0x08	TOC2_SMIF_CFG_STRUCT_ADDR	Null terminated table of pointers representing the SMIF configuration structure.
0x0C	TOC2_FIRST_USER_APP_ADDR	Address of CM0+ First User Application Object (such as HSM in Traveo II)
0x10	TOC2_FIRST_USER_APP_FORMAT	First Application Object Format. 0 - Basic 1 - Cypress standard 2 - Simplified
0x14	TOC2_SECOND_USER_APP_ADDR	Address of CM0+ Second User Application Object (0's if none)
0x18	TOC2_SECOND_USER_APP_FORMAT	Second Application Object Format 0 - Basic 1 - Cypress standard 2 - Simplified
0x1C	TOC2_FIRST_CM7_0_USER_APP_ADDR	Address of CM7 core0 First User Application Object
0x20	TOC2_SECOND_CM7_0_USER_APP_ADDR	Address of CM7 core0 Second User Application Object
0x24	TOC2_FIRST_CM7_1_USER_APP_ADDR	Address of CM7 core1 First User Application Object
0x28	TOC2_SECOND_CM7_1_USER_APP_ADDR	Address of CM7 core1 Second User Application Object
0x100	TOC2_SHASH_OBJECTS	Number of additional objects (not including objects for FACTORY_HASH) starting from offset 0x104 to be verified for SECURE_HASH
0x104	TOC2_SIGNATURE_VERIF_KEY	Address of signature verification key (0 if none). The object is signature scheme specific. It is the public key in case of RSA. The default value is an address of SFlash row 50.
0x108	TOC2_APP_PROTECTION_ADDR	Address of User SWPU object stored in SFlash. The default value is an address of SFlash row 59.
0x110-0x1F4	-	... (additional objects if needed or 0's if none)
0x1F8	TOC2_FLAGS	TOC2 configuration; see Table 38-8 for more details. If TOC2 is erased, Flash boot assumes TOC2_FLAGS = 0x0000_0242.
0x1FC	Reserved	Unused for Traveo II.

Note: If additional objects need to be added to the TOC2 structure the 32-bit address (in SFlash) and object size should be provided; otherwise, the SECURE_HASH calculation may fail.

Table 38-8. SFLASH_TOC2_FLAGS Description

Bits	Name	Default	Description
1:0	CLOCK_CONFIG	2	Indicates clock frequency configuration. The clock should stay the same after Flash boot execution. 0 = 8 MHz, IMO, no FLL 1 = 25 MHz, IMO + FLL 2 = 50 MHz, IMO + FLL (default) 3 = Use ROM boot clocks configuration (100 MHz)
4:2	LISTEN_WINDOW	0	Determines the Listen window to allow sufficient time to acquire debug port. 0 = 20 ms (Default) 1 = 10 ms 2 = 1 ms 3 = 0 ms (No Listen window) 4 = 100 ms
6:5	SWJ_PINS_CTL	2	Determines if SWJ pins are configured in SWJ mode by Flash boot. Note: SWJ pins may be enabled later in the user code. 0 = Do not enable SWJ pins in Flash boot. Listen window is skipped. 1 = Do not enable SWJ pins in Flash boot. Listen window is skipped. 2 = Enable SWJ pins in Flash boot (default). 3 = Do not enable SWJ pins in Flash boot. Listen window is skipped.
8:7	APP_AUTH_CTL	0	Determines if the application image digital signature verification (authentication) is performed: 0 = Authentication is enabled (default). 1 = Authentication is disabled. 2 = Authentication is enabled (recommended). 3 = Authentication is enabled.
10:9	FB_BOOTLOADER_CTL	1	Determine if the internal bootloader in Flash boot is disabled: 0 = Internal bootloader is disabled. 1 = Internal bootloader is launched if the other bootloader conditions are met (default). 2 = Internal bootloader is disabled. 3 = Internal bootloader is disabled.

38.3.5 Internal Bootloader

Bootloaders are a common part of the MCU system design. A bootloader enables product firmware update in the field. In a typical product, firmware is stored in the MCU's flash memory.

At the factory, initial programming of firmware into a product is typically done at PCB assembly time, using the MCU's JTAG or a SWD interface. However, these interfaces are not usually available in the field, and are generally not used for firmware updates.

A better way to update firmware in the field is to use an existing connection between the product and the outside world. The connection may be a standard communication port such as I²C, USB, or UART, a wireless channel such as Bluetooth low-energy (BLE), an automotive protocol such as CAN or LIN, or a custom protocol.

The flash boot contains an internal bootloader that may be used by OEMs for initial bootloading when the code flash is empty, besides SWD or JTAG initial programming. This may allow OEMs to repurpose SWD or JTAG pins, or completely disable the debugger.


The internal bootloader supports CAN and LIN communication interfaces. A bootloader is a separate region of the flash boot that receives data from CAN or LIN communication interfaces and places it into the RAM. The intention of the bootloader is to upload and launch the user application flash loader in the RAM. The flash loader programs a user application into code flash during the OEM serial production with no SWD or JTAG connection. After data is transferred, the flash boot validates the data and executes the code to start the second stage of bootloading or run the user application.

Note: The flash loader image does not require an encryption because the flash loader is uploaded into the device by the OEM on the factory setup.

38.3.5.1 Terms and Definitions

The product's embedded firmware must be able to use the communication port for two different purposes – normal operation and to update flash. The part of embedded firmware that knows how to update flash is called a bootloader (Figure 38-15).

Figure 38-15. Bootloader System



Typically, the system that provides data to update internal flash is called the host, and the system being updated is called the target. The host can be an external computer or another MCU on the same PCB as the target.

The act of transferring data from the host to the target flash is called bootloading, or bootload operation, or just bootload. Data placed in flash is called the application or firmware image.

38.3.5.2 Using Bootloader

The bootloader and the application typically share a communication port. The first step in using a bootloader is to manipulate the product so that the bootloader, and not the application, is executing. This can be done in response to an event such as pressing a button on the product, or by sending a command to the product. The application detects such an event and responds by transferring control to the bootloader.

After the bootloader starts running, the host can send a **Start Bootload** command over the communication channel. If the bootloader sends **OK** in response, bootloading can begin.

38.3.5.3 Bootloader Activation Conditions

The internal bootloader will activate if all the following conditions are met:

- Two words at the start of flash must be 0xFFFF_FFFF
- TOC2 is valid and TOC2_FLAGS bit
FB_BOOTLOADER_DISABLE should be 2'b01
(default). Otherwise, TOC2 is erased
- Protection mode is not SECURE or SECURE_DEAD
- No debugger connection occurs during a 1-second wait window

If any of these conditions are not met, the bootloader will not start.

38.3.5.4 Basic Bootloader Function Flow


During bootloading, the host reads the file for the new application, parses it into Flash Write commands, and sends those commands to the bootloader. After the entire file is received and installed in the target flash, the bootloader can pass control to the new application.

After a device reset, a bootloader typically executes first. It can then perform the following actions:

- Check the application's validity before transferring control to that application
- Manage the timing to start host communication
- Do the bootload/flash update operation
- Pass control to the application

The flash boot is designed to update the user application in the flash with the algorithm described in Figure 38-16.

Figure 38-16. Startup and Bootloading Sequence



(A) The flash boot checks if the internal bootloader (part of the flash boot) should be run.

(D) The internal bootloader is part of the flash boot firmware that downloads the flash loader into SRAM (C) and launch it (E).

(D) The flash loader requires neither a secure signature nor an encryption because it is uploaded into the device by the OEM on the factory setup.

The flash loader application format is the basic application format with CRC-32C appended to its end, the same format is used by the bootloader SDK non-secure applications.

The CRC-32C hash is used only to check the flash loader image integrity check.

The image may be stored anywhere in the SRAM, with the only limitation that it should not overlap with the flash boot stack that starts at the end of the SRAM minus 6KB.

The flash boot bootloader receives the start address and length of the application from the data of the Set App

Metadata bootloader command, which is the second bootloader command to be sent from the bootloading host to the device.

(F) The flash loader downloads a user application through the CAN and LIN communication and stores it into the code flash or work flash.


(G) The user application is verified for integrity by the flash loader.


If the user application signature verification fails, the flash loader tries to restart bootloading and receives a new image.

(H) The user application may or may not contain a bootloader. It is up to the user.

Note that only the flash boot part of the bootloading sequence (A) to (E) is developed as the flash boot firmware; the remaining sequence is developed by the user.

Figure 38-17. Internal Bootloader Flow





38.3.5.5 End-of-Line Programming

The internal bootloader is the part of flash boot firmware that has a goal to download a flash loader into the SRAM and launch it. The flash loader downloads the user application through the CAN or LIN communication interface and stores it into the code flash or work flash. The bootloader enables the end-of-line programming using only LIN or CAN.

The CAN and LIN interfaces are combined on the same device pins to minimize the number of connections for end-of-line programming.

First, the bootloader prepares the channel configuration for CAN and waits for the preconfigured time for the frame from the host. If there is a timeout, the channel is reconfigured for LIN and it again waits for the frame. This procedure is cyclically repeated until the frame from the host is received.

The frame receipt completion ensures that the bootloader is attached to the CAN or LIN pins, and bootloading continues with the current CAN or LIN channel configuration.

Bootloader Transport Layer Implementation Details

For the bootloader to transmit commands and responses the 8-byte packet format is chosen because it fits best with the CAN and LIN protocols. This approach significantly gains in the performance of big packet transmission between a host and a device. The proof of the selected approach (pack data into 8-byte frames).

The example calculations are done for LIN for the baud rate of 115200 kbps on the longest bootloader command – Program data (the command is 32-byte long).

Note that except the CAN or LIN protocols overhead, there is the bootloader protocol overhead, which consists of

service bytes. These bytes enclose the actual data (bootloader commands).

The best case is 8 bytes in a frame, then the LIN frame includes “Break, Sync, PID, data, checksum”, so:

- the overall bit count is 124 bit
- the payload bit count is 64 bit.

If there is 1 byte in a frame, the LIN frame still includes “Break, Sync, PID, data, checksum”, so:

- the overall bit count is 68 bit
- the payload bit count is 8 bit.

The bootloader command overhead is 8 bytes per command. The bootloader command program data is 32 bytes. It also has the specific overhead of 8 bytes (address of the row for programming that comes in every program data command).

The efficiency coefficients (versus overhead) are:

- for the LIN protocol:
 - $(64/124) = 0.52$, for 8-byte frames
 - $(8/68) = 0.12$, for 1-byte frames.
- for the bootloader program data command:
 - $(32 - 8 - 8) / 32 = 0.5$, for 8-byte frames
 - $(32 - 7 - 8) / 32 = 0.53$, for 8-byte frames.

The total time to transmit the program data command:

transmit time = command size / (baud rate × efficiency coefficients)

- through 8-byte frames: $32 / [(115200 \times 103 / 8) \times (0.5 \times 0.52)] = 8.55 \mu s$
- through 1-byte frames: $32 / [(115200 \times 103 / 8) \times (0.53 \times 0.12)] = 34.94 \mu s$.

Note that this time will be bigger due to inter frame intervals (4 for 8-byte frames and 32 for 1-byte frames), and the difference will be drastic.

Transmitting 8-byte frames four times more effective than using 1-byte frames.

Sending data as a byte sequence is not supported in the current implementation; however, this can be implemented on request in future.

CAN Transport Layer Implementation Details

The Classic CAN with the 500 kbps baud rate and 8-byte data size of the RX/TX buffers is used.

The CAN RX FIFO is used to receive long messages from the host. The implementation relies on the approach described above – the host should pack data into 8-byte packets (to have a smaller number of transfers). The bootloader protocol has commands with different sizes. Long commands (more than 8 bytes) are transmitted as series of 8-byte messages. If there is a remainder (less than 8 bytes), after sending full 8-byte messages, it should be sent by the last CAN message with a smaller data field size (DLC) equal to the remainder size (in bytes).

The `CAN_Transport_Read()` function extracts received data from the CAN RX FIFO (FIFO element is 8-byte wide) and pack it into a complete command. Then the received command is passed to the bootloader's internal byte buffer for future command processing. The bootloader handles data from this buffer as one bootloader command.

The response to the host is also formed as a byte buffer. The `CAN_TransportWrite()` function sends data from the bootloader's internal byte buffer in 8-byte messages. The remainder (if any) is sent as the last CAN message with a data size less than 8 bytes, with the corresponding DLC size.

Two IDs are used to communicate with the device through CAN:

- 0x1A1 ID is used to send a bootloader command from the host to device.
- 0x1B1 ID is used to send a response from the device to host.

LIN Transport Layer Implementation Details

The LIN protocol transport layer (according to LIN specification) is not used to minimize the protocol overhead and optimize the number of frames to be sent between the master and slave. Instead, specially allocated signals are used to transmit data. Also, the transport layer protocol (with packet length, start and stop flags, and so on) is implemented at the bootloader level.

The LIN signals size is 8 bytes. This means all the frames from the LIN master should be 8-byte long. When a meaningful data message is less than 8 bytes, it should be made artificially complement to 8 bytes. The same is true for

the device responses: they are arranged in 8-byte packs, when a response (or its part) is less than 8 bytes – it is artificially made complement to 8 bytes.

Each LIN frame has the frame length field (as a first frame byte), which indicates the number of “useful” data bytes in the frame to solve the naked LIN protocol gap (this does not have a field to indicate the frame data-byte number).

Two signals are used to communicate with the LIN Slave:

- Signal with PID equal to 45 (0x2D) is used to send a bootloader command from the host (LIN master) to device (Traveo II) which is the LIN slave.
- Signal with PID equal to 46 (0x2E) is used to obtain a response from the device.

The `LIN_Transport_Read()` function receives the first LIN frame and read the bootloader packet data length. Then the function pools the LIN frames (depends on the expected command length) and extracts received data out of them. The extracted data is packed into the bootloader's internal byte buffer as a complete command. Then, the received command is passed to the bootloader's internal byte buffer. The bootloader handles data from this buffer as one bootloader command.

The response to the host is also formed as a byte buffer. The `LIN_TransportWrite()` function sends data from the bootloader's internal byte buffer in 8-byte LIN frames. Note that each bootloader command implies a response from a device. Responses have an arbitrary size that depends on the command. The host (LIN master) recognizes the command it sends and should send the required amount of the LIN headers (with 0x2E PID) to obtain a full device response to a command. For some specific commands (such as Enter Bootloader or Verify Application), when a device response to the Bootloader command can be either less than 8 bytes (for example, error response is 7-byte long) or more. In this case, the host should send a number of LIN headers expected for the longest answer. Also, the host should consider the reasonable timeout for the answer to its LIN header.

When a device sends a shorter answer to the host bootloader command (such as, error happened) and the next LIN header was not answered by the device, then the host should exit on a timeout and “assume” that the previously received LIN response was a complete device answer (to the bootloader command from the host). Then the host should process the received response and act depending on its content.

Note that the end-of-line bootloader transport layer is designed for use with peer-to-peer connection. Only one master and one slave on a bus.

A device will accept only commands with the IDs matching the device IDs (see the above sections for the selected CAN/LIN IDs). Messages or frames with any other IDs will be ignored.

Revision History



Revision	Issue Date	Description of Change
**	July 27, 2018	Initial version of TRM
*A	September 27, 2018	Rework of chapters. Addition of SMIF chapter. Updates to System Clock diagram.
*B	September 06, 2019	Re-aligned "register" usage in accordance with Registers TRM Updated register tables as per Registers TRM Added SWPU section in "Protection Unit" chapter Added ECC details for DMA Added ECC error injection handling in Code-/Work-flash and made few corrections Added ECC generation scheme for SRAM and other register descriptions related to power handling Updated Programmable PPU's/Fixed PPU's/SWPU used by Boot ROM Updated clock path and made a few corrections in the Clock Subsystem chapter Corrected typos in Watchdog Timer chapter Updated the block diagram, corrected register names, and updated SCB[0] support for SPI master in SCB chapter Edited CAN FD chapter as per review comments Added DBG freeze control in TCPWM chapter Added analog calibration details and updated register explanations in SAR ADC chapter Added ETAS CAL support feature in Program and Debug chapter Removed unnecessary descriptions, IRQ0/1 details before system calls, operation status etc. for programming Addressed review comments from design and applications team Removed device-specific details from the TRM to the datasheet
*C	November 29, 2019	Removed clock diagram in the Clock Subsystem chapter Updated Protection Units chapter with correct faults Added critical notes in the SRAM chapter (specific to rev c) Updated BootROM chapter with correct SPMU being used for SRAM protection Updated resource availability in Power modes Realigned WDT "register" usage according to the Registers TRM Corrected statements in Introduction chapter related to DeepSleep and regulators Updated Bootloader activation conditions and TOC2 in Flashboot IO alignment of some registers and removal of unneeded references Corrected TCPWM clock references Added 272-BGA CPU board reference and updated IDE revisions in the Getting started section Updated Code flash register table Added Workflash 128 KB mapping and updated register table Added CPUSS unassigned peripheral address access notes Updated Device Power Modes chapter with reference to Hibernate
*D	September 21, 2020	See the PDF file attached with this TRM for the complete revision history.