



Floating-Point IP Cores User Guide

Updated for Intel® Quartus® Prime Design Suite: **20.1**



Subscribe



Send Feedback

UG-01058 | 2021.09.13

Latest document on the web: [PDF](#) | [HTML](#)

Contents

1. About Floating-Point IP Cores.....	6
1.1. List of Floating-Point IP Cores.....	6
1.2. Installing and Licensing Intel FPGA IP Cores.....	8
1.3. Design Flow.....	9
1.3.1. IP Catalog and Parameter Editor.....	9
1.3.2. Specifying the IP Core Parameters and Options (Intel Quartus Prime Pro Edition).....	11
1.3.3. Generating IP Cores (Intel Quartus Prime Standard Edition).....	15
1.4. Upgrading IP Cores.....	15
1.4.1. Migrating IP Cores to a Different Device.....	19
1.5. Floating-Point IP Cores General Features.....	20
1.6. IEEE-754 Standard for Floating-Point Arithmetic.....	20
1.6.1. Floating-Point Formats.....	21
1.6.2. Special Case Numbers.....	22
1.6.3. Rounding.....	22
1.7. Non-IEEE-754 Standard Format.....	23
1.8. Floating-Points IP Cores Output Latency.....	23
1.9. Floating-Point IP Cores Design Example Files.....	23
1.10. VHDL Component Declaration.....	25
1.11. VHDL LIBRARY-USE Declaration.....	25
2. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Core.....	26
2.1. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Features.....	26
2.2. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Output Latency.....	26
2.3. FP_ACC_CUSTOM Intel FPGA IP Resource Utilization and Performance.....	26
2.4. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Signals.....	27
2.5. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Parameters.....	29
3. ALTFP_ADD_SUB IP Core.....	30
3.1. ALTFP_ADD_SUB Features.....	30
3.2. ALTFP_ADD_SUB Output Latency.....	30
3.3. ALTFP_ADD_SUB Truth Table.....	30
3.4. ALTFP_ADD_SUB Resource Utilization and Performance.....	31
3.5. ALTFP_ADD_SUB Design Example: Addition of Double-Precision Format Numbers.....	32
3.5.1. ALTFP_ADD_SUM Design Example: Understanding the Simulation Results.....	32
3.6. ALTFP_ADD_SUB Signals.....	33
3.7. ALTFP_ADD_SUB Parameters.....	34
4. ALTFP_DIV IP Core.....	36
4.1. ALTFP_DIV Features.....	36
4.2. ALTFP_DIV Output Latency.....	36
4.3. ALTFP_DIV Truth Table.....	37
4.4. ALTFP_DIV Resource Utilization and Performance.....	37
4.5. ALTFP_DIV Design Example: Division of Single-Precision.....	38

4.5.1. ALTFP_DIV Design Example: Understanding the Simulation Results.....	38
4.6. ALTFP_DIV Signals.....	39
4.7. ALTFP_DIV Parameters.....	41
5. ALTFP_MULT IP Core.....	42
5.1. ALTFP_MULT IP Core Features.....	42
5.2. ALTFP_MULT Output Latency.....	42
5.3. ALTFP_MULT Truth Table.....	42
5.4. ALTFP_MULT Resource Utilization and Performance.....	43
5.5. ALTFP_MULT Design Example: Multiplication of Double-Precision Format Numbers.....	44
5.5.1. ALTFP_MULT Design Example: Understanding the Simulation Waveform.....	44
5.6. Parameters.....	45
5.7. ALTFP_MULT Signals.....	45
6. ALTFP_SQRT.....	47
6.1. ALTFP_SQRT Features.....	47
6.2. Output Latency.....	47
6.3. ALTFP_SQRT Truth Table.....	48
6.4. ALTFP_SQRT Resource Utilization and Performance.....	48
6.5. ALTFP_SQRT Design Example: Square Root of Single-Precision Format Numbers.....	49
6.5.1. ALTFP_SQRT Design Example: Understanding the Simulation Results.....	49
6.6. ALTFP_SQRT Signals.....	50
6.7. ALTFP_SQRT Parameters.....	51
7. ALTFP_EXP IP Core.....	52
7.1. ALTFP_EXP Features.....	52
7.2. Output Latency.....	52
7.3. ALTFP_EXP Truth Table.....	52
7.4. ALTFP_EXP Resource Utilization and Performance.....	53
7.5. ALTFP_EXP Design Example: Exponential of Single-Precision Format Numbers.....	53
7.5.1. ALTFP_EXP Design Example: Understanding the Simulation Results.....	53
7.6. ALTFP_EXP Signals.....	55
7.7. ALTFP_EXP Parameters.....	56
8. ALTFP_INV IP Core.....	57
8.1. ALTFP_INV Features.....	57
8.2. Output Latency.....	57
8.3. ALTFP_INV Truth Table.....	57
8.4. ALTFP_INV Resource Utilization and Performance.....	58
8.5. ALTFP_INV Design Example: Inverse of Single-Precision Format Numbers	58
8.5.1. ALTFP_INV Design Example: Understanding the Simulation Results.....	58
8.6. Ports.....	60
8.7. Parameters.....	60
9. ALTFP_INV_SQRT IP Core.....	62
9.1. ALTFP_INV_SQRT Features.....	62
9.2. Output Latency.....	62
9.3. ALTFP_INV_SQRT Truth Table.....	62
9.4. ALTFP_INV_SQRT Resource Utilization and Performance.....	63
9.5. ALTFP_INV_SQRT Design Example: Inverse Square Root of Single-Precision Format Numbers	63
9.5.1. ALTFP_INV_SQRT Design Example: Understanding the Simulation Results	63

9.6. Ports.....	65
9.7. Parameters.....	65
10. ALTFP_LOG.....	67
10.1. ALTFP_LOG Features.....	67
10.2. Output Latency.....	67
10.3. ALTFP_LOG Truth Table.....	67
10.4. ALTFP_LOG Resource Utilization and Performance.....	68
10.5. ALTFP_LOG Design Example: Natural Logarithm of Single-Precision Format Numbers .	68
10.5.1. ALTFP_LOG Design Example: Understanding the Simulation Results.....	69
10.6. Signals.....	70
10.7. Parameters.....	71
11. ALTFP_ATAN IP Core.....	72
11.1. Output Latency.....	72
11.2. ALTFP_ATAN Features.....	72
11.3. ALTFP_ATAN Resource Utilization and Performance.....	72
11.4. Ports.....	72
11.5. ALTFP_ATAN Parameters.....	73
12. ALTFP_SINCOS IP Core.....	74
12.1. ALTFP_SINCOS Features.....	74
12.2. Output Latency.....	74
12.3. ALTFP_SINCOS Resource Utilization and Performance.....	74
12.4. ALTFP_SINCOS Signals.....	75
12.5. ALTFP_SINCOS Parameters.....	75
13. ALTFP_ABS IP Core.....	77
13.1. ALTFP_ABS Features.....	77
13.2. ALTFP_ABS Output Latency.....	77
13.3. ALTFP_ABS Resource Utilization and Performance.....	77
13.4. ALTFP_ABS Design Example: Absolute Value of Multiplication Results.....	78
13.4.1. ALTFP_ABS Design Example: Understanding the Simulation Results.....	78
13.5. ALTFP_ABS Signals.....	79
13.6. ALTFP_ABS Parameters.....	80
14. ALTFP_COMPARE IP Core.....	82
14.1. ALTFP_COMPARE Features.....	82
14.2. ALTFP_COMPARE Output Latency.....	82
14.3. ALTFP_COMPARE Resource Utilization and Performance.....	82
14.4. ALTFP_COMPARE Design Example: Comparison of Single-Precision Format Numbers...	83
14.4.1. ALTFP_COMPARE Design Example: Understanding the Simulation Results	83
14.5. ALTFP_COMPARE Signals.....	84
14.6. ALTFP_COMPARE Parameters.....	85
15. ALTFP_CONVERT IP Core.....	86
15.1. ALTFP_CONVERT Features.....	86
15.2. ALTFP_CONVERT Conversion Operations.....	86
15.3. ALTFP_CONVERT Output Latency.....	87
15.4. ALTFP_CONVERT Resource Utilization and Performance.....	87
15.5. ALTFP_CONVERT Design Example: Convert Double-Precision Floating-Point Format Numbers.....	89
15.5.1. ALTFP_CONVERT Design Example: Understanding the Simulation Results.....	89

15.6. ALTFP_CONVERT Signals.....	90
15.7. ALTFP_CONVERT Parameters.....	92
16. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Core.....	94
16.1. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Features.....	96
16.2. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Output Latency.....	96
16.3. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Target Frequency.....	96
16.4. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Combined Target.....	96
16.5. FP_FUNCTIONS Intel FPGA IP Resource Utilization and Performance.....	97
16.6. FP_FUNCTIONS Intel FPGA IP Signals.....	110
16.7. FP_FUNCTIONS Intel FPGA IP Parameters.....	111
17. Floating-Point IP Cores User Guide Document Archives.....	115
18. Document Revision History for the Floating-Point IP Cores User Guide.....	116

1. About Floating-Point IP Cores

Attention: Intel has discontinued the support for the following IPs:

- ALTERA_FP_MATRIX_INV IP core
- ALTERA_FP_MATRIX_MULT IP core

Therefore, Intel does not recommend use of these IPs in new designs. For more information about Intel's current IP offering, refer to Intel's Intellectual Property website.

The Intel floating-point IP cores enable you to perform floating-point arithmetic in FPGAs through optimized parameterizable functions for Intel device architectures. You can customize the IP cores by configuring various parameters to accommodate your needs.

Related Information

- [Floating-Point IP Cores User Guide Document Archives](#) on page 115
Provides a list of user guides for previous versions of the Floating-Point IP Cores.
- [Digital Signal Processing Overview](#)

1.1. List of Floating-Point IP Cores

This table lists the floating-point IP cores.

Table 1. IP Cores Available in Intel® Quartus® Prime Standard Edition Software

IP Core Name	Function Overview	Supported Device
Operator Functions		
ALTFP_ADD_SUB	Adder/Subtractor	Arria® II GZ, Arria V, Arria V GZ, Intel® Cyclone® 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix® V, and Stratix IV
ALTFP_DIV	Divider	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_MULT	Multiplier	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_SQRT	Square Root	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
continued...		

IP Core Name	Function Overview	Supported Device
Algebraic and Transcendental Functions		
ALTFP_EXP	Exponential	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_INV	Inverse	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_INV_SQRT	Inverse Square Root	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_LOG	Natural Logarithm	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
Trigonometric Functions		
ALTFP_ATAN	Arctangent	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_SINCOS	Trigonometric Sine/Cosine	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
Other Functions		
ALTFP_ABS	Absolute value	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_COMPARE	Comparator	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
ALTFP_CONVERT	Converter	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Arria II GX, Cyclone IV GX, Stratix V, and Stratix IV
FP_ACC_CUSTOM Intel FPGA IP	An application specific accumulator	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Intel Arria 10, Arria II GX, Cyclone IV GX, Stratix V, Stratix IV, and Intel MAX [®] 10
FP_FUNCTIONS Intel FPGA IP	A collection of floating-point functions.	Arria II GZ, Arria V, Arria V GZ, Intel Cyclone 10 LP, Cyclone IV E, Cyclone V, Intel Arria 10, Arria II GX, Cyclone IV GX, Stratix V, Stratix IV, and Intel MAX 10

Table 2. IP Cores Available in Intel Quartus® Prime Pro Edition Software

IP Core Name	Function Overview	Supported Device
Floating Point Functions Intel FPGA IP	A collection of floating-point functions. This IP core replaces all other floating-point IP cores listed in the Quartus Prime Standard Edition table for devices available in Intel Quartus® Prime Pro Edition software.	Intel Stratix 10, Intel Arria 10, Intel Cyclone 10 GX, and Intel Agilex™
Floating Point Custom Accumulator Intel FPGA IP	An application specific accumulator	Intel Cyclone 10 GX and Intel Arria 10


Related Information

- [Introduction to Intel FPGA IP Cores](#)
Provides more information about Intel FPGA IP Cores
- [Generating a Combined Simulator Setup Script \(Intel Quartus Prime Pro Edition\)](#)
Create simulation scripts that do not require manual updates for software or IP version upgrades.

1.2. Installing and Licensing Intel FPGA IP Cores

The Intel Quartus Prime software installation includes the Intel FPGA IP library. This library provides many useful IP cores for your production use without the need for an additional license. Some Intel FPGA IP cores require purchase of a separate license for production use. The Intel FPGA IP Evaluation Mode allows you to evaluate these licensed Intel FPGA IP cores in simulation and hardware, before deciding to purchase a full production IP core license. You only need to purchase a full production license for licensed Intel IP cores after you complete hardware testing and are ready to use the IP in production.

The Intel Quartus Prime software installs IP cores in the following locations by default:

Figure 1. IP Core Installation Path

Table 3. IP Core Installation Locations

Location	Software	Platform
<drive>:\intelFPGA_pro\quartus\ip\altera	Intel Quartus Prime Pro Edition	Windows*
<drive>:\intelFPGA\quartus\ip\altera	Intel Quartus Prime Standard Edition	Windows
<home directory>:/intelFPGA_pro/quartus/ip/altera	Intel Quartus Prime Pro Edition	Linux*
<home directory>:/intelFPGA/quartus/ip/altera	Intel Quartus Prime Standard Edition	Linux

Note: The Intel Quartus Prime software does not support spaces in the installation path.

1.3. Design Flow

Use the IP Catalog and parameter editor to define and instantiate complex IP cores. Using the GUI ensures that you set all IP core ports and parameters properly.

If you are an expert user, and choose to configure the IP core directly through parameterized instantiation in your design, refer to the port and parameter details. The details of these ports and parameters are hidden in the parameter editor.

1.3.1. IP Catalog and Parameter Editor


The IP Catalog displays the IP cores available for your project, including Intel FPGA IP and other IP that you add to the IP Catalog search path. Use the following features of the IP Catalog to locate and customize an IP core:

- Filter IP Catalog to **Show IP for active device family** or **Show IP for all device families**. If you have no project open, select the **Device Family** in IP Catalog.
- Type in the Search field to locate any full or partial IP core name in IP Catalog.
- Right-click an IP core name in IP Catalog to display details about supported devices, to open the IP core's installation folder, and for links to IP documentation.
- Click **Search for Partner IP** to access partner IP information on the web.

The parameter editor prompts you to specify an IP variation name, optional ports, and output file generation options. The parameter editor generates a top-level Intel Quartus Prime IP file (.qip) for an IP variation in Intel Quartus Prime Pro Edition projects. This file represents the IP variation in the project, and stores parameterization information.⁽¹⁾

⁽¹⁾ The parameter editor generates a top-level Quartus IP file (.qip) for an IP variation in Intel Quartus Prime Standard Edition projects.

Figure 2. Example IP Parameter Editor



1.3.1.1. The Parameter Editor

The parameter editor helps you to configure IP core ports, parameters, and output file generation options. The basic parameter editor controls include the following:

- Use the **Presets** window to apply preset parameter values for specific applications (for select cores).
- Use the **Details** window to view port and parameter descriptions, and click links to documentation.
- Click **Generate** ► **Generate Testbench System** to generate a testbench system (for select cores).
- Click **Generate** ► **Generate Example Design** to generate an example design (for select cores).
- Click **Validate System Integrity** to validate a system's generic components against companion files. (Platform Designer systems only)
- Click **Sync All System Info** to validate a system's generic components against companion files. (Platform Designer systems only)

The IP Catalog is also available in Platform Designer (**View** ► **IP Catalog**). The Platform Designer IP Catalog includes exclusive system interconnect, video and image processing, and other system-level IP that are not available in the Intel Quartus Prime IP Catalog. Refer to *Creating a System with Platform Designer* or *Creating a System with Platform Designer (Standard)* for information on use of IP in Platform Designer (Standard) and Platform Designer, respectively.

Related Information

- [Creating a System with Platform Designer](#)
- [Creating a System with Platform Designer \(Standard\)](#)


1.3.2. Specifying the IP Core Parameters and Options (Intel Quartus Prime Pro Edition)

Quickly configure Intel FPGA IP cores in the Intel Quartus Prime parameter editor. Double-click any component in the IP Catalog to launch the parameter editor. The parameter editor allows you to define a custom variation of the IP core. The parameter editor generates the IP variation synthesis and optional simulation files, and adds the .ip file representing the variation to your project automatically.

Follow these steps to locate, instantiate, and customize an IP core in the parameter editor:

1. Create or open an Intel Quartus Prime project (.qpf) to contain the instantiated IP variation.
2. In the IP Catalog (**Tools > IP Catalog**), locate and double-click the name of the IP core to customize. To locate a specific component, type some or all of the component's name in the IP Catalog search box. The New IP Variation window appears.
3. Specify a top-level name for your custom IP variation. Do not include spaces in IP variation names or paths. The parameter editor saves the IP variation settings in a file named <your_ip>.ip. Click **OK**. The parameter editor appears.

Figure 3. IP Parameter Editor (Intel Quartus Prime Pro Edition)



4. Set the parameter values in the parameter editor and view the block diagram for the component. The **Parameterization Messages** tab at the bottom displays any errors in IP parameters:
 - Optionally, select preset parameter values if provided for your IP core. Presets specify initial parameter values for specific applications.
 - Specify parameters defining the IP core functionality, port configurations, and device-specific features.
 - Specify options for processing the IP core files in other EDA tools.


Note: Refer to your IP core user guide for information about specific IP core parameters.

5. Click **Generate HDL**. The **Generation** dialog box appears.
6. Specify output file generation options, and then click **Generate**. The synthesis and simulation files generate according to your specifications.
7. To generate a simulation testbench, click **Generate ► Generate Testbench System**. Specify testbench generation options, and then click **Generate**.
8. To generate an HDL instantiation template that you can copy and paste into your text editor, click **Generate ► Show Instantiation Template**.
9. Click **Finish**. Click **Yes** if prompted to add files representing the IP variation to your project.
10. After generating and instantiating your IP variation, make appropriate pin assignments to connect ports.

Note: Some IP cores generate different HDL implementations according to the IP core parameters. The underlying RTL of these IP cores contains a unique hash code that prevents module name collisions between different variations of the IP core. This unique code remains consistent, given the same IP settings and software version during IP generation. This unique code can change if you edit the IP core's parameters or upgrade the IP core version. To avoid dependency on these unique codes in your simulation environment, refer to *Generating a Combined Simulator Setup Script*.

1.3.2.1. IP Core Generation Output (Intel Quartus Prime Pro Edition)

The Intel Quartus Prime software generates the following output file structure for individual IP cores that are not part of a Platform Designer system.

Figure 4. Individual IP Core Generation Output (Intel Quartus Prime Pro Edition)**Table 4. Output Files of Intel FPGA IP Generation**


File Name	Description
<your_ip>.ip	Top-level IP variation file that contains the parameterization of an IP core in your project. If the IP variation is part of a Platform Designer system, the parameter editor also generates a .qsys file.
<your_ip>.cmp	The VHDL Component Declaration (.cmp) file is a text file that contains local generic and port definitions that you use in VHDL design files.
<your_ip>_generation.rpt	IP or Platform Designer generation log file. Displays a summary of the messages during IP generation.
continued...	

File Name	Description
<your_ip>.qgsimc (Platform Designer systems only)	Simulation caching file that compares the .qsys and .ip files with the current parameterization of the Platform Designer system and IP core. This comparison determines if Platform Designer can skip regeneration of the HDL.
<your_ip>.qgsynth (Platform Designer systems only)	Synthesis caching file that compares the .qsys and .ip files with the current parameterization of the Platform Designer system and IP core. This comparison determines if Platform Designer can skip regeneration of the HDL.
<your_ip>.qip	Contains all information to integrate and compile the IP component.
<your_ip>.csv	Contains information about the upgrade status of the IP component.
<your_ip>.bsf	A symbol representation of the IP variation for use in Block Diagram Files (.bdf).
<your_ip>.spd	Input file that ip-make-simscript requires to generate simulation scripts. The .spd file contains a list of files you generate for simulation, along with information about memories that you initialize.
<your_ip>.ppf	The Pin Planner File (.ppf) stores the port and node assignments for IP components you create for use with the Pin Planner.
<your_ip>_bb.v	Use the Verilog blackbox (_bb.v) file as an empty module declaration for use as a blackbox.
<your_ip>_inst.v or _inst.vhd	HDL example instantiation template. Copy and paste the contents of this file into your HDL file to instantiate the IP variation.
<your_ip>.regmap	If the IP contains register information, the Intel Quartus Prime software generates the .regmap file. The .regmap file describes the register map information of master and slave interfaces. This file complements the .sopcinfo file by providing more detailed register information about the system. This file enables register display views and user customizable statistics in System Console.
<your_ip>.svd	Allows HPS System Debug tools to view the register maps of peripherals that connect to HPS within a Platform Designer system. During synthesis, the Intel Quartus Prime software stores the .svd files for slave interface visible to the System Console masters in the .sof file in the debug session. System Console reads this section, which Platform Designer queries for register map information. For system slaves, Platform Designer accesses the registers by name.
<your_ip>.v <your_ip>.vhd	HDL files that instantiate each submodule or child IP core for synthesis or simulation.
mentor/	Contains a msim_setup.tcl script to set up and run a ModelSim* simulation.
aldec/	Contains a Riviera-PRO* script rivierapro_setup.tcl to setup and run a simulation.
/synopsys/vcs /synopsys/vcsmx	Contains a shell script vcs_setup.sh to set up and run a VCS* simulation. Contains a shell script vcsmx_setup.sh and synopsys_sim.setup file to set up and run a VCS MX simulation.
/cadence	Contains a shell script ncsim_setup.sh and other setup files to set up and run an NCSim simulation.
/xcelium	Contains an Xcelium* Parallel simulator shell script xcelium_setup.sh and other setup files to set up and run a simulation.
/submodules	Contains HDL files for the IP core submodule.
<IP submodule>/	Platform Designer generates /synth and /sim sub-directories for each IP submodule directory that Platform Designer generates.

1.3.3. Generating IP Cores (Intel Quartus Prime Standard Edition)

This topic describes parameterizing and generating an IP variation using a legacy parameter editor in the Intel Quartus Prime Standard Edition software.

Figure 5. Legacy Parameter Editors




1. In the IP Catalog (**Tools > IP Catalog**), locate and double-click the name of the IP core to customize. The parameter editor appears.
2. Specify a top-level name and output HDL file type for your IP variation. This name identifies the IP core variation files in your project. Click **OK**. Do not include spaces in IP variation names or paths.
3. Specify the parameters and options for your IP variation in the parameter editor. Refer to your IP core user guide for information about specific IP core parameters.
4. Click **Finish** or **Generate** (depending on the parameter editor version). The parameter editor generates the files for your IP variation according to your specifications. Click **Exit** if prompted when generation is complete. The parameter editor adds the top-level .qip file to the current project automatically.

Note: For devices released prior to Intel Arria 10 devices, the generated .qip and .sip files must be added to your project to represent IP and Platform Designer systems. To manually add an IP variation generated with legacy parameter editor to a project, click **Project > Add/Remove Files in Project** and add the IP variation .qip file.

1.4. Upgrading IP Cores

Any Intel FPGA IP variations that you generate from a previous version or different edition of the Intel Quartus Prime software, may require upgrade before compilation in the current software edition or version. The Project Navigator displays a banner indicating the IP upgrade status. Click **Launch IP Upgrade Tool** or **Project > Upgrade IP Components** to upgrade outdated IP cores.

Figure 6. IP Upgrade Alert in Project Navigator



Icons in the **Upgrade IP Components** dialog box indicate when IP upgrade is required, optional, or unsupported for an IP variation in the project. Upgrade IP variations that require upgrade before compilation in the current version of the Intel Quartus Prime software.

Note: Upgrading IP cores may append a unique identifier to the original IP core entity names, without similarly modifying the IP instance name. There is no requirement to update these entity references in any supporting Intel Quartus Prime file, such as the Intel Quartus Prime Settings File (.qsf), Synopsys* Design Constraints File (.sdc), or Signal Tap File (.stp), if these files contain instance names. The Intel Quartus Prime software reads only the instance name and ignores the entity name in paths that specify both names. Use only instance names in assignments.


Table 5. IP Core Upgrade Status

IP Core Status	Description
IP Upgraded 	Indicates that your IP variation uses the latest version of the Intel FPGA IP core.
IP Component Outdated 	Indicates that your IP variation uses an outdated version of the IP core.
IP Upgrade Optional 	Indicates that upgrade is optional for this IP variation in the current version of the Intel Quartus Prime software. You can upgrade this IP variation to take advantage of the latest development of this IP core. Alternatively, you can retain previous IP core characteristics by declining to upgrade. Refer to the Description for details about IP core version differences. If you do not upgrade the IP, the IP variation synthesis and simulation files are unchanged and you cannot modify parameters until upgrading.
IP Upgrade Required 	Indicates that you must upgrade the IP variation before compiling in the current version of the Intel Quartus Prime software. Refer to the Description for details about IP core version differences.
continued...	

IP Core Status	Description
	
IP Upgrade Unsupported 	Indicates that upgrade of the IP variation is not supported in the current version of the Intel Quartus Prime software due to incompatibility with the current version of the Intel Quartus Prime software. The Intel Quartus Prime software prompts you to replace the unsupported IP core with a supported equivalent IP core from the IP Catalog. Refer to the Description for details about IP core version differences and links to Release Notes.
IP End of Life 	Indicates that Intel designates the IP core as end-of-life status. You may or may not be able to edit the IP core in the parameter editor. Support for this IP core discontinues in future releases of the Intel Quartus Prime software.
IP Upgrade Mismatch Warning 	Provides warning of non-critical IP core differences in migrating IP to another device family.
IP has incompatible subcores 	Indicates that the current version of the Intel Quartus Prime software does not support compilation of your IP variation, because the IP has incompatible subcores
Compilation of IP Not Supported 	Indicates that the current version of the Intel Quartus Prime software does not support compilation of your IP variation. This can occur if another edition of the Intel Quartus Prime software, such as the Intel Quartus Prime Standard Edition, generated this IP. Replace this IP component with a compatible component in the current edition.

Note: Beginning with the Intel Quartus Prime Pro Edition software version 19.1, IP upgrade supports migration of IP released within one year of the Intel Quartus Prime Pro Edition software version, as the following chart defines:


Figure 7. Intel Quartus Prime Pro Edition IP Version Upgrade Paths



Follow these steps to upgrade IP cores:

1. In the latest version of the Intel Quartus Prime software, open the Intel Quartus Prime project containing an outdated IP core variation. The **Upgrade IP Components** dialog box automatically displays the status of IP cores in your project, along with instructions for upgrading each core. To access this dialog box manually, click **Project > Upgrade IP Components**.
2. To upgrade one or more IP cores that support automatic upgrade, ensure that you turn on the **Auto Upgrade** option for the IP cores, and click **Auto Upgrade**. The **Status** and **Version** columns update when upgrade is complete. Example designs that any Intel FPGA IP core provides regenerate automatically whenever you upgrade an IP core.
3. To manually upgrade an individual IP core, select the IP core and click **Upgrade in Editor** (or simply double-click the IP core name). The parameter editor opens, allowing you to adjust parameters and regenerate the latest version of the IP core.

Figure 8. Upgrading IP Cores (Intel Quartus Prime Pro Edition Example)



Note: Intel FPGA IP cores older than Intel Quartus Prime software version 12.0 do not support upgrade. Intel verifies that the current version of the Intel Quartus Prime software compiles the previous two versions of each IP core. The *Intel FPGA IP Core Release Notes* reports any verification exceptions for Intel FPGA IP cores. Intel does not verify compilation for IP cores older than the previous two releases.

Related Information

[Intel FPGA IP Release Notes](#)

1.4.1. Migrating IP Cores to a Different Device

Migrate an Intel FPGA IP variation when you want to target a different (often newer) device. Most Intel FPGA IP cores support automatic migration. Some IP cores require manual IP regeneration for migration. A few IP cores do not support device migration, requiring you to replace them in the project. The **Upgrade IP Components** dialog box identifies the migration support level for each IP core in the design.

1. To display the IP cores that require migration, click **Project > Upgrade IP Components**. The **Description** field provides migration instructions and version differences.
2. To migrate one or more IP cores that support automatic upgrade, ensure that the **Auto Upgrade** option is turned on for the IP cores, and click **Perform Automatic Upgrade**. The **Status** and **Version** columns update when upgrade is complete.
3. To migrate an IP core that does not support automatic upgrade, double-click the IP core name, and click **OK**. The parameter editor appears. If the parameter editor specifies a **Currently selected device family**, turn off **Match project/default**, and then select the new target device family.

4. Click **Generate HDL**, and confirm the **Synthesis** and **Simulation** file options. Verilog HDL is the default output file format. If you specify VHDL as the output format, select **VHDL** to retain the original output format.
5. Click **Finish** to complete migration of the IP core. Click **OK** if the software prompts you to overwrite IP core files. The **Device Family** column displays the new target device name when migration is complete.
6. To ensure correctness, review the latest parameters in the parameter editor or generated HDL.

Note: IP migration may change ports, parameters, or functionality of the IP variation. These changes may require you to modify your design or to re-parameterize your IP variant. During migration, the IP variation's HDL generates into a library that is different from the original output location of the IP core. Update any assignments that reference outdated locations. If a symbol in a supporting Block Design File schematic represents your upgraded IP core, replace the symbol with the newly generated `<my_ip>.bsf`. Migration of some IP cores requires installed support for the original and migration device families.

Related Information

[Intel FPGA IP Release Notes](#)

1.5. Floating-Point IP Cores General Features

All Intel FPGA floating-point IP cores offer the following features:

- Support for floating-point formats.
- Input support for not-a-number (NaN), infinity, zero, and normal numbers.
- Optional asynchronous input ports including asynchronous clear (`ac_l_r`) and clock enable (`clk_en`).
- Support for round-to-nearest-even rounding mode.
- Compute results of any mathematical operations according to the IEEE-754 standard compliance with a maximum of 1 unit in the last place (u.l.p.) error. This assumption is applied to all floating-point IP cores.

Intel FPGA floating-point IP cores do not support denormal number inputs. If the input is a denormal value, the IP core forces the value to zero and treats the value as a zero before going through any operation.

Related Information

[FFT IP Core: User Guide](#)

Intel also offers the single-precision floating-point option in the FFT IP core.

1.6. IEEE-754 Standard for Floating-Point Arithmetic

The floating-point IP cores implement the following representations in the IEEE-754 standard:

- Floating-point numbers
- Special values (zero, infinity, denormal numbers, and NaN bit combinations)
- Single-precision, double-precision, and single-extended precision formats for floating-point numbers


1.6.1. Floating-Point Formats

All floating-point formats have binary patterns. In the following figure, S represents a sign bit, E represents an exponent field, and M is the mantissa (part of a logarithm, or fraction) field.

For a normal floating-point number, a leading 1 is always implied, for example, binary 1.0011 or decimal 1.1875 is stored as 0011 in the mantissa field. This format saves the mantissa field from using an extra bit to represent the leading 1. However, the leading bit for a denormal number can be either 0 or 1. For zero, infinity, and NaN, the mantissa field does not have an implied leading 1 nor any explicit leading bit.

Figure 9. IEEE-754 Floating-Point Format

This figure shows a floating-point format.



1.6.1.1. Single-Precision Format


The single-precision format contains the following binary patterns:

- The MSB holds the sign bit.
- The next 8 bits hold the exponent bits.
- 23 LSBs hold the mantissa.

The total width of a floating-point number in the single-precision format is 32 bits. The bias for the single-precision format is 127.

Figure 10. Single-Precision Representation

This figure shows a single-precision representation.



1.6.1.2. Double-Precision Format


The double-precision format contains the following binary patterns:

- The MSB holds the sign bit.
- The next 11 bits hold the exponent bits.
- 52 LSBs hold the mantissa.

The total width of a floating-point number in the double-precision format is 64 bits. The bias for the double-precision format is 1023.

Figure 11. Double-Precision Representation

This figure shows a double-precision representation.



1.6.1.3. Single-Extended Precision Format

The single-extended precision format contains the following binary patterns:

- The MSB holds the sign bit.
- The exponent and mantissa fields do not have fixed widths.
- The minimum exponent field width is 11 bits and must be less than the width of the mantissa field.
- The width of the mantissa field must be a minimum of 31 bits.

The sum of the widths of the sign bit, exponent field, and mantissa field must be a minimum of 43 bits and a maximum of 64 bits. The bias for the single-extended precision format is unspecified in the IEEE-754 standard. In these IP cores, a bias of $2^{(\text{WIDTH_EXP}-1)}-1$ is assumed for the single-extended precision format.

1.6.2. Special Case Numbers

The following table lists the special case numbers defined by the IEEE-754 standard and the data bit representations.

Table 6. Special Case Numbers in IEEE-754 Representation

Meaning	Sign Field	Exponent Field	Mantissa Field
Zero	Don't care	All 0's	All 0's
Positive Denormalized	0	All 0's	Non-zero
Negative Denormalized	1	All 0's	Non-zero
Positive Infinity	0	All 1's	All 0's
Negative Infinity	1	All 1's	All 0's
Not-a-Number (NaN)	Don't care	All 1's	Non-zero

1.6.3. Rounding

The IEEE-754 standard defines four types of rounding modes, which are:

- round-to-nearest-even
- round-toward-zero
- round-toward-positive-infinity
- round-toward-negative-infinity

Intel floating-point IP cores support only the most commonly used rounding mode, which is the round-to-nearest-even mode (TO_NEAREST). With round-to-nearest-even, the IP core rounds the result to the nearest floating-point number. If the result is exactly halfway between two floating-point numbers, the IP core rounds the result so that the LSB becomes a zero, which is even.

1.7. Non-IEEE-754 Standard Format

Only the ALTFP_CONVERT and FP_FUNCTIONS Intel FPGA IP (when the convert function is selected) support the fixed point format.

The fixed-point data type is similar to the conventional integer data type, except that the fixed-point data carries a predetermined number of fractional bits. If the width of the fraction is 0, the data becomes a normal signed integer.

The notation for fixed-point format numbers in this user guide is $Q_m.f$, where Q designates that the number is in Q format notation, m is the number of bits used to indicate the integer portion of the number, and f is the number of bits used to indicate the fractional portion of the number.


For example, $Q_{4.12}$ describes a number with 4 integer bits and 12 fractional bits in a 16-bit word.

The following figures show the difference between the signed-integer format and the fixed-point format for a 32-bit number.

Figure 12. Signed-Integer Format



Figure 13. Fixed-Point Format



1.8. Floating-Points IP Cores Output Latency

The IP cores measure the output latency in clock cycles and is different for each IP core. In some IP cores, the precision modes determine the number of clock cycles between the input and output result. When you select a mode, the options for latency are fixed for that mode.

For specific details about latency options, refer to the Output *Latency* section of your selected IP core in this user guide.

1.9. Floating-Point IP Cores Design Example Files

The design examples for each IP core in this user guide use the IP Catalog and parameter editor to define custom IP variations.

Simulate the designs in the ModelSim - Intel FPGA Edition software to generate a waveform display of the device behavior. You must be familiar with the ModelSim - Intel FPGA Edition software before trying out the design examples.

Table 7. Design Files for Floating-Point IP Cores

Floating-Point IP Cores	Design Files
ALTFP_ADD_SUB	- altfp_add_sub_DesignExample.zip (Intel Quartus Prime design files) - altfp_add_sub_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_DIV	- altfp_div_DesignExample.zip (Intel Quartus Prime design files) - altfp_div_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_MULT	- altfp_mult_DesignExample.zip (Intel Quartus Prime design files) - altfp_mult_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_SQRT	- altfp_sqrt_DesignExample.zip (Intel Quartus Prime design files) - altfp_sqrt_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_EXP	- altfp_exp_DesignExample.zip (Intel Quartus Prime design files) - altfp_exp_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_INV	- altfp_inv_DesignExample.zip (Intel Quartus Prime design files) - altfp_inv_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_INV_SQRT	- altfp_inv_sqrt_DesignExample.zip (Intel Quartus Prime design files) - altfp_inv_sqrt_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_LOG	- altfp_log_DesignExample.zip (Intel Quartus Prime design files) - altfp_log_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_ATAN	Not Available
ALTFP_SINCOS	Not Available
ALTFP_ABS	- altfp_mult_abs_DesignExample.zip (Intel Quartus Prime design files) - altfp_mult_abs_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_COMPARE	- altfp_compare_DesignExample.zip (Intel Quartus Prime design files) - altfp_compare_ex_msim.zip (ModelSim - Intel FPGA Edition files)
ALTFP_CONVERT	- altfp_convert_DesignExample.zip (Intel Quartus Prime design files) - altfp_convert_float2int_msim.zip (ModelSim - Intel FPGA Edition files)
FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP	Not Available
FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP	Not Available

Related Information

- [ALTFP_ADD_SUB Design Example: Addition of Double-Precision Format Numbers on page 32](#)
- [ALTFP_DIV Design Example: Division of Single-Precision on page 38](#)
- [ALTFP_MULT Design Example: Multiplication of Double-Precision Format Numbers on page 44](#)
- [ALTFP_SQRT Design Example: Square Root of Single-Precision Format Numbers on page 49](#)

- [ALTFP_EXP Design Example: Exponential of Single-Precision Format Numbers](#) on page 53
- [ALTFP_INV Design Example: Inverse of Single-Precision Format Numbers](#) on page 58
- [ALTFP_INV_SQRT Design Example: Inverse Square Root of Single-Precision Format Numbers](#) on page 63
- [ALTFP_LOG Design Example: Natural Logarithm of Single-Precision Format Numbers](#) on page 68
- [ALTFP_ABS Design Example: Absolute Value of Multiplication Results](#) on page 78
- [ALTFP_COMPARE Design Example: Comparison of Single-Precision Format Numbers](#) on page 83
- [ALTFP_CONVERT Design Example: Convert Double-Precision Floating-Point Format Numbers](#) on page 89
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

1.10. VHDL Component Declaration

The VHDL component declaration is located in the <Intel Quartus Prime installation directory>\libraries\vhdl\altera_mf\altera_mf_components.vhd

1.11. VHDL LIBRARY-USE Declaration

The VHDL LIBRARY-USE declaration is not required if you use the VHDL Component Declaration.

```
LIBRARY altera_mf;  
  
USE altera_mf.altera_mf_components.all;
```

2. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Core

This IP core performs floating-point accumulation and allows you to restrict the range of inputs and maximum accumulated value to save resources. The core uses device latency models to generate RTL to meet a target FMax at the cost of latency.

Table 8. Floating Point Custom Accumulator Intel FPGA IP Release Information

Item	Description
Version	19.1
Intel Quartus Prime Version	20.1
Release Date	2020.04.13

2.1. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Features

The FPACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP core offers the following features:

- Supports frequency driven cores.
- Supports VHDL RTL generation.
- Supports customization of the required range of the input and output values.

2.2. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Output Latency

The amount of latency is driven by the target frequency and the selected device family. You must set the desired frequency and the target device before generating the IP core. The IP core reports the latency when you set the parameters and when you generate the IP core. Then, use the reported latency to incorporate the IP core into your design.

2.3. FP_ACC_CUSTOM Intel FPGA IP Resource Utilization and Performance

Table 9. FP_ACC_CUSTOM Intel FPGA IP Resource Utilization and Performance

This table lists the resource utilization and performance information for the FP_ACC_CUSTOM Intel FPGA IP core. The information was derived using the Intel Quartus Prime software version 13.1.

Device Family	Input Data		Accumulator Size		Target Frequency (MHz)	Latency	ALMs	DSP Blocks	Logic Registers		M10K	M20K	f _{MAX}
	Floating Point Format	MaxM SBX	MSBA	LSBA					Primary	Secondary			
Arria V (5AGX FB3H4 F40C5)	Double	24	40	-52	270	15	866	0	1,166	106	0	--	265
Cyclone V (5CGX FC7D6 F31C7)	Double	24	40	-52	230	15	830	0	1,102	32	0	--	198
Stratix V (5SGX EA7K2 F40C2)	Double	24	40	-52	400	15	968	0	1,655	27	--	0	426
Arria V (5AGX FB3H4 F40C5)	Single	12	20	-26	270	12	337	0	588	52	0	--	309
Cyclone V (5CGX FC7D6 F31C7)	Single	12	20	-26	230	12	383	0	494	28	0	--	225
Stratix V (5SGX EA7K2 F40C2)	Single	12	20	-26	400	13	475	0	903	20	--	0	450

Related Information

Fitter Resources Reports

Provides information about Intel Quartus Prime resource utilization

2.4. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Signals

Figure 14. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP




Table 10. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Input Ports

Port Name	Required	Description
clk	Yes	All input signals, otherwise explicitly stated, must be synchronous to this clock
areset	Yes	Asynchronous active-high reset. Deassert this signal synchronously to the input clock to avoid metastability issues.
en	No	Global enable signal. This port is optional.
x	Yes	Data input port.
n	Yes	Boolean port which signals the beginning of a new data set to be accumulated. This should go high together with the first element in the new data set and should go low the next cycle. The data sets may be of variable length and a new data set may be started at any time. The accumulation result for an input is available after the reported latency.

Table 11. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Output Ports

Port Name	Required	Description
r	Yes	The running value of the accumulation.
x0	Yes	The overflow flag for port x. The signal goes high when the exponent of the input x is larger than maxMSBX. The signal remains high for the entire data set. This flag invalidates port r. You should consider increasing maxMSBX. This flag also indicate infinity and NaN.
xu	Yes	The underflow flag for port x. The signal goes high when the exponent of the input x is smaller than LSBA. The signal remains high for the entire data set. This flag does not invalidate port r. You should consider lowering LSBA.
ao	Yes	The overflow flag for Accumulator. The signal goes high when the exponent of the accumulated value is larger than MSBA. The signal remains high for the entire data set. This flag invalidates port r. You should consider increasing MSBA.

2.5. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Parameters

Table 12. FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP Parameters

Category	Parameter	Values	Description
Input Data	Floating point format	single, double	Choose the floating point format of the input data values. The output data values of the accumulator is in the same format. The default is single .
	maxMSBX	—	The maximum weight of the MSB of an input. For example, when adding probabilities in the 0 to 1 range set this weight to $\text{ceil}(\log_2(1))=0$. The xo output signal goes high when the MSB of an input value has a weight larger than maxMSBX. The result of the accumulation is then invalid. If you are unsure about the range of the inputs, then set the maxMSBX parameter to MSBA, at the possible expense of increased resource usage. The default value is 12 .
Accumulator Size	MSBA	—	The weight of the MSB of the accumulator. For example, in a financial simulation, if the value of a stock cannot exceed 100,000 dollars, use a value of $\text{ceil}(\log_2(100000))=17$. In a circuit simulation where the circuit adds numbers in the 0 to 1 range, for one year, at 400 MHz, use a value of $\text{ceil}(\log_2(365 \times 60 \times 60 \times 24 \times 400 \times 10^6))=54$. The ao output signal goes high when the MSB of the accumulated value has a weight larger than MSBA. The result of the accumulation is then invalid. Intel recommends adding a few guard bits to avoid possible accumulator overflow. A few guard bits have little impact on the accumulator size. The default value is 20 .
	LSBA	—	The weight of the LSB of the accumulator and the accuracy of the accumulator. Because an N term accumulation can invalidate the $\log_2(N)$ LSBs of the accumulator, you must consider the length of the accumulation and the range of the inputs when setting this parameter. For example, if a 2^{-30} accuracy is required over an accumulation of 1024 numbers, then set the LSBA to: $(-30 - \log_2(1024)) = -40$. Any input $2^e \times 1.F$, where F is the mantissa and e is less than the LSBA will be shifted out of the accumulator. The au output signal goes high to indicate this situation. The default value is -26 .
Required Performance	Target frequency	Any positive integer value.	Choose the frequency in MHz at which this core is expected to run. This together with the target device family determines the amount of pipelining in the core. The default value is 200 MHz.
Optional	Generate an enable port	—	Choose if the accumulator should have an enable signal. This parameter is disabled by default.
Report	—	—	Reports the latency of the device, which is the number of cycles it takes for an accumulation to propagate through the block from input to output.

3. ALTFP_ADD_SUB IP Core

This IP core allows you to perform floating-point addition or subtraction between two inputs dynamically.

3.1. ALTFP_ADD_SUB Features

The ALTFP_ADD_SUB IP core offers the following features:

- Dynamically configurable adder and subtracter functions.
- Optional exception handling output ports such as zero, overflow, underflow, and nan.
- Optimization of speed and area.

3.2. ALTFP_ADD_SUB Output Latency

The output latency options for the ALTFP_ADD_SUB IP core are the same for all three precision formats—single, double, and single-extended. The options available are 7, 8, 9, 10, 11, 12, 13, and 14 clock cycles.

3.3. ALTFP_ADD_SUB Truth Table

Table 13. Truth Table for Addition/Subtraction Operations

DATAA[]	DATAB[]	SIGN BIT	RESULT[]	Overflow	Underflow	Zero	NaN
Normal	Normal	0	Zero	0	0	1	0
Normal	Normal	0/1	Normal	0	0	0	0
Normal	Normal	0/1	Denormal	0	1	1	0
Normal	Normal	0/1	Infinity	1	0	0	0
Normal	Denormal	0/1	Normal	0	0	0	0
Normal	Zero	0/1	Normal	0	0	0	0
Normal	Infinity	0/1	Infinity	1	0	0	0
Normal	NaN	X	NaN	0	0	0	1
Denormal	Normal	0/1	Normal	0	0	0	0
Denormal	Denormal	0/1	Normal	0	0	0	0
Denormal	Zero	0/1	Zero	0	0	1	0
Denormal	Infinity	0/1	Infinity	1	0	0	0

continued...

Intel Corporation. All rights reserved. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

**ISO
9001:2015
Registered**

DATAA[]	DATAB[]	SIGN BIT	RESULT[]	Overflow	Underflow	Zero	NaN
Denormal	NaN	X	NaN	0	0	0	1
Zero	Normal	0/1	Normal	0	0	0	0
Zero	Denormal	0/1	Zero	0	0	1	0
Zero	Zero	0/1	Zero	0	0	1	0
Zero	Infinity	0/1	Infinity	1	0	0	0
Zero	NaN	X	NaN	0	0	0	1
Infinity	Normal	0/1	Infinity	1	0	0	0
Infinity	Denormal	0/1	Infinity	1	0	0	0
Infinity	Zero	0/1	Infinity	1	0	0	0
Infinity	Infinity	0/1	Infinity	1	0	0	0
Infinity	NaN	X	NaN	0	0	0	1
NaN	Normal	X	NaN	0	0	0	1
NaN	Denormal	X	NaN	0	0	0	1
NaN	Zero	X	NaN	0	0	0	1
NaN	Infinity	X	NaN	0	0	0	1
NaN	NaN	X	NaN	0	0	0	1

3.4. ALTFP_ADD_SUB Resource Utilization and Performance

The following lists the resource utilization and performance information for the ALTFP_ADD_SUB IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 14. ALTFP_ADD_SUB Resource Utilization and Performance for the Stratix Series of Devices

Device Family	Precision	Optimization	Output latency	Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	f _{MAX} (MHz)
Stratix IV	single	speed	7	594	376	385	228
			14	674	686	498	495
		area	7	576	345	375	227
			14	596	603	421	484
	double	speed	7	1,198	687	824	187
			14	997	1,607	1,080	398
		area	7	1,106	630	762	189
			14	904	1,518	1,013	265

3.5. ALTFP_ADD_SUB Design Example: Addition of Double-Precision Format Numbers

This design example uses the ALTFP_ADD_SUB IP core to perform the addition of double-precision format numbers using the parameter editor in the Intel Quartus Prime software.


Related Information

- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

3.5.1. ALTFP_ADD_SUM Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

Figure 15. ALTFP_ADD_SUB Simulation Waveform



This design example implements a floating-point adder for the addition of double-precision format numbers. All the optional input ports (clk_en and aclr) and optional output ports (overflow, underflow, zero, and nan) are enabled.

In this example, the output latency of the multiplier is set to 7 clock cycles. Every addition result appears at the result[] port 7 clock cycles after the input values are captured on the dataa[] and datab[] ports.

The following lists the inputs and corresponding outputs obtained from the simulation waveform.

Table 15. Summary of Input Values and Corresponding Outputs

Time	Event
0 ns, start-up	dataa[] value: 0000 0000 0000 0000h datab[] value: 7FF0 0000 0000 0000h Output value: All values seen on the output port before the 7th clock cycle are merely due to the behavior of the system during startup and should be disregarded.
4250 ns	Output value: 7FF0 0000 0000 0000h Exception handling ports: overflow asserts The addition of zero at the input port dataa[], and infinity value at the input port datab[] results in infinity value.
40,511 ns	dataa[] value: 0000 0000 0000 0000h

continued...

Time	Event
	datab[] value: 0000 0000 1000 0123h The is the addition of a zero and a denormal value.
43,750 ns	Output value: 0000 0000 0000 0000h Exception handling ports: zero remains asserted. Denormal inputs are not supported and are forced to zero before addition takes place. This results in a zero.

3.6. ALTFP_ADD_SUB Signals

Figure 16. ALTFP_ADD_SUB




Table 16. ALTFP_ADD_SUB Input Ports

Port Name	Required	Description
aclr	No	Asynchronous clear input for floating-point adder or subtractor. The source is asynchronously reset when the aclr signal is asserted high.
add_sub	No	Optional input port to enable dynamic switching between the adder and subtractor functions. The add_sub port must be used when the DIRECTION parameter is set to VARIABLE. When the add_sub port is high, result[] = dataa[] + datab[], otherwise, result[] = dataa[] - datab[].
clk_en	No	Clock enable to the floating-point adder or subtractor. This port allows addition or subtraction to occur when asserted high. When asserted low, no operations occur and the outputs are unchanged.
clock	Yes	Clock input to the IP core.
dataa[]	Yes	Data input to the floating-point adder or subtractor. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa bits. The size of this port is the total width of the sign bit, the exponent bits, and the mantissa bits.
datab[]	Yes	Data input to the floating-point adder or subtractor. This port is configured in the same way as dataa[].

Table 17. ALTFP_ADD_SUB Output Ports

Port Name	Required	Description
nan	Yes	NaN exception output. Asserted when an illegal addition or subtraction occurs, such as infinity minus infinity. When an invalid addition or subtraction occurs, a NaN value is output to the result[] port. Any adding or subtracting involving NaN values also produces a NaN value.
overflow	Yes	Overflow exception port. Asserted when the result of the addition or subtraction, after rounding, exceeds or reaches infinity. Infinity is defined as a number in which the exponent exceeds $2^{\text{WIDTH_EXP}} - 1$.
result[]	Yes	Floating-point output result. Like the input values, the MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits, and mantissa bits.
underflow	Yes	Underflow port for the adder or subtractor. Asserted when the result of the addition or subtraction, after rounding, the value is zero and the inputs are not equal. The underflow port is also asserted when the result is a denormalized number.
zero	No	Zero port for the adder or subtractor. Asserted when the result[] port is zero.

3.7. ALTFP_ADD_SUB Parameters

Table 18. ALTFP_ADD_SUB Parameters

Parameter Name	Type	Required	Description
DIRECTION	String	Yes	Specifies addition or subtraction operations. Values are ADD, SUB, or VARIABLE. If this parameter is not specified, the default is ADD. When the value is VARIABLE, the add_sub port determines whether the operation is addition or subtraction. The add_sub port must be connected if the DIRECTION parameter is set to VARIABLE. If the value is ADD or SUB, the add_sub port is ignored.
PIPELINE	Integer	No	Specifies the latency in clock cycles used in the ALTFP_ADD_SUB IP core. The PIPELINE parameter supports values of 7 through 14. If this parameter is not specified, the default value is 11. In general, a higher pipeline value produces better f_{MAX} performance.
ROUNDING	String	Yes	Specifies the rounding mode. The default value is TO_NEAREST. Other rounding modes are currently not supported.
OPTIMIZE	String	No	Defines the design preference, whether the design is optimized for speed (faster f_{MAX}), or optimized for area (lower resource count). Values are SPEED and AREA. If this parameter is not specified, the default is SPEED.
WIDTH_EXP	Integer	No	Specifies the precision of the exponent. The bias of the exponent is always set to $2^{\text{WIDTH_EXP}-1} - 1$ (that is, 127 for single-precision format and 1023 for double-precision format). The WIDTH_EXP parameter must be 8 for the single-precision mode and 11 for the double-precision mode, or a minimum of 11 for the single-extended precision mode. The WIDTH_EXP parameter must be less than the

continued...

Parameter Name	Type	Required	Description
			WIDTH_MAN parameter. The sum of WIDTH_EXP and the WIDTH_MAN parameters must be less than 64. If this parameter is not specified, the default is 8.
WIDTH_MAN	Integer	No	Specifies the precision of the mantissa. The WIDTH_MAN parameter must be 23 (to comply with the IEEE-754 standard for the single-precision mode) when the WIDTH_EXP parameter is 8. Otherwise, the WIDTH_MAN parameter must have a value that is greater than or equal to 31. The WIDTH_MAN parameter must be greater than the WIDTH_EXP parameter. The sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64. If this parameter is not specified, the default is 23.

4. ALTFP_DIV IP Core

This IP core performs floating-point division operation.

4.1. ALTFP_DIV Features

The ALTFP_DIV IP core offers the following features:

- Division functions.
- Optional exception handling output ports such as zero, division_by_zero, overflow, underflow, and nan.
- Optimization of speed and area.
- Low latency option.

4.2. ALTFP_DIV Output Latency

The output latency options for the ALTFP_DIV IP core differs depending on the precision selected, the width of the mantissa, or both. You have the choice of selecting the smaller figures of clock cycles delay in your design if the low latency option is desired.

Table 19. Latency Options for Each Operation

Precision	Mantissa Width	Latency (in clock cycles)
Single	23	6, 14, 33
Double	52	10, 24, 61
Single Extended	31 – 32	8, 18, 41
	33 – 34	8, 18, 43
	35 – 36	8, 18, 45
	37 – 38	8, 18, 47
	39 – 40	8, 18, 49
	41	10, 24, 41
	42	10, 24, 51
	43 – 44	10, 24, 53
	45 – 46	10, 24, 55
	47 – 48	10, 24, 57
	49 – 50	10, 24, 59
	51 – 52	10, 24, 61

4.3. ALTFP_DIV Truth Table

Table 20. Truth Table for Division Operations

DATAA[]	DATAB[]	SIGN BIT	RESULT[]	Overflow	Underflow	Zero	Division-by-zero	NaN
Normal	Normal	0/1	Normal	0	0	0	0	0
Normal	Normal	0/1	Denormal	0	0	1	0	0
Normal	Normal	0/1	Infinity	1	0	0	0	0
Normal	Normal	0/1	Zero	0	1	1	0	0
Normal	Denormal	0/1	Infinity	0	0	0	1	0
Normal	Zero	0/1	Infinity	0	0	0	1	0
Normal	Infinity	0/1	Zero	0	0	1	0	0
Normal	NaN	X	NaN	0	0	0	0	1
Denormal	Normal	0/1	Zero	0	0	1	0	0
Denormal	Denormal	0/1	NaN	0	0	0	0	1
Denormal	Zero	0/1	NaN	0	0	0	0	1
Denormal	Infinity	0/1	Zero	0	0	1	0	0
Denormal	NaN	X	NaN	0	0	0	0	1
Zero	Normal	0/1	Zero	0	0	1	0	0
Zero	Denormal	0/1	NaN	0	0	0	0	1
Zero	Zero	0/1	NaN	0	0	0	0	1
Zero	Infinity	0/1	Zero	0	0	1	0	0
Zero	NaN	X	NaN	0	0	0	0	1
Infinity	Normal	0/1	Infinity	0	0	0	0	0
Infinity	Denormal	0/1	Infinity	0	0	0	0	0
Infinity	Zero	0/1	Infinity	0	0	0	0	0
Infinity	Infinity	0/1	NaN	0	0	0	0	1
Infinity	NaN	X	NaN	0	0	0	0	1
NaN	Normal	X	NaN	0	0	0	0	1
NaN	Denormal	X	NaN	0	0	0	1	1
NaN	Zero	X	NaN	0	0	0	1	1
NaN	Infinity	X	NaN	0	0	0	0	1
NaN	NaN	X	NaN	0	0	0	0	1

4.4. ALTFP_DIV Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_DIV IP core. The information was derived using the Quartus II software version 10.0.

Table 21. ALTFP_DIV Resource Utilization and Performance for Stratix IV Devices

Device family	Precision	Optimizati on	Output latency	Logic Usage				f _{MAX} (MHz)
				Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-bit DSP	
Stratix IV	Single	Speed	33	3,593	3,351	2,500	—	313
		Area	33	1,646	2,074	1,441	—	308
	Double	Speed	61	13,867	13,143	10,196	—	292
		Area	61	5,125	7,360	4,842	—	267
Low Latency Option								
Stratix IV	Single	—	6	207	304	212	16	154
		—	14	253	638	385	16	358
	Double	—	10	714	1,077	779	44	151
		—	24	765	2,488	1,397	44	238

4.5. ALTFP_DIV Design Example: Division of Single-Precision

This design example uses the ALTFP_DIV IP core to implement a floating-point divider for the division of single-precision format numbers with low latency. This example uses the parameter editor to define the core.

Related Information


- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

4.5.1. ALTFP_DIV Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

Figure 17. ALTFP_DIV Simulation Waveform

This figure shows the expected simulation results in the ModelSim - Intel FPGA Edition software.



This design example implements a floating-point divider for the division of single-precision numbers with a low latency option. The output latency is 6, hence every division generates the output result 6 clock cycles later.

Table 22. Summary of Input Values and Corresponding Outputs

This table lists the inputs and corresponding outputs obtained from the simulation in the waveform.

Time	Event
0 ns, start-up	dataa[] value: 0000 0000h datab[] value: 0000 0000h Output value: The undefined value is seen on the result[] port, which is ignored. All values seen on the output port before the 6th clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
17600 ns	Output value: 7FC0 0000h Exception handling ports: nan asserts The division of zeros result in a NaN.
2000 ns	dataa[] value: 2D0B 496Ah datab[] value: 3A5A FC26h Both inputs hold normal values.
20800 ns	Output result: 321F 6EC6h Exception output ports: nan deasserts The division of two normal value results in a normal value.
11000 ns	dataa[] value: 046E 78BCh datab[] value: 6798 698Bh Both inputs hold normal values.
27200 ns	Output value: 0h Exception handling ports: underflow and zero asserts The division of the two normal values results in a denormal value. As denormal values are not supported, the result is zero and the underflow port asserts. The zero port is also asserted to indicate that the result is zero.
2600 ns	dataa[] value: 0D72 54A8h datab[] value: 0070 0000h The input port dataa[] holds a normal value while the input port datab[] holds a denormal value.
36800 ns	Output value: 7F80 0000h Exception handling ports: division_by_zero asserts Denormal numbers are forced-zero values, therefore, attempts to divide a normal value with a zero result in an infinity value.

4.6. ALTFP_DIV Signals

Figure 18. ALTFP_DIV Signals




Table 23. ALTFP_DIV Input Signals

Port Name	Required	Description
aclr	No	Asynchronous clear input for the floating-point divider. The source is asynchronously reset when the aclr signal is asserted high.
clock	Yes	Clock input to the IP core.
clk_en	No	Clock enable to the floating-point divider. This port enables division. This signal is active high. When this signal is low, no division takes place and the outputs remain the same.
dataa[]	Yes	Numerator data input. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits and mantissa bits.
datab[]	Yes	Denominator data input. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits and mantissa bits.

Table 24. ALTFP_DIV Output Signals

Port Name	Required	Description
result[]	Yes	Divider output port. The division result (after rounding). As with the input values, the MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits, and mantissa bits.
overflow	No	Overflow port for the divider. Asserted when the result of the division (after rounding) exceeds or reaches infinity. Infinity is defined as a number in which the exponent exceeds 2WIDTH_EXP-1.
underflow	No	Underflow port for the divider. Asserted when the result of the division (after rounding) is zero even though neither of the inputs to the divider is zero, or when the result is a denormalized number.
zero	No	Zero port for the divider. Asserted when the value of result[] is zero.
division_by_zero	No	Division-by-zero output port for the divider. Asserted when the value of datab[] is a zero.
nan	No	NaN port. Asserted when an invalid division occurs, such as infinity dividing infinity or zero dividing zero. A NaN value appears as output at the result[] port. Any division of a NaN value causes the nan output port to be asserted.

4.7. ALTFP_DIV Parameters

Table 25. ALTFP_DIV Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to $(2^{(WIDTH_EXP - 1)}) - 1$, that is, 127 for single precision and 1023 for double precision. The value of WIDTH_EXP must be 8 for single precision, 11 for double precision, and a minimum of 11 for single extended precision. The value of WIDTH_EXP must be less than the value of WIDTH_MAN, and the sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the precision of the mantissa. If this parameter is not specified, the default is 23. When WIDTH_EXP is 8 and the floating-point format is the single-precision format, the WIDTH_MAN value must be 23. Otherwise, the value of WIDTH_MAN must be a minimum of 31. The value of WIDTH_MAN must be greater than the value of WIDTH_EXP, and the sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
ROUNDING	String	Yes	Specifies the rounding mode. The default value is TO_NEAREST. The floating-point divider does not support other rounding modes.
OPTIMIZE	String	No	Specifies whether to optimize for area or for speed. Values are AREA and SPEED. A value of AREA optimizes the design using less total logic utilization or resources. A value of SPEED optimizes the design for better performance. If this parameter is not specified, the default value is SPEED.
PIPELINE	Integer	No	Specifies the number of clock cycles needed to produce the result. For the single-precision format, the latency options are 33, 14 or 6. For the double-precision format, the latency options are 61, 24 or 10. For the single-extended precision format, the value ranges from a minimum of 41 to a maximum of 61. For the low-latency option, the latency is determined from the mantissa width. For a mantissa width of 31 to 40 bits, the value is 8 or 18. For a mantissa width of 41 bits or more, the value is 10 or 24.

5. ALTFP_MULT IP Core

This IP core performs floating-point multiplication operation.

5.1. ALTFP_MULT IP Core Features

The ALTFP_MULT IP core offers the following features:

- Multiplication functions.
- Optional exception handling output ports such as zero, overflow, underflow, and nan.
- Optional dedicated multiplier circuitries in Cyclone and Stratix series.

5.2. ALTFP_MULT Output Latency

The output latency options for the ALTFP_MULT IP core are similar for all precisions.

Table 26. Latency Options for Each Precision Format

Precision	Mantissa Width	Latency (in clock cycles)
Single	23	5, 6, 10,11
Double	52	5, 6, 10,11
Single-Extended	31-52	5, 6, 10,11

5.3. ALTFP_MULT Truth Table

Table 27. Truth Table for Multiplier Operations

DATAA[]	DATAB[]	RESULT[]	Overflow	Underflow	Zero	NaN
Normal	Normal	Normal	0	0	0	0
Normal	Normal	Denormal	0	1	1	0
Normal	Normal	Infinity	1	0	0	0
Normal	Normal	Zero	0	1	1	0
Normal	Denormal	Zero	0	0	1	0
Normal	Zero	Zero	0	0	1	0
Normal	Infinity	Infinity	1	0	0	0
Normal	NaN	NaN	0	0	0	1
continued...						

DATAA[]	DATAB[]	RESULT[]	Overflow	Underflow	Zero	NaN
Denormal	Normal	Zero	0	0	1	0
Denormal	Denormal	Zero	0	0	1	0
Denormal	Zero	Zero	0	0	1	0
Denormal	Infinity	NaN	0	0	0	1
Denormal	NaN	NaN	0	0	0	1
Zero	Normal	Zero	0	0	1	0
Zero	Denormal	Zero	0	0	1	0
Zero	Zero	Zero	0	0	1	0
Zero	Infinity	NaN	0	0	0	1
Zero	NaN	NaN	0	0	0	1
Infinity	Normal	Infinity	1	0	0	0
Infinity	Denormal	NaN	0	0	0	1
Infinity	Zero	NaN	0	0	0	1
Infinity	Infinity	Infinity	1	0	0	0
Infinity	NaN	NaN	0	0	0	1
NaN	Normal	NaN	0	0	0	1
NaN	Denormal	NaN	0	0	0	1
NaN	Zero	NaN	0	0	0	1
NaN	Infinity	NaN	0	0	0	1
NaN	NaN	NaN	0	0	0	1

5.4. ALTFP_MULT Resource Utilization and Performance

The following tables list the resource utilization and performance information for the ALTFP_MULT IP core. The information was derived using the Quartus II software version 10.0.

Table 28. ALTFP_MULT Resource Utilization and Performance for Stratix IV Devices with Dedicated Multiplier Circuitry

Device Family	Precision	Output latency	Logic usage				f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-bit DSP	
Stratix IV	Single	5	138	148	100	4	274
		11	185	301	190	4	445
	Double	5	306	367	272	10	255
		11	419	523	348	10	395

5.5. ALTFP_MULT Design Example: Multiplication of Double-Precision Format Numbers

This design example uses the ALTFP_MULT IP core to compute the multiplication results of two double-precision format numbers. This example uses the parameter editor GUI to define the core.

Related Information


- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

5.5.1. ALTFP_MULT Design Example: Understanding the Simulation Waveform

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

Figure 19. ALTFP_MULT Simulation Waveform

This figure shows the expected simulation results in the ModelSim - Intel FPGA Edition software.



This design example implements a floating-point multiplier for the multiplication of double-precision format numbers. All the optional input ports (`clk_en` and `aclr`) and output ports (`overflow`, `underflow`, `zero`, and `nan`) are enabled.

In this example, the latency is set to 6 clock cycles. Therefore, every multiplication result appears at the result port 6 clock cycles later.

Table 29. Summary of Input Values and Corresponding Outputs

This table lists the inputs and corresponding outputs obtained from the simulation in the waveform.

Time	Event
0 ns, start-up	<code>dataa[]</code> value: 0000 0000 0000 0000h <code>datab[]</code> value: 4037 742C 3C9E ECC0h Output value: All values seen on the output port before the 6th clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
110 ns	Output value: 0000 0000 0000 0000h
continued...	

Time	Event
	Exception handling ports: zero asserts The multiplication of zero at the input port dataa[], and a non-zero value at the input port datab[] results in a zero.
600 ns	dataa[] value: 7FF0 0000 0000 0000h datab[] value: 4037 742C 3C9E ECC0h This is the multiplication of an infinity value and a normal value.
710 ns	Output value: 7FF0 0000 0000 0000h Exception handling ports: overflow asserts The multiplication of an infinity value and a normal value results in infinity. All multiplications with an infinity value results in infinity except when infinity is multiplied with a zero.

5.6. Parameters

Table 30. ALTFP_MULT Intel FPGA IP core Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	No	Specifies the value of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always $2^{(WIDTH_EXP - 1)} - 1$ (that is, 127 for the single-precision format and 1023 for the double-precision format). WIDTH_EXP must be 8 for the single-precision format or a minimum of 11 for the double-precision format and the single-extended precision format. WIDTH_EXP must be less than WIDTH_MAN. The sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
WIDTH_MAN	Integer	No	Specifies the value of the mantissa. If this parameter is not specified, the default is 23. When WIDTH_EXP is 8 and the floating-point format is single-precision, the WIDTH_MAN value must be 23; otherwise, the value of WIDTH_MAN must be a minimum of 31. The WIDTH_MAN value must always be greater than the WIDTH_EXP value. The sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
DEDICATED_MULTIPLIER_CIRCUITRY	String	No	Specifies whether to use dedicated multiplier circuitry. Values are AUTO, YES, or NO. If this parameter is not specified, the default is AUTO. If a device does not have dedicated multiplier circuitry, the DEDICATED_MULTIPLIER_CIRCUITRY parameter has no effect and defaults to NO.
PIPELINE	Integer	No	Specifies the number of clock cycles needed to produce the multiplied result. Values are 5, 6, 10, and 11. If this parameter is not specified, the default is 5.

5.7. ALTFP_MULT Signals

Table 31. ALTFP_MULT IP Core Input Signals

Port Name	Required	Description
clock	Yes	Clock input to the IP core.
clk_en	No	Clock enable. Allows multiplication to take place when asserted high. When signal is asserted low, no multiplication occurs and the outputs remain unchanged.
aclr	No	Synchronous clear. Source is asynchronously reset when asserted high.
dataa[]	Yes	Floating-point input data input to the multiplier. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of sign bit, exponent bits, and mantissa bits.
datab[]	Yes	Floating-point input data to the multiplier. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of sign bit, exponent bits, and mantissa bits.

Table 32. ALTFP_MULT IP Core Output Signals

Port Name	Required	Description
result[]	Yes	Output port for the multiplier. The floating-point result after rounding. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa.
overflow	No	Overflow port for the multiplier. Asserted when the result of the multiplication, after rounding, exceeds or reaches infinity. Infinity is defined as a number in which the exponent exceeds $2^{\text{WIDTH_EXP}} - 1$.
underflow	No	Underflow port for the multiplier. Asserted when the result of the multiplication (after rounding) is 0 while none of the inputs to the multiplication is 0, or asserted when the result is a denormalized number.
zero	No	Zero port for the multiplier. Asserted when the value of result[] is 0.
nan	No	NaN port for the multiplier. This port is asserted when an invalid multiplication occurs, such as the multiplication of infinity and zero. In this case, a NaN value is the output generated at the result[] port. The multiplication of any value and NaN produces NaN.

6. ALTFP_SQRT

This IP core performs square root calculation based on the input provided. You can use the ports and parameters available to customize the ALTFP_SQRT IP core according to your application.

6.1. ALTFP_SQRT Features

The ALTFP_SQRT IP core offers the following features:

- Square root functions.
- Optional exception handling output ports such as zero, overflow, and nan.

6.2. Output Latency

The output latency options for the ALTFP_SQRT IP core differs depending on the precision selected, the width of the mantissa, or both.

Table 33. Latency Options for Each Precision Format

Precision	Mantissa Width	Latency (in clock cycles)
Single	23	16, 28
Double	52	30, 57
Single-extended	31	20, 36
	32	20, 37
	33	21, 38
	34	21, 39
	35	22, 40
	36	22, 41
	37	23, 42
	38	23, 43
	39	24, 44
	40	24, 45
	41	25, 46
	42	25, 47
	43	26, 48
continued...		

Precision	Mantissa Width	Latency (in clock cycles)
	44	26, 49
	45	27, 50
	46	27, 51
	47	28, 52
	48	28, 53
	49	29, 54
	50	29, 55
	51	30, 56

6.3. ALTFP_SQRT Truth Table

Truth Table for Square Root Operations

DATA[]	SIGN BIT	RESULT[]	NaN	Overflow	Zero
Normal	0	Normal	0	0	0
Denormal	0/1	Zero	0	0	1
Positive Infinity	0	Infinity	0	1	0
Negative Infinity	1	All 1's	1	0	0
Positive NaN	0	All 1's	1	0	0
Negative NaN	1	All 1's	1	0	0
Zero	0/1	Zero	0	0	1
Normal	1	All 1's	1	0	0

6.4. ALTFP_SQRT Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_SQRT IP core. The information was derived using the Quartus II software version 10.0.

Table 34. ALTFP_SQRT Resource Utilization and Performance for Stratix IV Devices

Device Family	Precision	Output latency	Logic usage			f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	
Stratix IV	Single	28	502	932	528	472
	Double	57	2,177	3,725	2,202	366

6.5. ALTFP_SQRT Design Example: Square Root of Single-Precision Format Numbers

This design example uses the ALTFP_SQRT IP core to compute the square root of single-precision format numbers. This example uses the parameter editor in the Intel Quartus Prime software.

Related Information

- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

6.5.1. ALTFP_SQRT Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

These figures show the expected simulation results in the ModelSim - Intel FPGA Edition software.

Figure 20. ALTFP_SQRT Simulation Waveform (Input Data)





Figure 21. ALTFP_SQRT Simulation Waveform (Output Data)



This design example implements a floating-point square root function for single-precision format numbers with all the exception output ports instantiated. The output ports include overflow, zero, and nan.

The output latency is 28 clock cycles. Every square root computation generates the output result 28 clock cycles later.

Table 35. Summary of Input Values and Corresponding Outputs

This table lists the inputs and corresponding outputs obtained from the simulation in the waveforms.

Time	Event
0 ns, start-up	Output value: All values seen on the output port before the 28th clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
2 000 ns	data[] value: 2D0B 496Ah

continued...

Time	Event
	The data input is a normal number.
84 000 ns	Output value: 363C D4EBh The square root computation of a normal input results in a normal output.
14 000 ns	data[] value: 0000 0000h
96 000 ns	Output value: 0000 0000h Exception handling ports: zero asserts The square root computation of zero results in a zero.
23 000 ns	data[] value: 7F80 0000h The input is infinity.
105 000 ns	Output value: 7F80 0000h Exception handling ports: overflow asserts

6.6. ALTFP_SQRT Signals

Figure 22. ALTFP_SQRT Signals




Table 36. ALTFP_SQRT IP Core Input Signals

Port Name	Required	Description
clock	Yes	Clock input to the IP core.
clk_en	No	Clock enable that allows square root operations when the port is asserted high. When the port is asserted low, no operation occurs and the outputs remain unchanged.
aclr	No	Asynchronous clear. When the aclr port is asserted high, the function is asynchronously reset.
	Yes	Floating-point input data. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of sign bit, exponent bits, and mantissa bits.

Table 37. ALTFP_SQRT IP Core Output Signals

Port Name	Required	Description
result[]	Yes	Square root output port for the floating-point result. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits, and mantissa bits.
overflow	Yes	Overflow port. Asserted when the result of the square root (after rounding) exceeds or reaches infinity. Infinity is defined as a number in which the exponent exceeds $2^{\text{WIDTH_EXP}} - 1$.
zero	Yes	Zero port. Asserted when the value of the result[] port is 0.
nan	Yes	NaN port. Asserted when an invalid square root occurs, such as negative numbers or NaN inputs.

6.7. ALTFP_SQRT Parameters

Table 38. ALTFP_SQRT Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to $2^{\text{WIDTH_EXP}} - 1$, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of the WIDTH_EXP parameter must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of the WIDTH_EXP parameter must be less than the value of the WIDTH_MAN parameter, and the sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the value of the mantissa. If this parameter is not specified, the default is 23. When the WIDTH_EXP parameter is 8 and the floating-point format is single-precision, the WIDTH_MAN parameter value must be 23. Otherwise, the value of the WIDTH_MAN parameter must be a minimum of 31. The value of the WIDTH_MAN parameter must be greater than the value of the WIDTH_EXP parameter. The sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
ROUNDING	String	Yes	Specifies the rounding mode. The default value is TO_NEAREST. Other rounding modes are not supported.
PIPELINE	Integer	Yes	Specifies the number of clock cycles for the square root results of the result[] port. Values are WIDTH_MAN + 5 and $((\text{WIDTH_MAN} + 5/2) + 2)$ as specified by truncating the radix point.

7. ALTFP_EXP IP Core

This IP core performs exponential calculation based on the input provided.

7.1. ALTFP_EXP Features

The ALTFP_EXP IP core offers the following features:

- Exponential value of a given input.
- Optional exception handling output ports such as zero, overflow, underflow, and nan.

7.2. Output Latency

The output latency options for the ALTFP_EXP IP core differs depending on the precision selected, the width of the mantissa, or both.

Precision	Mantissa Width	Latency (in clock cycles)
Single	23	17
Double	52	25
Single-extended	31 – 38	22
	39 – 52	25

7.3. ALTFP_EXP Truth Table

Table 39. Truth Table for Exponential Operations

DATAA[]	Calculation	RESULT[]	NaN	Overflow	Underflow	Zero
Normal	edata	Normal	0	0	0	0
Normal	edata	Infinity	0	1	0	0
Normal (numbers of small magnitude)	edata	1	0	0	1	0
Normal (negative numbers of large magnitude)	edata	0	0	0	1	0
continued...						

Intel Corporation. All rights reserved. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

DATAA[]	Calculation	RESULT[]	NaN	Overflow	Underflow	Zero
Denormal	e0	1	0	0	0	0
Zero	e0	1	0	0	0	0
Infinity (+)	e+	Infinity	0	0	0	0
Infinity (-)	e-	0	0	0	0	1
NaN	—	NaN	1	0	0	0

7.4. ALTFP_EXP Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_EXP IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 40. ALTFP_EXP Resource Utilization and Performance for Stratix IV Devices

Device Family	Precision	Output Latency	Logic usage				f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-bit DSP	
Stratix IV	Single	17	631	521	448	19	284
	Double	25	4,104	2,007	2,939	46	279

7.5. ALTFP_EXP Design Example: Exponential of Single-Precision Format Numbers

This design example uses the ALTFP_EXP IP core to compute the exponential value of single-precision format numbers. This example uses the parameter editor in the Intel Quartus Prime software.

Related Information

- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

7.5.1. ALTFP_EXP Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

These figures show the expected simulation results in the ModelSim - Intel FPGA Edition software.

(2) Any denormal input is treated as a zero before going through the exponential process.

Figure 23. ALTFP_EXP Simulation Waveform (Input Data)





Figure 24. ALTFP_EXP Simulation Waveform (Output Data)



This design example implements a floating-point exponential for the single-precision format numbers. The optional input ports (`clk_en` and `ack_r`) and all four exception handling output ports (`nan`, `overflow`, `underflow`, and `zero`) are enabled.

For single-precision format numbers, the latency is fixed at 17 clock cycles. Therefore, every exponential operation outputs the results 17 clock cycles later.

Table 41. Summary of Input Values and Corresponding Outputs

This table lists the inputs and corresponding outputs obtained from the simulation in the waveforms.

Time	Event
0 ns, start-up	<code>data[]</code> value: 1A03 568Ch Output value: An undefined value is seen on the <code>result[]</code> port, which is ignored. All values seen on the output port before the 17th clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
82.5 ns	Output value: 3F80 0000h As the input value of 1A03568Ch is a very small number, it is seen as a value that is approaching zero, and the result approaches 1 (which is represented by 3F800000). Exponential operations carried out on numbers of very small magnitudes result in a 1 and assert the <code>underflow</code> flag. Exception handling ports: <code>underflow</code> asserts
30 ns	<code>data[]</code> value: F3FC DEFFh This is a normal negative value of a very large magnitude.
112.5 ns	Output value: 0000 0000h The outcome of exponential operations on negative numbers of very large magnitudes approaches zero. Exception handling ports: <code>underflow</code> remains asserted
60 ns	<code>data[]</code> value: 7F80 0000h This is a positive infinite value.
142.5 ns	Output value: 7F80 0000h The operation on positive infinite values results in infinity. Exception handling ports: <code>underflow</code> deasserts, <code>overflow</code> asserts
90 ns	<code>data[]</code> value: 7FC0 0000h This is a NaN.
172.5 ns	Output value: 7FC0 0000h

continued...

Time	Event
	The exponential of a NaN results in a NaN. Exception handling ports: nan asserts
120 ns	data[] value: C1D4 49BAh This is a normal value.
202.5 ns	Output value: 2C52 5981h The result is a normal value. Exception handling ports: nan deasserts

7.6. ALTFP_EXP Signals

Figure 25. ALTFP_EXP Signals




Table 42. ALTFP_EXP IP Core Input Signals

Port Name	Required	Description
aclr	No	Asynchronous clear. When the aclr port is asserted high the function is asynchronously reset.
clk_en	No	Clock enable. When the clk_en port is asserted high, an exponential value operation takes place. When this signal is asserted low, no operation occurs and the outputs remain unchanged.
clock	Yes	Clock input to the IP core.
data[]	Yes	Floating-point input data. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of the sign bit, exponent bits, and mantissa bits.

Table 43. ALTFP_EXP IP Core Output Signals

Port Name	Required	Description
result[]	Yes	The floating-point exponential result of the value at data[]. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits, and mantissa bits.
overflow	No	Overflow exception output. Asserted when the result of the operation (after rounding) is infinite.
underflow	No	Underflow exception output. Asserted when the result of the exponential approaches 1 (from numbers of very small magnitude), or when the result approaches 0 (from negative numbers of very large magnitudes).
zero	No	Zero exception output. Asserted when the value in the result[] port is zero.
nan	No	NaN exception output. Asserted when an invalid operation occurs. Any operation involving NaN also asserts the nan port.

7.7. ALTFP_EXP Parameters

Table 44. ALTFP_EXP IP Core Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to $2^{(WIDTH_EXP - 1)} - 1$, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of the WIDTH_EXP parameter must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of the WIDTH_EXP parameter must be less than the value of the WIDTH_MAN parameter, and the sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the value of the mantissa. If this parameter is not specified, the default is 23. When the WIDTH_EXP parameter is 8 and the floating-point format is single-precision, the WIDTH_MAN parameter value must be 23. Otherwise, the value of the WIDTH_MAN parameter must be a minimum of 31. The value of the WIDTH_MAN parameter must be greater than the value of the WIDTH_EXP parameter. The sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
PIPELINE	Integer	Yes	Specifies the amount of latency, expressed in clock cycles, used in the ALTFP_EXP IP core. Acceptable pipeline values are 17, 22, and 25 cycles of latency. Create the ALTFP_EXP IP core with the MegaWizard Plug-In Manager to calculate the value for this parameter.
ROUNDING	String	Yes	Specifies the rounding mode. The default value is T0_NEAREST. Other rounding modes are not supported.

8. ALTFP_INV IP Core

This IP core performs the function of $1/a$ where a is the given input.

8.1. ALTFP_INV Features

The ALTFP_INV IP core offers the following features:

- Inverse value of a given input.
- Optional exception handling output ports such as zero, division_by_zero, underflow, and nan.

8.2. Output Latency

The output latency options for the ALTFP_INVIP core differs depending on the precision selected, the width of the mantissa, or both.

Precision	Mantissa Width	Latency (in clock cycles)
Single	23	20
Double	52	27
Single Extended	31 – 39	20
	40 – 52	27

8.3. ALTFP_INV Truth Table

Table 45. Truth Table for Inverse Operations

DATA[]	SIGN BIT	RESULT[]	Underflow	Zero	Division_by_zero	NaN
Normal	0/1	Normal	0	0	0	0
Normal	0/1	Denormal	1	1	0	0
Normal	0/1	Infinity	0	0	0	0
Normal	0/1	Zero	1	1	0	0
Denormal	0/1	Infinity	0	0	1	0
continued...						

- (3) Any calculated or computed **denormal** output is replaced by a zero and asserts the zero and underflow flags.

DATA[]	SIGN BIT	RESULT[]	Underflow	Zero	Division_by_zero	NaN
Zero	0/1	Infinity	0	0	1	0
Infinity	0/1	Zero	0	1	0	0
NaN	X	NaN	0	0	0	1

8.4. ALTFP_INV Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_INV IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 46. ALTFP_INV Resource Utilization and Performance for Stratix IV Devices

Device Family	Precision	Output Latency	Logic usage				f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-Bit DSP	
Stratix IV	Single	20	401	616	373	16	412
	Double	27	939	1,386	912	48	203

8.5. ALTFP_INV Design Example: Inverse of Single-Precision Format Numbers

This design example uses the ALTFP_INV IP core to compute the inverse of single-precision format numbers. This example uses the parameter editor in the Intel Quartus Prime software.

Related Information

- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

8.5.1. ALTFP_INV Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

These figures show the expected simulation results in the ModelSim - Intel FPGA Edition software.

(4) Any denormal input is treated as a zero before going through the inverse process.

Figure 26. ALTFP_INV Simulation Waveform (Input Data)





Figure 27. ALTFP_INV Simulation Waveform (Output Data)



This design example implements a floating-point inverse for single-precision format numbers. The optional input ports (`clk_en` and `aclr`) and all four exception handling output ports (`division_by_zero`, `nan`, `zero`, and `underflow`) are enabled.

The latency is fixed at 20 clock cycles; therefore, every inverse operation outputs results 20 clock cycles later.

This table lists the inputs and corresponding outputs obtained from the simulation in the waveforms.

Table 47. Summary of Input Values and Corresponding Outputs

Time	Event
0 ns, start-up	data[] value: 34A2 E42Fh Output value: An undefined value is seen on the result[] port, which is ignored. All values seen on the output port before the 20th clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
97.5 ns	Output value: 4A49 2A2Fh Exception handling ports: division_by_zero deasserts The inverse of a normal number results in a normal value.
10 ns	data[] value: 7F80 0000h This is an infinity value.
107.5 ns	Output value: 0000 0000h Exception handling ports: zero asserts The inverse of an infinity value produces a zero.
60 ns	data[] value: 7FC0 0000h This is a NaN.
157.5 ns	Output value: 7FC0 0000h Exception handling ports: nan asserts The inverse of a NaN results in a NaN
70 ns	data[] value: 0000 1000h This is a denormal number.
167.5 ns	Output value: 7F80 0000h
continued...	

Time	Event
	Exception handling ports: nan deasserts, division_by_zero asserts Denormal numbers are forced-zero values, therefore, the inverse of a zero results in infinity.

8.6. Ports

Table 48. ALTFP_INV IP core Input Ports

Port Name	Required	Description
ac1r	No	Asynchronous clear. When the ac1r port is asserted high, the function is asynchronously cleared.
clk_en	No	Clock enable. When the clk_en port is asserted high, an inversion value operation takes place. When signal is asserted low, no operation occurs and the outputs remain unchanged.
clock	Yes	Clock input to the IP core.
data[]	Yes	Floating-point input data. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of the sign bit, exponent bits, and mantissa bits.

Table 49. ALTFP_INV IP core Output Ports

Port Name	Required	Description
result[]	Yes	The floating-point inverse result of the value at the data[] input port. The MSB is the sign, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits, and mantissa bits.
underflow	No	Underflow exception output. Asserted when the result of the inversion (after rounding) is a denormalized number.
zero	No	Zero exception output. Asserted when the value at the result[] port is a zero.
division_by_zero	No	Division-by-zero exception output. Asserted when the denominator input is a zero.
nan	No	NaN exception output. Asserted when an invalid inversion occurs, such as the inversion of NaN. In this case, a NaN value is output to the result[] port. Any operation involving NaN also asserts the nan port.

8.7. Parameters

Table 50. ALTFP_INV IP core Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to $2^{(WIDTH_EXP - 1)} - 1$, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of the WIDTH_EXP parameter must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of the WIDTH_EXP
continued...			

Parameter Name	Type	Required	Description
			parameter must be less than the value of the WIDTH_MAN parameter, and the sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the value of the mantissa. If this parameter is not specified, the default is 23. When the WIDTH_EXP parameter is 8 and the floating-point format is single-precision, the WIDTH_MAN parameter value must be 23. Otherwise, the value of the WIDTH_MAN parameter must be a minimum of 31. The value of the WIDTH_MAN parameter must be greater than the value of the WIDTH_EXP parameter. The sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
PIPELINE	Integer	Yes	Specifies the amount of latency in clock cycles used in the ALTFP_INV IP core. Create the ALTFP_INV IP core to calculate the value for this parameter.
ROUNDING	String	No	Specifies the rounding mode. The default value is TO_NEAREST. Other rounding modes are not supported.

9. ALTFP_INV_SQRT IP Core

This IP core performs inverse square root value of a given input.

9.1. ALTFP_INV_SQRT Features

The ALTFP_INV_SQRT IP core offers the following features:

- Inverse square root value of a given input.
- Optional exception handling output ports such as zero, division_by_zero, and nan.

9.2. Output Latency

The output latency options for the ALTFP_INV_SQRT IP core differs depending on the precision selected, the width of the mantissa, or both.

Table 51. Latency Options for Each Precision Format

Precision	Mantissa Width	Latency (in clock cycles)
Single	23	26
Double	52	36
Single-Extended	31– 39	26
	40 – 52	36

9.3. ALTFP_INV_SQRT Truth Table

Table 52. Truth Table for Inverse Square Root Operations

DATA[]	SIGN BIT	RESULT[]	Zero	Division_by_zero	NaN
Normal	0	Normal	0	0	0
Normal	1	NaN	0	0	1
Denormal	0/1	Infinity	0	1	0
continued...					

⁽⁵⁾ Any denormal input is treated as a zero before going through the inverse process.

DATA[]	SIGN BIT	RESULT[]	Zero	Division_by_zero	NaN
Zero	0/1	Infinity	0	1	0
Infinity	0/1	Zero	1	0	0
NaN	X	NaN	0	0	1

9.4. ALTFP_INV_SQRT Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_INV_SQRT IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 53. ALTFP_INV_SQRT Resource Utilization and Performance for Stratix IV Devices

Device Family	Precision	Output Latency	Logic usage				f _{MAX} (MHz)
			Adaptive Look-up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-Bit DSP	
Stratix IV	Single	26	502	658	430	22	413
	Double	36	1,324	1,855	1,209	78	209

9.5. ALTFP_INV_SQRT Design Example: Inverse Square Root of Single-Precision Format Numbers

This design example uses the ALTFP_INV_SQRT IP core to compute the inverse square root of single-precision format numbers. This example uses the parameter editor GUI to define the core.

Related Information

- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

9.5.1. ALTFP_INV_SQRT Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

These figures show the expected simulation results in the ModelSim - Intel FPGA Edition software.

Figure 28. ALTFP_INV_SQRT Simulation Waveform (Input Data)





Figure 29. ALTFP_INV_SQRT Simulation Waveform (Output Data)



This design example implements a floating-point inverse square root for single-precision format numbers. The optional input ports (`clk_en` and `aclr`) and all three exception handling output ports (`division_by_zero`, `nan`, and `zero`) are enabled.

The latency is fixed at 26 clock cycles. Therefore, every inverse square root operation outputs the results 26 clock cycles later.

This table lists the inputs and corresponding outputs obtained from the simulation in the waveforms.

Table 54. Summary of Input Values and Corresponding Outputs

Time	Event
0 ns, start-up	data[] value: 05AE 470Bh Output value: An undefined value is seen on the result[] port, which can be ignored. All values seen on the output port before the 26th clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
127.5 ns	Output value: 5C5B 64CEh The inverse square root of a normal number results in a normal value.
10 ns	data[] value: E8A7 E93Dh This is a negative normal value.
137.5 ns	Output value: FFC0 0000h Exception handling ports: nan asserts The inverse square root of a negative value produces a NaN.
20 ns	data[] value: 0000 0004h The is a denormal value.
147.5 ns	Output value: 7F80 0000h Denormal numbers are forced-zero values, therefore the inverse square root of zero results in infinity. Exception handling ports: nan deasserts, division_by_zero asserts
50 ns	data[] value: 7F80 0000h This is an infinity value.
177.5 ns	Output value: 0000 0000h The inverse square root of an infinity value produces a zero. Exception handling ports: zero asserts

9.6. Ports

Figure 30. ALTFP_INV_SQRT Signals




Table 55. ALTFP_INV_SQRT IP Core Input Signals

Port Name	Required	Description
aclr	No	Asynchronous clear. When the aclr port is asserted high, the function is asynchronously cleared.
clk_en	No	Clock enable. When the clk_en port is asserted high, an inversion value operation takes place. When signal is asserted low, no operation occurs and the outputs remain unchanged.
clock	Yes	Clock input to the IP core.
data[]	Yes	Floating-point input data. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of the sign bit, exponent bits, and mantissa bits.

Table 56. ALTFP_INV_SQRT IP Core Output Signals

Port Name	Required	Description
result[]	Yes	The floating-point inverse result of the value at the data[] input port. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. The size of this port is the total width of the sign bit, exponent bits, and mantissa bits.
zero	No	Zero exception output. Asserted when the value at the result[] port is a zero.
division_by_zero	No	Division-by-zero exception output. Asserted when the denominator input is a zero.
nan	No	NaN exception output. Asserted when an invalid inversion of square root occurs, such as the square root of a negative number. In this case, a NaN value is output to the result[] output port. Any operation involving a NaN produces a NaN.

9.7. Parameters

Table 57. ALTFP_INV_SQRT IP core Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to $2^{(WIDTH_EXP - 1)} - 1$, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of the WIDTH_EXP parameter must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of the WIDTH_EXP parameter must be less than the value of the WIDTH_MAN parameter, and the sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the value of the mantissa. If this parameter is not specified, the default is 23. When the WIDTH_EXP parameter is 8 and the floating-point format is single-precision, the WIDTH_MAN parameter value must be 23. Otherwise, the value of the WIDTH_MAN parameter must be a minimum of 31. The value of the WIDTH_MAN parameter must be greater than the value of the WIDTH_EXP parameter. The sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
PIPELINE	Integer	Yes	Specifies the amount of latency, expressed in clock cycles, used in the ALTFP_INV_SQRT IP core. Create the ALTFP_INV_SQRT IP core to calculate the value for this parameter.
ROUNDING	String	No	Specifies the rounding mode. The default value is TO_NEAREST. Other rounding modes are not supported.

10. ALTFP_LOG

This IP core performs natural logarithm function. You can use the ports and parameters available to customize the ALTFP_LOG IP core according to your application.

10.1. ALTFP_LOG Features

The ALTFP_LOG IP core offers the following features:

- Natural logarithm functions.
- Optional exception handling output ports such as zero and nan.

10.2. Output Latency

The output latency options for the ALTFP_LOG IP core differs depending on the precision selected, the width of the mantissa, or both.

Table 58. Latency Options for Each Precision Format

Precision	Mantissa Width	Latency (in clock cycles)
Single	23	21
Double	52	34
Single Extended	31–36	25
	37–42	28
	43–48	31
	49–52	34

10.3. ALTFP_LOG Truth Table

This table lists the truth table for the natural logarithm operation.

Table 59. Truth Table for Natural Logarithm Operations

DATA[]	SIGN BIT	RESULT[]	Zero	NaN
Normal	0	Normal	0	0
Normal	1	NaN ⁽⁶⁾	0	1
continued...				

⁽⁶⁾ The natural logarithm of a negative value is invalid. Therefore, the output produced is a NaN.

DATA[]	SIGN BIT	RESULT[]	Zero	NaN
1 ⁽⁷⁾	0	Zero	1	0
Denormal ⁽⁸⁾	0	Negative Infinity	0	0
Zero ⁽⁹⁾	0/1	Negative Infinity	0	0
Infinity	0	Positive Infinity	1	0
NaN	X	NaN	0	1

10.4. ALTFP_LOG Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_LOG IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 60. ALTFP_LOG Resource Utilization and Performance for Stratix IV Devices

Device Family	Precision	Output Latency	Logic usage				f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-Bit DSP	
Stratix IV	Single	21	1,950	1,864	1,378	8	385
	Double	34	5,451	6,031	4,151	64	211

10.5. ALTFP_LOG Design Example: Natural Logarithm of Single-Precision Format Numbers

This design example uses the ALTFP_LOG IP core to compute the natural logarithm of single-precision format numbers. This example uses the parameter editor GUI to define the core.

Related Information

- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

⁽⁷⁾ The "1" in this case is equivalent to $\ln 1$.

⁽⁸⁾ The value of positive denormalized numbers is a value that approximates zero, and the output produced is a negative infinity number.

⁽⁹⁾ The zero in this case represents zero special case of the IEEE standard. It is not equivalent to $\ln 0$, but instead approximates to it.

Time	Event
	As denormal numbers are not supported, the input is forced to zero before going through the logarithm function. The natural logarithm of zero is negative infinity.
45 ns	data[] value: 0080 0000h This is the smallest valid input. All the input bits are 0 except the LSB of the exponent field.
147.5 ns	Output value: C2AE AC50h
60 ns	data[] value: 3F80 0000h The input value 3F80 0000h is equivalent to the actual value, $1.0 \times 2^0 = 1$.
152.5 ns	Output value: 0000 0000h Exception handling ports: zero asserts Since $\ln 1$ results in zero, it produces an output of zero.

10.6. Signals

Figure 33. ALTFP_LOG Signals




Table 62. ALTFP_LOG IP Core Input Signals

Port Name	Required	Description
aclr	No	Asynchronous clear. When the aclr port is asserted high, the function is asynchronously cleared.
clk_en	No	Clock enable. When the clk_en port is asserted high, a natural logarithm operation takes place. When signal is asserted low, no operation occurs and the outputs remain unchanged. Deasserting clk_en halts operation until it is asserted again. Assert the clk_en signal for the number of clock cycles equivalent to the required output latency (PIPELINE parameter value) for the results to be shown at the output.
clock	Yes	Clock input to the IP core.
data[]	Yes	Floating-point input data. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of the sign bit, exponent bits, and mantissa bits.
continued...		

Port Name	Required	Description
		For single precision, the width is fixed to 32 bits. For double precision, the width is fixed to 64 bits. For single extended precision, you can choose a width in the range from 43 to 64 bits.

Table 63. ALTFP_LOG IP Core Output Signals

Port Name	Required	Description
result[]	Yes	The natural logarithm of the value on input data. The natural logarithm of the data[] input port, shown in floating-point format. The widths of the result[] output port and data[] input port are the same.
zero	No	Zero exception output. Asserted when the exponent and mantissa of the output port are zero. This occurs when the actual input value is 1 because $\ln 1 = 0$.
nan	No	NaN exception output. Asserted when the exponent and mantissa of the output port are all 1's and non-zero, respectively. This occurs when the input is a negative number or NaN.

10.7. Parameters

Table 64. ALTFP_LOG IP core Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to $2^{(WIDTH_EXP - 1)} - 1$, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of the WIDTH_EXP parameter must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of the WIDTH_EXP parameter must be less than the value of the WIDTH_MAN parameter, and the sum of the WIDTH_EXP and WIDTH_MAN parameters must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the precision of the mantissa. If this parameter is not specified, the default is 23. The value of WIDTH_MAN must be 23 for the single-precision format, and 52 for the double-precision format. For the single-extended precision format, the valid value ranges from 31 to 52. The value of WIDTH_MAN must be greater than the value of WIDTH_EXP, and the sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
PIPELINE	Integer	Yes	Specifies the amount of latency in clock cycles used in the ALTFP_LOG IP core. Create the ALTFP_LOG IP core to calculate the value for this parameter.

11. ALTFP_ATAN IP Core

This IP core performs arctangent calculation. You can use the ports and parameters available to customize the ALTFP_ATAN IP core according to your application.

11.1. Output Latency

The output latency option for the ALTFP_ATAN Intel FPGA IP core have a fixed latency level for single-precision format.

Table 65. Latency Option

Trigonometric Function	Precision	Mantissa Width	Latency (in clock cycles)
Arctangent	Single	23	34

11.2. ALTFP_ATAN Features

The ALTFP_ATAN IP core offers the following features:

- Arctangent value of a given angle, θ in unit radian.
- Support for single-precision floating point format.
- Support for optional input ports such as asynchronous clear (aclr) and clock enable (clk_en) ports.

11.3. ALTFP_ATAN Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_ATAN IP core. The information was derived using the Intel Quartus Prime software version 11.0.

Table 66. ALTFP_ATAN Resource Utilization and Performance

Device Family	Function	Precision	Output Latency	Logic usage				f _{MAX} (MHz)
				Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-Bit DSP	
Stratix V	ArcTangent	Single	36	2,454	1,010	1,303	27	255.49

11.4. Ports

Table 67. ALTFP_ATAN IP core Input Ports

Port Name	Required	Description
ac1r	No	Asynchronous clear. When the ac1r port is asserted high, the function is asynchronously cleared.
clk_en	No	Clock enable. When the clk_en port is asserted high, division takes place. When the signal is deasserted, no operation occurs and the outputs remain unchanged.
clock	Yes	Clock input to the IP core.
data[]	Yes	Floating-point input data. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of the sign bit, exponent bits, and mantissa bits.

Port Name	Required	Description
result[]	Yes	The result of the trigonometric function in floating-point format. The widths of the result[] output port and data[] input port are the same.

11.5. ALTFP_ATAN Parameters

Table 68. ALTFP_ATAN Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. The bias of the exponent is always set to $2(WIDTH_EXP-1) - 1$ (that is, 127 for single-precision format). The value of WIDTH_EXP must be 8 for single-precision format. The default value for WIDTH_EXP is 8.
WIDTH_MAN	Integer	Yes	Specifies the precision of the mantissa. The value of WIDTH_MAN must be 23 when WIDTH_EXP is 8. The default value for WIDTH_MAN is 23.
PIPELINE	Integer	Yes	The number of pipeline is fixed for the mantissa width and some internal parameter. For the correct settings, refer to Table 12-1 on page 12-2.
ROUNDING	Integer	No	Specifies the rounding mode. The default value is TO_NEAREST. Other rounding modes are not supported.

12. ALTFP_SINCOS IP Core

This IP core perform trigonometric Sine/Cosine functions. You can use the ports and parameters available to customize the ALTFP_SINCOS IP core according to your application.

12.1. ALTFP_SINCOS Features

The ALTFP_SINCOS IP core offers the following features:

- Implements sine and cosine calculations.
- Support for single-precision floating point format.
- Support for optional input ports such as asynchronous clear (ac1r) and clock enable (clk_en) ports.

12.2. Output Latency

The output latency options for the ALTFP_SINCOS IP core have a fixed latency level for sine and cosine functions.

Trigonometric Function	Precision	Mantissa Width	Latency (in clock cycles)
Sine	Single	23	36
Cosine	Single	23	36

Related Information

[ALTFP_SINCOS Parameters](#) on page 75

12.3. ALTFP_SINCOS Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_SINCOS IP core. The information was derived using the Intel Quartus Prime software version 10.1.

Table 69. ALTFP_SINCOS Resource Utilization and Performance

Device Family	Function	Precision	Output Latency	Logic usage				f _{MAX} (MHz)
				Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	18-Bit DSP	
Stratix IV	Sine	Single	36	2,859	2,190	1,830	16	292.96
	Cosine	Single	35	2,753	2,041	1,745	16	258.26

Intel Corporation. All rights reserved. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

12.4. ALTFP_SINCOS Signals

Figure 34. ALTFP_SINCOS Signals




Table 70. ALTFP_SINCOS IP Core Input Signals

Port Name	Required	Description
aclr	No	Asynchronous clear. When the aclr port is asserted high, the function is asynchronously cleared.
clk_en	No	Clock enable. When the clk_en port is asserted high, sine or cosine operation takes place. When the signal is asserted low, no operation occurs and the outputs remain unchanged.
clock	Yes	Clock input to the Intel FPGA IP core.
data[]	Yes	Floating-point input data. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of the sign bit, exponent bits, and mantissa bits.

Table 71. ALTFP_SINCOS IP Core Output Signals

Port Name	Required	Description
result[]	Yes	The trigonometric of the data[] input port in floating-point format. The widths of the result[] output port and data[] input port are the same.

12.5. ALTFP_SINCOS Parameters

Table 72. ALTFP_SINCOS IP Core Parameters

Parameter Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. The bias of the exponent is always set to $2^{(WIDTH_EXP-1)} - 1$ (that is, 127 for single-precision format). The value of WIDTH_EXP must be 8 for single-precision format and must be less than WIDTH_MAN. The available value for WIDTH_EXP is 8.
WIDTH_MAN	Integer	Yes	Specifies the precision of the mantissa. The value of WIDTH_MAN must be 23 when WIDTH_EXP is 8. Otherwise, WIDTH_MAN must be a minimum of 31. The value of WIDTH_MAN must be greater than WIDTH_EXP. The available value for WIDTH_MAN is 23.
PIPELINE	Integer	Yes	The number of pipeline is fixed for the mantissa width and some internal parameter. For the correct settings, refer to Output Latency.

Related Information

[Output Latency](#) on page 74

13. ALTFP_ABS IP Core

This IP core performs absolute value calculation for the given input.

13.1. ALTFP_ABS Features

The ALTFP_ABS IP core offers the following features:

- Absolute value of a given input.
- Optional exception handling output ports such as zero, division_by_zero, overflow, underflow, and nan.
- Carry-through exception ports from other floating-point modules that act as inputs to the ALTFP_ABS IP core.

13.2. ALTFP_ABS Output Latency

The output latency options for the ALTFP_ABS IP core are the same for all three precision formats—single, double, and single-extended. The options available are zero without pipeline, and 1 clock cycle.

13.3. ALTFP_ABS Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_ABS IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 73. ALTFP_ABS Resource Utilization and Performance for the Stratix III Device Family

Precision	Output Latency	Logic usage				f _{MAX} (MHz)
		Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	18-Bit DSP	Memory	
Single	0	0	0	0	0	The f _{MAX} of this IP core depends on the speed of the selected device
	1	0	36	0	0	
Double	0	0	0	0	0	
	1	0	68	0	0	

13.4. ALTFP_ABS Design Example: Absolute Value of Multiplication Results

This design example uses the ALTFP_ABS IP core to compute the absolute value of the multiplication result of single-precision format numbers. This example incorporates the ALTFP_MULT IP core and uses the parameter editor in the Intel Quartus Prime software.


Related Information

- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

13.4.1. ALTFP_ABS Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

Figure 35. ALTFP_ABS Simulation Waveform



This design example produces a floating-point absolute value function for the multiplication results of single-precision format numbers. All the optional input ports (clk_en and ackr) and optional output ports (overflow, underflow, zero, division_by_zero, and nan) are enabled.

In this example, the latency of the multiplier is set to five clock cycles, while none is being set for the absolute value function. Thus, the absolute value result only appears at the result[] port five cycles after the input values are captured on the input ports.

The dataaa[] and datab[] values in the simulation waveform above portray the two input values that are being fed to the multiplier. The value in the result[] port depicts the multiplication result that has gone through the absolute value operation.

This table lists the inputs and corresponding outputs obtained from the simulation.

Table 74. Summary of Input Values and Corresponding Outputs

Time	Event
0 ns, start-up	dataaa[] value: C080 0000h datab[] value: 4000 0000h
continued...	

Time	Event
	Output value: All values seen on the output port before the 5th clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
22.5 ns	Output value: 4100 0000h The multiplication of a negative number with a positive number results in a negative number. The absolute value of the result is reflected on the result[] port.
20 ns	dataa[] value: 579D F479h datab[] value: 7F80 0000h The value of dataa[] is normal while the value of datab[] is infinity.
42.5 ns	Output value: 7F80 0000h Exception handling ports: overflow asserts The multiplication of a normal value with infinity results in infinity and sets the overflow port in the multiplier. The absolute value of the output is infinity and the overflow port is also set as this assertion of the port is being carried through from the corresponding overflow port in the multiplier.

13.5. ALTFP_ABS Signals

Figure 36. ALTFP_ABS Signals




Table 75. ALTFP_ABS Input Signals

Port Name	Required	Description
aclr	No	Asynchronous clear. When the aclr port is asserted high, the function is asynchronously cleared.
clk_en	No	Clock enable. When the clk_en port is asserted high, an absolute value operation takes place. When the signal is asserted low, no operation occurs and the outputs remain unchanged.
clock	Yes	Clock input to the IP core.
data[]	Yes	Floating-point input data. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of sign bit, exponent bits, and mantissa bits.
continued...		

Port Name	Required	Description
zero_in	No	Zero exception input. Carry-through exception input port from other floating-point modules.
nan_in	No	NaN exception input. Carry-through exception input port from other floating-point modules.
overflow_in	No	Overflow exception input. Carry-through exception input port from other floating-point modules.
underflow_in	No	Underflow exception input. Carry-through exception input port from other floating-point modules.
division_by_zero_in	No	Division-by-zero exception input. Carry-through exception input port from other floating-point modules.

Table 76. ALTFP_ABS Output Signals

Port Name	Required	Description
result[]	Yes	The absolute value result of the input data. The size of this port corresponds to the size of the input data[] port.
zero	No	Zero exception output carried from the input. Asserted if the corresponding carry-through port from the input is asserted.
nan	No	NaN output carried from the input. Asserted if the corresponding carry-through port from the input is asserted.
overflow	No	Overflow exception output carried from the input. Asserted if the corresponding carry-through port from the input is asserted.
underflow	No	Underflow exception output carried from the input. Asserted if the corresponding carry-through port from the input is asserted.
division_by_zero	No	Division-by-zero exception output carried from the input. Asserted if the corresponding carry-through port from the input is asserted.

13.6. ALTFP_ABS Parameters

Table 77. ALTFP_ABS Parameters

Port Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to $2^{(WIDTH_EXP - 1)} - 1$, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of WIDTH_EXP must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of WIDTH_EXP must be less than the value of WIDTH_MAN, and the sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the precision of the mantissa. If this parameter is not specified, the default is 23. When WIDTH_EXP is 8 and the floating-point format is single-precision, the WIDTH_MAN value must be 23. Otherwise, the value of WIDTH_MAN must be a

continued...

Port Name	Type	Required	Description
			minimum of 31. The value of WIDTH_MAN must be greater than the value of WIDTH_EXP, and the sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
PIPELINE	Integer	Yes	Specifies the amount of latency, expressed in clock cycles, used in the ALTFP_ABS IP core. Create the ALTFP_ABS IP core with the parameter editor to calculate the value for this parameter.

14. ALTFP_COMPARE IP Core

This IP core performs comparison functions between two inputs.

14.1. ALTFP_COMPARE Features

The ALTFP_COMPARE IP core offers the following features:

- Comparison functions between two inputs.
- Seven status output ports:
 - aeb (input A is equal to input B).
 - aneb (input A is not equal to input B).
 - agb (input A is greater than input B).
 - ageb (input A is greater than or equal to input B).
 - alb (input A is less than input B).
 - aleb (input A is less than or equal to input B).
 - unordered (used as an output to flag if one or both input ports are NaN).

14.2. ALTFP_COMPARE Output Latency

The output latency options for the ALTFP_COMPARE IP core are the same for all three precision formats—single, double, and single-extended. The options available are 1, 2, and 3 clock cycles.

14.3. ALTFP_COMPARE Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_COMPARE IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 78. ALTFP_COMPARE Resource Utilization and Performance for Stratix IV Devices

Device Family	Precision	Output Latency	Logic Usage			f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Look-Up Modules (ALMs)	
Stratix IV	single	3	68	33	47	794
	double	3	121	47	87	680

14.4. ALTFP_COMPARE Design Example: Comparison of Single-Precision Format Numbers

This design example uses the ALTFP_COMPARE IP core to implement the comparison of single-precision format numbers using the parameter editor in the Intel Quartus Prime software.

Related Information


- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

14.4.1. ALTFP_COMPARE Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

This figure shows the expected simulation results in the ModelSim - Intel FPGA Edition software.

Figure 37. ALTFP_COMPARE Simulation Waveform



This design example implements a floating-point comparator for single-precision numbers. Both optional input ports (clk_en and aclr) and all seven output ports (ageb, aeb, agb, aneb, alb, aleb, and unordered) are enabled.

The chosen output latency is 3. Therefore, the comparison operation generates the output result 3 clock cycles later.

This table lists the inputs and corresponding outputs obtained from the simulation in the waveform.

Table 79. Summary of Input Values and Corresponding Outputs

Time	Event
0 ns, start-up	dataaa[] value: 619B CE11h datab[] value: 9106 CA22h Output value: An undefined value is seen on the result[] port, which is ignored. All values seen on the output port before the 3rd clock cycle are merely due to the behavior of the system during start-up and should be disregarded.
25 ns	Output ports: ageb, aneb, and agb assert
350 ns	dataaa[] value: 0060 0000h
continued...	

Time	Event
	datab[] value: 0070 0000h Both input values are denormal numbers.
375 ns	Output ports: aeb, ageb, and aleb assert Denormal inputs are not supported and are forced to zero before comparison takes place, which results in the dataa[] value being equal to datab[].
460 ns	The aclr signal is set for 1 clock cycle.
495.5 ns	The comparisons of subsequent data inputs are performed 3 clock cycles after the aclr signal deasserts.

14.5. ALTFP_COMPARE Signals

Figure 38. ALTFP_COMPARE Signals




Table 80. ALTFP_COMPARE Input Signals

Port Name	Required	Description
aclr	No	Asynchronous clear. The source is asynchronously reset when asserted high.
clk_en	No	Clock enable. When this port is asserted high, a compare operation takes place. When signal is asserted low, no operation occurs and the outputs remain unchanged.
clock	Yes	Clock input to the IP core.
dataa[]	Yes	Data input. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of sign bit, exponent bits, and mantissa bits.
datab[]	Yes	Data input. The MSB is the sign bit, the next MSBs are the exponent, and the LSBs are the mantissa. This input port size is the total width of sign bit, exponent bits, and mantissa bits.

Table 81. ALTFP_COMPARE Output Signals

Port Name	Required	Description
aeb	Yes	Output port for the comparator. Asserted if the value of the dataa[] port equals the value of the datab[] port.
agb	Yes	Output port for the comparator. Asserted if the value of the dataa[] port is greater than the value of the datab[] port.
ageb	Yes	Output port for the comparator. Asserted if the value of the dataa[] port is greater than or equal to the value of the datab[] port.
alb	Yes	Output port for the comparator. Asserted if the value of the dataa[] port is less than the value of the datab[] port.
aleb	Yes	Output port for the comparator. Asserted if the value of the dataa[] port is less than or equal to the value of the datab[] port.
aneb	Yes	Output port for the comparator. Asserted if the value of the dataa[] port is not equal to the value of the datab[] port.
unordered	Yes	Output port for the comparator. Asserted when either the dataa[] port and the datab[] port is set to NaN, or if both the dataa[] port and the datab[] port are set to NaN.

14.6. ALTFP_COMPARE Parameters

Table 82. ALTFP_COMPARE Parameters

Port Name	Type	Required	Description
WIDTH_EXP	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to 2 (WIDTH_EXP - 1) - 1, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of WIDTH_EXP must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of WIDTH_EXP must be less than the value of WIDTH_MAN, and the sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
WIDTH_MAN	Integer	Yes	Specifies the precision of the mantissa. If this parameter is not specified, the default is 23. When WIDTH_EXP is 8 and the floating-point format is single-precision, the WIDTH_MAN value must be 23. Otherwise, the value of WIDTH_MAN must be a minimum of 31. The value of WIDTH_MAN must be greater than the value of WIDTH_EXP, and the sum of WIDTH_EXP and WIDTH_MAN must be less than 64.
PIPELINE	Integer	Yes	Specifies the latency in clock cycles used in the ALTFP_COMPARE IP core. The pipeline values are 1, 2, and 3 latency in clock cycles.

15. ALTFP_CONVERT IP Core

This IP core performs conversion functions for various formats.

15.1. ALTFP_CONVERT Features

The ALTFP_CONVERT IP core offers the following features:

- Conversion functions for the following formats:
 - Integer-to-Float
 - Float-to-Integer
 - Float-to-Float
 - Fixed-to-Float
 - Float-to-Fixed
- Support for signed and unsigned integer
- Optional exception handling output ports such as overflow, underflow, and nan

Table 83. Supported Operations and Exception Ports

Operation	Supported Exception Ports
Integer-to-Float	Not supported
Float-to-Integer	overflow, underflow, and nan
Float-to-Float	overflow, underflow, and nan
Fixed-to-Float	Not supported
Float-to-Fixed	overflow, underflow, and nan

15.2. ALTFP_CONVERT Conversion Operations

This table lists the features of each conversion operation.

Table 84. ALTFP_CONVERT Conversion Operations

Operation	Features
Integer-to-Float Conversion	• Converts integers to the IEEE-754 standard floating-point representation. • Supports conversions of signed integers to floating-point numbers in single, double, and single-extended precision formats.
Float-to-Integer Conversion	• Converts IEEE-754 standard floating-point representations to the integer-bit format. • Supports conversions of single, double, and single-extended precision formats to signed integers.
<i>continued...</i>	

Intel Corporation. All rights reserved. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

Operation	Features
Float-to-Float Conversion	- Converts between IEEE-754 standard floating-point representations. - Supports conversions of between single double, and single-extended precision formats. - This operation offers the following modes: - Single-precision format to single-extended precision format or double-precision format. - Double-precision format to single-precision format or single-extended precision format. - Single-extended precision format to single-precision or double-precision format.
Fixed-to-Float Conversion	- Converts fixed-point format data to the IEEE-754 standard floating-point representation. - Supports conversions of fixed-point format data to floating-point numbers in single, double, and single-extended precision formats.
Float-to-Fixed Conversion	- Converts IEEE-754 standard floating-point representations to the fixed-point format. - Supports conversion of floating-point numbers in single, double, and single-extended precision formats.

15.3. ALTFP_CONVERT Output Latency

The output latency options for the all the conversion operations in the ALTFP_CONVERT IP core are fixed, except for the Float-to-Float operation.

Table 85. Latency Options for Each Operation

Operation	Conversion From	Latency (in clock cycles)
Integer-to-Float	N/A	6
Float-to-Integer	N/A	6
Float-to-Float	Single-precision format	2
	Double-precision format	3
	Single-extended precision format	3
Fixed-to-Float	N/A	6
Float-to-Fixed	N/A	6

15.4. ALTFP_CONVERT Resource Utilization and Performance

This table lists the resource utilization and performance information for the ALTFP_CONVERT IP core. The information was derived using the Intel Quartus Prime software version 10.0.

Table 86. ALTFP_CONVERT Resource Utilization and Performance for Stratix III Devices

Operation	Format	Pipeline	Logic Usage			f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	
Integer-to-Float	32-bit integer to single-precision	6	182	238	157	515
	32-bit integer to double-precision	6	150	139	123	510
	64-bit integer to single-precision	6	385	371	296	336
	64-bit integer to single-precision	6	393	461	344	336
Float-to-Integer	Single-precision to 32-bit integer	6	256	255	176	455
	Single-precision to 64-bit integer	6	417	361	257	311
	Double-precision to 32-bit integer	6	406	387	273	409
	Double-precision to 64-bit integer	6	535	480	362	309
Float-to-Float	Single-precision to double-precision	2	44	73	40	868
	Double-precision to single-precision	3	103	140	89	520
Fixed-to-Float	16.16 fixed-point to double-precision	6	182	238	155	519
	16.16 fixed-point to double-precision	6	150	139	122	513
	32.32 fixed-point to single-precision	6	384	371	296	334
	32.32 fixed-point to single-precision	6	393	461	336	333
continued...						

Operation	Format	Pipeline	Logic Usage			f _{MAX} (MHz)
			Adaptive Look-Up Tables (ALUTs)	Dedicated Logic Registers (DLRs)	Adaptive Logic Modules (ALMs)	
Float-to-Fixed	Single-precision to 16.16 fixed-point	6	319	261	210	438
	Single-precision to 32.32 fixed-point	6	469	367	288	315
	Double-precision to 16.16 fixed-point	6	579	393	402	365
	Double-precision to 32.32 fixed-point	6	695	486	474	306

15.5. ALTFP_CONVERT Design Example: Convert Double-Precision Floating-Point Format Numbers

This design example uses the ALTFP_CONVERT IP core to convert double-precision floating-point format numbers to 64-bit integers. This design example uses the parameter editor in the Intel Quartus Prime software.

Related Information


- [Floating-Point IP Cores Design Example Files](#) on page 23
- [Floating-Point IP Cores Design Examples](#)
Provides the design example files for the Floating-Point IP cores

15.5.1. ALTFP_CONVERT Design Example: Understanding the Simulation Results

The simulation waveform in this design example is not shown in its entirety. Run the design example files in the ModelSim - Intel FPGA Edition software to see the complete simulation waveforms.

This figure shows the expected simulation results in the ModelSim - Intel FPGA Edition software.

Figure 39. ALTFP_CONVERT Simulation Waveform



This design example implements a float-to-integer converter for converting double-precision floating-point format numbers to 64-bit integers. In this operation, the optional exception ports of `overflow`, `underflow`, and `nan` are available apart from the `result[]` port.

The latency for the float-to-integer operation is six clock cycles. Therefore, each conversion generates the output result six clock cycles after receiving the input value.

This table lists the inputs and corresponding outputs obtained from the simulation in the waveform.

Table 87. Summary of Input Values and Corresponding Outputs

Time	Event
0 ns, start-up	dataa[] value: C394 AD22 761B 9EE5h Output value: The <code>result[]</code> port displays 0 regardless of what the input value is. This value seen on the output port before the 6th clock cycle is merely due to the behavior of the system during start-up and should be disregarded.
55 ns	Output value: FAD4 B762 7918 46C0h
150 ns	dataa[] value: 000F 0000 5555 1111h This value is a denormal number.
205 ns	Denormal inputs are not supported and are forced to zero before conversion takes place.
300 ns	dataa[] value: 5706 40CF 0EC6 1176h
355 ns	Output value: 7FFF FFFF FFFF FFFFh Exception handling ports: <code>overflow</code> asserts. The <code>overflow</code> flag is triggered because the width of the resulting integer is more than the maximum width allowed, and the value seen on the <code>result[]</code> port is the standard value used to represent a positive overflow number.
350 ns	dataa[] value: C728 3147 8444 1F75h
405 ns	Output value: 8000 0000 0000 0000h Exception handling ports: <code>overflow</code> remains asserted. This is a standard value to represent a negative overflow number.
400 ns	dataa[] value: 145A 257C 895A B309h
455 ns	Output value: 0000 0000h Exception handling ports: <code>underflow</code> asserts. The input value triggers the <code>underflow</code> port because the exponent of the input value is less than the exponent bias of 1023.
500 ns	dataa[] value: FFFF 0000 DDDD 5555h This value is a NaN.
555 ns	Output value: 0000 0000h Exception handling ports: <code>nan</code> asserts.

15.6. ALTFP_CONVERT Signals

Figure 40. ALTFP_CONVERT Signals




Table 88. ALTFP_CONVERT Input Signals

Port Name	Required	Description
clock	Yes	The clock input to the ALTFP_CONVERT IP core.
clk_en	No	Clock enable that allows conversions to take place when asserted high. When asserted low, no operation occurs and the outputs are unchanged.
aclr	No	Asynchronous clear. The source is asynchronously reset when the aclr signal is asserted high.
dataa[]	Yes	Data input. The size of this input port depends on the WIDTH_DATA parameter value. If the operation mode value is INT2FLOAT or FIXED2FLOAT, the data on the input bus is an integer. If the operation mode value is FLOAT2INT or FLOAT2FIXED, the input bus is the IEEE floating-point representation. In the single-precision format, the input bus width value is 32. In the double-precision format, the input bus width value is 64. In the single-extended precision format, the input bus range is from 43 to 64. If the operation mode value is FLOAT2FLOAT, the input bus value is the IEEE floating-point representation. In the single-precision format, the input bus width value is 32. In the double-precision format, the input bus width value is 64. In the single-extended precision format, the input bus range is from 43 to 64.

Table 89. ALTFP_CONVERT Output Signals

Port Name	Required	Description
result[]	Yes	Output for the floating-point converter. The size of this output port depends on the WIDTH_RESULT parameter value. If the operation mode value is FLOAT2INT or FLOAT2FIXED, the output bus is an IEEE floating-point representation. If the operation mode is FLOAT2INT, the output bus is an integer representation. If the selected precision is the single-precision format, the output bus width value is 32. If the selected precision is the double-precision format, the output bus width value is 64. If the selected precision is the single-extended precision format, the input bus range is from 43 to 64.

continued...

Port Name	Required	Description
		If the operation mode value is FLOAT2FLOAT, the output bus is an IEEE floating-point representation. If the selected precision is the single-precision format, the output bus is in the 64-bit double-precision format. If the selected precision is the double-precision format, the output bus is in the 32-bit single-precision format. If the selected precision is the single-extended precision format, the output bus ranges from 43 to 64.
overflow	No	Optional overflow exception output. This port is available only when the operation mode values are FLOAT2FIXED, FLOAT2INT, or FLOAT2FLOAT. Asserted when the result of the conversion (after rounding), exceeds the maximum width of the result[] port, or when the dataa[] input is infinity.
underflow	No	Optional underflow exception output. This port is available only when the operation mode values are FLOAT2FIXED, FLOAT2INT, or FLOAT2FLOAT. Asserted when the result of the conversion, after rounding, is fractional. In FLOAT2INT operations, this port is asserted when the exponent value of the floating-point input is smaller than the exponent bias. In FLOAT2FLOAT operations, this port is asserted when the floating-point input has a value smaller than the lowest exponent limit of the target floating-point format.
nan	No	Optional NaN exception output. This port is available only when the operation mode values are FLOAT2INT, FLOAT2FLOAT, or FLOAT2FIXED. Asserted when the input port is a NaN representation. If the operation mode value is FLOAT2INT or FLOAT2FIXED, the result[] port is set to zero. If the operation mode value is FLOAT2FLOAT, the result[] port is set to a NaN representation.

15.7. ALTFP_CONVERT Parameters

Table 90. ALTFP_CONVERT Parameters

Port Name	Type	Required	Description
WIDTH_EXP_INPUT	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to 2 (WIDTH_EXP - 1) - 1, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of WIDTH_EXP_INPUT must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of WIDTH_EXP_INPUT must be less than the value of WIDTH_MAN_INPUT, and the sum of WIDTH_EXP_INPUT and WIDTH_MAN_INPUT must be less than 64. These settings apply only to the FLOAT2FIXED, FLOAT2INT, and FLOAT2FLOAT operation modes.
WIDTH_MAN_INPUT	Integer	Yes	Specifies the precision of the mantissa. If this parameter is not specified, the default is 23. When WIDTH_EXP_INPUT is 8 and the floating-point format is single-precision, the WIDTH_MAN_INPUT value must be 23. Otherwise, the value of WIDTH_MAN_INPUT must be a minimum of 31. The value of WIDTH_MAN_INPUT must be greater than the value of WIDTH_EXP_INPUT, and the sum of WIDTH_EXP_INPUT and WIDTH_MAN_INPUT must be less than 64. These settings apply only to the FLOAT2FIXED, FLOAT2INT, and FLOAT2FLOAT operation modes.
WIDTH_INT	Integer	Yes	Specifies the integer width. If the operation is FIXED2FLOAT or INT2FLOAT, this parameter defines the integer width on the input side. If the operation is FLOAT2INT or FLOAT2FIXED, this parameter defines the result width on the output side. The available settings are 32 bits, 64 bits or n bits. For n bits settings, the range is from 4 bits to 64 bits.

continued...

Port Name	Type	Required	Description
			If unspecified, the default setting for WIDTH_INT is 32 bits.
WIDTH_DATA	Integer	Yes	Specifies the input data width. If the operation is INT2FLOAT, the WIDTH_DATA is also WIDTH_INT. If the operation is FIXED2FLOAT, the data width value is WIDTH_INT + fractional width. If the operation is FLOAT2FIXED, FLOAT2INT or FLOAT2FLOAT, the data width value is WIDTH_EXP_INPUT + WIDTH_MAN_INPUT + 1. The available settings are 32 bits, 64 bits or n bits. For n bits settings, the range is from 4 bits to 64 bits. If unspecified, the default setting for WIDTH_DATA is 32 bits.
WIDTH_EXP_OUTPUT	Integer	Yes	Specifies the precision of the exponent. If this parameter is not specified, the default is 8. The bias of the exponent is always set to 2 (WIDTH_EXP - 1) - 1, that is, 127 for the single-precision format and 1023 for the double-precision format. The value of WIDTH_EXP_OUTPUT must be 8 for the single-precision format, 11 for the double-precision format, and a minimum of 11 for the single-extended precision format. The value of WIDTH_EXP_OUTPUT must be less than the value of WIDTH_MAN_OUTPUT, and the sum of WIDTH_EXP_OUTPUT and WIDTH_MAN_OUTPUT must be less than 64. These settings apply only to the FLOAT2FIXED, FLOAT2INT, and FLOAT2FLOAT operation modes.
WIDTH_MAN_OUTPUT	Integer	Yes	Specifies the precision of the mantissa. If this parameter is not specified, the default is 23. When WIDTH_EXP_OUTPUT is 8 and the floating point format is single-precision, the WIDTH_MAN_OUTPUT value must be 23. Otherwise, the value of WIDTH_MAN_OUTPUT must be a minimum of 31. The value of WIDTH_MAN_OUTPUT must be greater than the value of WIDTH_EXP_OUTPUT, and the sum of WIDTH_EXP_OUTPUT and WIDTH_MAN_OUTPUT must be less than 64. These settings apply only to the FLOAT2FIXED, FLOAT2INT, and FLOAT2FLOAT operation modes.
WIDTH_RESULT	Integer	Yes	Specifies the width of the output result. In an INT2FLOAT, FLOAT2FLOAT, or FIXED2FLOAT operation, the result width is WIDTH_EXP_OUTPUT + WIDTH_MAN_OUTPUT + 1. In a FLOAT2INT operation, the result width is the value of the WIDTH_INT parameter. In a FLOAT2FIXED operation, this parameter is the result width. The available settings are 32 bits, 64 bits or n bits. For n bits settings, the range is from 4 bits to 64 bits.
ROUNDING	Integer	Yes	Specifies the rounding mode. The default value is TO_NEAREST. Other modes are not supported.
OPERATION	Integer	Yes	Specifies the operating mode. Values are INT2FLOAT, FLOAT2INT, FLOAT2FLOAT, FLOAT2FIXED, and FIXED2FLOAT. If this parameter is not specified, the default value is INT2FLOAT. When set to INT2FLOAT, the conversion of an integer input to an IEEE floating-point representation output takes place. When set to FLOAT2INT, the conversion of an IEEE floating-point representation input to an integer output takes place. When set to FLOAT2FLOAT, the conversion between IEEE floating-point representations input and output takes place. When set to FIXED2FLOAT, the conversion of a fixed point input to an IEEE floating-point representation output takes place. When set to FLOAT2FIXED, the IEEE floating-point input conversion to fixed point representation output takes place.

16. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Core

This IP core replaces all the functions supported by the existing floating point IP cores shown in the previous chapters in this document for devices in the Intel Quartus Prime Pro Edition software.

Table 91. Floating Point Functions Intel FPGA IP Release Information

Item	Description
Version	19.1
Intel Quartus Prime Version	20.1
Release Date	2020.04.13

Table 92. List of Functions Supported by FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP

Function	Description
Arithmetic	
Add	Two input addition
Sub	Two input subtraction
Add/Sub	Two input addition and or subtraction supporting two configurations. By default, the IP core computes both the sum (output port q) and the difference (output port s) of the two inputs. You may also select to generate an IP core that provides a single shared output for both the sum and the difference of inputs. In this configuration an additional input port (opSel) allows you to select the desired operation at run-time. When opSel is low (0), the IP core computes the sum of the two inputs. When opSel is high (1), it computes their difference. Enable this configuration by turning on Use Select Signal .
Multiply	Two input multiplication
Divide	Two input division
Reciprocal	Performs the function of $1/a$ where a is the input. <i>Note:</i> This function replaces the ALTFP_INV IP core in Intel Arria 10 devices.
Absolute	Generates absolute value of the input
Scalar Product	Performs the scalar product between two vectors with the dimensions that you set
Multiply-Accumulate	Two input multiplication followed by a single cycle accumulation
Accumulate	Perform single input accumulation in a single cycle
Multiply-Add	Performs two input multiplication followed by addition
<i>continued...</i>	

Function	Description
Complex-Multiply	Performs multiplication of two complex values
Roots	
Square Root	Performs square root to the input value
Reciprocal Square Root	Performs the function of $1/\sqrt{a}$ where a is the input <i>Note:</i> This function replaces the ALTFP_INV_SQRT IP core in Intel Arria 10 devices.
Cube Root	Performs cube root to the input value
3D Hypotenuse	Performs the function of $Q = \sqrt{a^2 + b^2 + c^2}$
Conversions	
Fixed-Point to Floating-Point	Converts a fixed point input to floating point representation
Floating-Point to Fixed-Point	Converts a floating-point input to fixed point representation
Floating to Floating Point	Converts a floating-point input to floating-point representation of a different precision
Comparisons	
Minimum	Compares and produces the smallest value of two inputs
Maximum	Compares and produces the biggest value of two inputs
Less Than	Compares and returns true if input a is less than input b
Less Than or Equal	Compares and returns true if input a is less than or equal to input b
Equal	Compares and returns true if input a is equal to input b
Greater Than	Compares and returns true if input a is greater than input b
Greater Than or Equal	Compares and returns true if input a is greater than or equal to input b
Not Equal	Compares and returns true if input a is not equal to input b
Exp/Log/Pow	
Exponent	Performs the function of e^a where a is the input
Exponent base 2	Performs the function of 2^a where a is the input
Exponent base 10	Performs the function of 10^a where a is the input
Log	Performs the function of $\log_e(a)$ where a is the input
$\log_2$	Performs the function of $\log_2(a)$ where a is the input
$\log_{10}$	Performs the function of $\log_{10}(a)$ where a is the input
$\log(1+x)$	Performs the function of $\log_e(1+a)$ where a is the input
Power	
LdExp	Sets the exponential value of a floating-point input Performs the function $a * 2^b$ where a is a floating-point input and b is an integer input. You specify both the format of input a and the width of input b
Trigonometry	
Sin	Performs sine function of a single input
Cos	Performs cosine function of a single input
Tan	Performs tangent function of a single input
continued...	

Function	Description
Arcsin	Performs arc sine function of a single input
Arccos	Performs arc cosine function of a single input
Arctan	Performs arc tangent function of a single input
Arctan2	Performs the function of arctan (b/a) where a and b are the inputs

16.1. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Features

The FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP core offers the following features:

- Supports both latency and frequency driven cores.
- Supports VHDL code generation.

16.2. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Output Latency

If you require a specific latency, follow these steps:

1. In the FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP parameter editor, click the **Basic** tab.
2. Under the Performance category, in the **Goal** option, select **latency**.
3. In the **Target** field, set your desired latency (cycles).
4. Then, click **Check Performance**.

16.3. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Target Frequency

If you require a specific frequency, follow these steps:

1. In the FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP parameter editor, click **Basic** tab.
2. Under the Performance category, in the **Goal** option, select **frequency**.
3. In the **Target** field, set your desired frequency (MHz).
4. The IP core reports the latency for the instance that it generates in the Report category.

Note: You must verify the frequency by running the Timing Analyzer.

16.4. FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP Combined Target

If you require a combined target of latency and frequency, follow these steps:

1. In the FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP parameter editor, click the **Basic** tab.
2. Under the Performance category, in the **Goal** option, select **Combined**.
3. In the **Target** field, set your desired frequency (MHz).
4. In the **Target** field, set your desired latency (cycles).
5. Then, click **Finish**.

16.5. FP_FUNCTIONS Intel FPGA IP Resource Utilization and Performance

These tables list the resource utilization and performance information for the FP_Functions Intel FPGA IP core. The information was derived using the Intel Quartus Prime software version 14.1. The frequency target was set to 200 MHz.

Table 93. Arithmetic

Family	Function	Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
									Primary	Secondary
Arria V (5AGXFB 3H4F40 C5)	Abs	Single	0	—	33	0	—	0	0	0
		Double	0	—	65	0	—	0	0	0
	Add	Single	9	233.1	360	0	—	0	507	29
		Double	12	251.95	886	0	—	0	1064	61
	AddSubtract	Single	9	249.31	477	0	0	0	651	63
		Double	12	252.46	1161	0	0	0	1713	91
	Cube Root	Single	9	275.18	132	6	—	2	132	20
		Double	24	185.77	634	17	—	10	1297	58
	Divide	Single	18	249	456	5	—	4	771	100
		Double	35	185.29	1409	39	—	15	3035	138
	Exp base 10	Single	16	212.72	547	3	—	2	675	18
		Double	31	185.77	2194	0	—	10	2626	56
Arria V (5AGXFB 3H4F40 C5)	Exp base 2	Single	7	236.41	345	0	—	2	214	19
		Double	21	185.84	932	0	—	10	1324	51
	Exp base e	Single	14	217.96	718	0	—	2	597	46
		Double	28	185.87	2134	0	—	10	2398	46
	Reciprocal	Single	12	253.16	210	4	—	3	294	26
		Double	30	185.29	877	9	—	14	1764	105
	Reciprocal Square Root	Single	7	267.52	118	4	—	2	141	14
		Double	20	185.74	539	13	—	9	1210	52
	LDExp	Single	2	367.92	69	0	—	0	85	0
		Double	2	359.32	100	0	—	0	146	0

continued...

Family	Function	Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
									Primary	Secondary
Arria V (5AGXFB 3H4F40 C5)	Log base 10	Single	16	250	379	4	--	3	622	65
		Double	34	186.12	1,380	40	--	11	3,025	143
Arria V (5AGXFB 3H4F40 C5)	Log(1+x)	Single	21	222.77	766	4	--	3	1,171	82
		Double	43	185.94	2,361	40	--	14	4,702	183
	Log base 2	Single	16	232.29	350	4	--	3	584	64
		Double	37	185.32	1,342	13	--	17	3,156	121
	Log base e	Single	16	248.57	379	4	--	3	616	57
		Double	35	185.84	1,422	40	--	13	3,066	147
	Multiply	Single	5	281.14	156	0	--	1	152	6
		Double	7	186.01	339	0	--	4	549	13
Arria V (5AGXFB 3H4F40 C5)	Power	Single	45	201.82	1,347	11	--	14	2,410	165
		Double	82	185.43	4,195	20	--	38	8,149	266
	Square Root	Single	8	261.92	119	3	--	2	174	13
		Double	21	185.94	548	8	--	9	1,225	44
	Subtract	Single	9	232.67	363	0	--	0	505	32
		Double	12	257.07	884	0	--	0	1,064	61
Cyclone V (5CGXFC 7D6F31 C7)	Abs	Single	0	--	33	0	--	0	0	0
		Double	0	--	65	0	--	0	0	0
	Add	Single	12	225.94	403	0	--	0	562	35
		Double	20	208.99	932	0	--	0	1,813	72
	AddSubtract	Single	12	224.67	509	0	--	0	805	65
		Double	20	211.55	1,197	0	--	0	2,647	120
	Cube Root	Single	10	230.47	131	6	--	2	213	11
		Double	34	212.49	890	17	--	10	1,991	54
	Divide	Single	20	232.61	466	5	--	4	991	62
		Double	51	201.01	1,782	41	--	15	4,317	165
Cyclone V (5CGXFC 7D6F31 C7)	Exp base 10	Single	20	217.58	552	3	--	2	905	32
		Double	52	212.77	2,317	0	--	10	4,287	122
	Exp base 2	Single	9	211.33	352	0	--	2	314	13
		Double	36	219.3	1,128	0	--	10	2,364	87
	Exp base e	Single	17	207.68	698	0	--	2	860	31
		Double	50	198.85	2,309	0	--	10	4,300	126
	Reciprocal	Single	14	230.95	245	4	--	3	378	26
		Double	44	207.43	1,201	9	--	14	2,694	94
continued...										

Family	Function	Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
									Primary	Secondary
	Reciprocal Square Root	Single	9	233.37	137	4	--	2	223	25
		Double	30	250	782	13	--	9	1,932	46
Cyclone V (5CGXFC7D6F31C7)	LDExp	Single	2	346.02	69	0	--	0	87	1
		Double	3	357.91	104	0	--	0	215	0
	Log base 10	Single	22	203.33	486	4	--	3	1,066	47
		Double	49	196.97	1,888	40	--	11	4,483	153
	Log(1+x)	Single	29	191.5	944	4	--	3	1,844	105
		Double	62	168.27	3,012	40	--	14	6,899	210
	Log base 2	Single	20	202.1	413	4	--	3	918	50
		Double	54	194.21	1,898	13	--	17	4,732	151
Cyclone V (5CGXFC7D6F31C7)	Log base e	Single	22	181.42	482	4	--	3	1,058	45
		Double	50	196.27	1,941	40	--	13	4,611	197
	Multiply	Single	6	268.6	159	0	--	1	223	2
		Double	11	205.17	431	0	--	4	970	18
	Power	Single	62	181.19	1,778	11	--	14	3,562	154
		Double	127	186.53	5,411	22	--	38	12,361	325
	Square Root	Single	8	219.15	126	3	--	2	205	12
		Double	31	250	822	8	--	9	2,056	55
	Subtract	Single	12	232.07	399	0	--	0	566	42
		Double	20	204.25	918	0	--	0	1,839	60
Stratix V (5SGXE A7K2F40C2)	Abs	Single	0	--	33	--	0	0	0	0
		Double	0	--	65	--	0	0	0	0
	Add	Single	5	364.83	366	--	0	0	299	19
		Double	7	329.49	834	--	0	0	801	53
	AddSubtract	Single	5	354.74	489	--	0	0	411	29
		Double	7	338.41	1,106	--	0	0	1,039	134
	Cube Root	Single	8	420.17	114	--	5	2	124	11
		Double	20	277.7	520	--	11	10	997	17
	Divide	Single	13	363.5	377	--	3	4	591	71
		Double	23	270.86	1,091	--	20	15	2,274	120
	Exp base 10	Single	11	292.4	486	--	3	2	417	12
		Double	22	271.74	2,033	--	0	10	1,761	48
Stratix V (5SGXE A7K2F40C2)	Exp base 2	Single	5	387.3	351	--	0	2	160	1
		Double	17	279.56	897	--	0	10	995	27
continued...										

Family	Function	Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
									Primary	Secondary
	Exp base e	Single	8	284.09	653	--	0	2	350	18
		Double	23	268.38	2,043	--	0	10	1,710	44
	Reciprocal	Single	9	279.33	199	--	3	3	211	13
		Double	22	241.31	764	--	9	14	1,391	49
	Reciprocal Square Root	Single	6	420.52	105	--	3	2	129	9
		Double	17	271.37	449	--	8	9	1,009	47
	LDExp	Single	0	--	67	--	0	0	0	0
		Double	0	717.36	99	--	0	0	66	0
	Log base 10	Single	11	359.58	358	--	3	3	443	29
		Double	23	271.96	1,077	--	20	11	2,252	101
Stratix V (5SGXE A7K2F40 C2)	Log(1+x)	Single	15	338.64	748	--	3	3	905	55
		Double	27	280.98	1,911	--	20	13	3,301	122
	Log base 2	Single	11	340.37	304	--	3	3	392	15
		Double	27	258.33	1,053	--	8	16	2,241	110
	Log base e	Single	11	351.86	359	--	3	3	439	35
		Double	23	270.49	1,071	--	20	13	2,210	94
	Multiply	Single	3	399.52	136	--	0	1	72	1
		Double	4	250.75	312	--	0	4	237	5
	Power	Single	31	261.23	1,171	--	8	12	1,492	83
		Double	60	267.81	3,555	--	13	37	5,347	244
	Square Root	Single	6	393.7	112	--	3	2	129	7
		Double	17	274.12	458	--	8	9	1,019	41
	Subtract	Single	5	320.41	360	--	0	0	299	14
		Double	7	338.52	835	--	0	0	801	51
Arria 10 (10AX11 5H4F34I 3SP)	Abs	Single	0	--	33	--	0	0	0	0
		Double	0	--	65	--	0	0	0	0
	Add	Single	4	296.4	49	--	0	1	0	0
		Double	7	296.3	840	--	0	0	779	67
	AddSubtract	Single	5	319.39	483	--	0	0	408	37
		Double	7	289.77	1,106	--	0	0	1,006	156
	Cube Root	Single	10	432.9	126	--	5	2	121	0
		Double	24	282.09	594	--	11	10	1,155	29
	Divide	Single	16	347.34	394	--	3	4	561	66
		Double	30	258.26	1,208	--	20	15	2,175	136
continued..										

Family	Function	Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
									Primary	Secondary
Arria 10 (10AX11 5H4F34I 3SP)	Exp base 10	Single	14	271.37	502	--	3	2	432	40
		Double	29	242.42	2,185	--	0	10	1,683	90
	Exp base 2	Single	7	317.86	370	--	0	2	124	9
		Double	22	251.45	906	--	0	10	1,172	47
	Exp base e	Single	26	365.36	298	--	3	6	137	11
		Double	28	260.42	2,156	--	0	10	1,724	93
	Reciprocal	Single	12	278.94	225	--	3	3	172	3
		Double	27	260.89	824	--	9	14	1,448	100
	Reciprocal Square Root	Single	8	418.94	117	--	3	2	130	1
		Double	22	243.43	523	--	8	9	950	37
	LDExp	Single	0	--	68	--	0	0	0	0
		Double	0	--	99	--	0	0	66	0
	Log base 10	Single	15	293.69	364	--	3	3	441	42
		Double	28	272.03	1,158	--	20	11	2,095	214
	Log(1+x)	Single	18	301.3	747	--	3	3	882	79
		Double	32	251.95	2,018	--	20	13	3,019	248
Arria 10 (10AX11 5H4F34I 3SP)	Log base 2	Single	14	275.79	316	--	3	3	402	3
		Double	32	271.96	1,173	--	8	16	2,372	132
	Log base e	Single	29	378.07	297	--	3	9	315	6
		Double	29	256.54	1,219	--	20	13	2,338	152
	Multiply	Single	3	288.4	49	--	0	1	0	0
		Double	5	288.35	312	--	0	4	236	26
	Power	Single	40	262.12	1,335	--	8	14	1,523	127
		Double	73	237.7	3,957	--	13	37	5,362	305
	Square Root	Single	8	432.9	124	--	3	2	118	8
		Double	22	249.25	539	--	8	9	1,000	34
	Subtract	Single	4	296.9	49	--	0	1	0	0
		Double	7	296.82	842	--	0	0	783	76

Table 94. Trigonometry

Family	Function	Precision	Scale By Pi	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
Arria V (5AGX)	Arccos	Single	0	35	217.7 7	768	9	--	8	1,289	94
		Single	1	39	216.4 5	819	9	--	9	1,383	92
continued...											

Family	Function	Precision	Scale By Pi	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
FB3H4 F40C5)		Double	0	76	185.7	2,917	27	--	37	6,489	230
		Double	1	83	184.2	3,120	27	--	40	6,899	198
	Arcsin	Single	0	29	215.8	652	9	--	8	1,069	93
		Single	1	34	222.32	747	9	--	9	1,178	80
		Double	0	66	185.32	2,762	29	--	41	6,365	171
		Double	1	72	184.16	2,963	29	--	44	6,696	200
		Single	0	27	232.29	603	7	--	6	937	77
		Single	1	31	230.73	664	7	--	7	1,034	89
	Arctan	Double	0	65	185.7	2,047	23	--	31	4,535	164
		Double	1	71	185.6	2,229	23	--	34	4,854	174
	Arctan2	Single	0	43	230.2	1,013	11	--	9	1,719	128
		Single	1	43	230.2	1,013	11	--	9	1,719	128
		Double	0	92	184.2	3,195	44	--	43	6,822	285
		Double	1	92	184.2	3,195	44	--	43	6,822	285
Arria V (5AGX FB3H4 F40C5)	Cos	Single	0	25	205.3	768	5	--	6	1,563	120
		Single	1	12	242.13	490	0	--	3	475	36
		Double	0	45	184.2	2,879	34	--	33	5,973	244
		Double	1	29	185.87	1,719	0	--	13	2,499	92
	Sin	Single	0	26	223.51	964	5	--	6	1,439	110
		Single	1	12	240.56	585	0	--	3	563	66
		Double	0	46	184.16	3,019	36	--	33	6,308	249
		Double	1	29	185.77	1,748	0	--	14	2,699	92
	Tan	Single	0	38	221.78	1,368	12	--	12	2,625	163
continued...											

Family	Function	Precision	Scale By Pi	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
Cyclone V (5CGX FC7D6 F31C7)		Single	1	25	231.4 3	1,29 7	4	--	10	1,512	140
		Double	0	68	185.5 6	5,21 1	56	--	65	10,670	530
		Double	1	52	184.1 6	3,87 4	26	--	43	6,896	238
	Arccos	Single	0	42	217.2	857	9	--	8	1,701	107
		Single	1	47	196.2 7	943	9	--	9	1,887	104
		Double	0	113	196.5	3,95 7	31	--	37	9,739	343
		Double	1	123	210.1 7	4,21 8	31	--	40	10,353	333
	Arcsin	Single	0	35	222.6 2	757	9	--	8	1,464	65
		Single	1	40	215.1 9	844	9	--	9	1,627	111
		Double	0	101	201.6 5	3,81 6	31	--	41	9,709	334
		Double	1	112	197.7 1	4,04 6	31	--	44	10,383	260
	Arctan	Single	0	33	227.5 8	706	7	--	6	1,266	92
		Single	1	38	206.3 1	787	7	--	7	1,434	90
		Double	0	98	188.7 9	2,92 0	24	--	31	7,154	297
		Double	1	109	180.4 1	3,19 6	24	--	34	7,875	297
Cyclone V (5CGX FC7D6 F31C7)	Arctan2	Single	0	51	206.1 4	1,15 3	11	--	9	2,013	149
		Single	1	51	206.1 4	1,15 3	11	--	9	2,013	149
		Double	0	144	191.1 7	4,58 9	46	--	43	10,740	417
		Double	1	144	191.1 7	4,58 9	46	--	43	10,740	417
	Cos	Single	0	32	174.6 7	959	5	--	6	2,258	110
		Single	1	15	212.9 9	517	0	--	3	702	25
		Double	0	75	182.7 5	3,75 1	34	--	33	9,177	352
		Double	1	50	212.6 8	1,98 5	0	--	13	3,914	169

continued...

Family	Function	Precision	Scale By Pi	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
	Sin	Single	0	33	191.6 1	1,08 6	5	--	6	2,394	132
		Single	1	14	207.8 1	579	0	--	3	783	39
		Double	0	75	196.3 9	3,78 7	38	--	33	9,545	284
		Double	1	49	206.5 3	2,16 5	0	--	14	4,336	177
	Tan	Single	0	46	185.7 4	1,65 5	12	--	12	3,738	200
		Single	1	29	205.4 7	1,28 3	4	--	10	2,142	102
		Double	0	112	194.7	7,05 2	58	--	65	16,793	607
		Double	1	89	197.2 4	5,32 7	26	--	43	11,741	376
Strati x V (5SGX EA7K2 F40C2)	Arccos	Single	0	23	291.4 6	753	--	9	8	801	34
		Single	1	27	288.4 3	823	--	9	9	891	27
		Double	0	53	247.4 6	2,38 0	--	27	37	4,435	145
		Double	1	58	233.1 5	2,57 0	--	27	40	4,717	121
	Arcsin	Single	0	20	290.6 1	598	--	9	8	698	19
		Single	1	23	294.9 9	678	--	9	9	800	21
		Double	0	47	237.2 5	2,23 5	--	27	40	4,407	89
		Double	1	52	240.3 3	2,41 1	--	27	43	4,621	134
	Arctan	Single	0	20	293.6	544	--	6	6	646	53
		Single	1	23	290.7	620	--	6	7	715	50
		Double	0	47	241.7 2	1,83 7	--	18	30	3,424	145
		Double	1	52	247.3 4	2,00 2	--	18	33	3,654	126
	Arctan2	Single	0	31	288.3 5	890	--	9	9	1,277	71
		Single	1	31	288.3 5	890	--	9	9	1,277	71
		Double	0	69	239.2 3	2,98 3	--	29	42	5,530	212

continued...

Family	Function	Precision	Scale By Pi	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
		Double	1	69	239.2 3	2,98 3	--	29	42	5,530	212
Strati x V (5SGX EA7K2 F40C2)	Cos	Single	0	17	267.0 2	711	--	5	6	890	48
		Single	1	8	364.8 3	452	--	0	3	368	20
		Double	0	33	242.2 5	2,52 9	--	17	31	4,510	187
		Double	1	21	275.4 8	1,60 8	--	0	13	1,799	54
	Sin	Single	0	18	309.6 9	856	--	5	6	917	74
		Single	1	8	317.4 6	538	--	0	3	382	10
		Double	0	34	257.6 7	2,71 4	--	19	31	4,766	229
		Double	1	22	260.8 9	1,86 4	--	0	14	1,898	104
	Tan	Single	0	27	272.2 6	1,31 2	--	11	12	1,675	117
		Single	1	17	295.6 8	1,16 4	--	3	10	1,175	65
		Double	0	52	268.3 8	4,61 2	--	30	60	8,152	264
		Double	1	41	260.8 9	3,88 6	--	13	43	5,294	193
Arria 10 (10AX 115H4 F34I3 SP)	Arccos	Single	0	28	270.4 2	703	--	9	8	656	29
		Single	1	31	261.2 3	705	--	9	9	624	19
		Double	0	63	257.8	2,62 6	--	27	37	4,917	241
		Double	1	69	255.6 2	2,81 3	--	27	40	5,127	268
	Arcsin	Single	0	25	249.6 3	665	--	9	8	659	18
		Single	1	28	254.1 9	673	--	9	9	649	30
		Double	0	57	255.6 2	2,44 0	--	29	40	4,750	213
		Double	1	62	251.5 1	2,59 6	--	29	43	4,985	190
	Arctan	Single	0	26	271.3	600	--	6	6	578	32
		Single	1	29	274.2	594	--	6	7	583	22

continued...

Family	Function	Precision	Scale By Pi	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
	Arctan2	Double	0	57	254.1 3	1,86 6	--	22	30	3,654	171
		Double	1	63	258.3 3	2,04 3	--	22	33	3,726	253
		Single	0	40	248.1 4	1,00 2	--	9	9	1,258	85
		Single	1	40	248.1 4	1,00 2	--	9	9	1,258	85
		Double	0	84	255.1	3,02 5	--	33	42	5,675	328
		Double	1	84	255.1	3,02 5	--	33	42	5,675	328
Arria 10 (10AX 115H4 F34I3 SP)	Cos	Single	0	21	336.9 3	786	--	5	6	979	154
		Single	1	11	310.3 7	512	--	0	3	297	22
		Double	0	39	263.9 2	2,70 2	--	17	33	3,697	375
		Double	1	29	242.1 9	1,69 8	--	0	13	2,030	62
	Sin	Single	0	22	311.3 3	876	--	5	6	1,003	116
		Single	1	11	279.0 2	585	--	0	3	330	19
		Double	0	41	265.2 5	2,79 1	--	19	33	3,902	334
		Double	1	29	259.6 1	1,91 8	--	0	14	1,943	72
	Tan	Single	0	34	265.6	1,35 9	--	11	12	1,756	155
		Single	1	23	265.8 9	1,26 8	--	3	10	1,065	94
		Double	0	64	248.1 4	5,10 7	--	30	65	7,578	458
		Double	1	53	251.7	4,00 2	--	17	43	5,619	343

Table 95. FPFXP

Family	Input Precision	Output Width	Output Fraction	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
Arria V (5AGXF B3H4F4 0C5)	Single	32	0	2	277.93	168	0	--	0	75	1
		32	16	2	266.1	169	0	--	0	75	0
		32	32	2	277.93	168	0	--	0	75	1
continued...											

Family	Input Precision	Output Width	Output Fraction	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
		64	0	3	226.4	291	0	--	0	172	0
		64	16	3	226.4	291	0	--	0	172	0
		64	32	3	226.4	291	0	--	0	172	0
	Double	32	0	3	332.12	197	0	--	0	115	0
		32	16	3	344.12	197	0	--	0	115	0
		32	32	3	332.12	197	0	--	0	115	0
		64	0	3	256.28	326	0	--	0	205	4
		64	16	3	256.28	326	0	--	0	205	4
		64	32	3	256.28	326	0	--	0	205	4
		64	32	3	256.28	326	0	--	0	205	4
Cyclone V (5CGXF C7D6F3 1C7)	Single	32	0	3	245.04	171	0	--	0	110	0
		32	16	3	245.04	171	0	--	0	110	0
		32	32	3	245.04	171	0	--	0	110	0
		64	0	4	190.62	244	0	--	0	269	0
		64	16	4	190.62	244	0	--	0	269	0
		64	32	4	190.62	244	0	--	0	269	0
	Double	32	0	4	291.63	209	0	--	0	160	1
		32	16	4	302.94	209	0	--	0	160	1
		32	32	4	291.63	209	0	--	0	160	1
		64	0	5	207.25	329	0	--	0	347	2
		64	16	5	207.25	329	0	--	0	347	2
		64	32	5	207.25	329	0	--	0	347	2
Stratix V (5SGXE A7K2F4 0C2)	Single	32	0	0	717.36	168	--	0	0	38	0
		32	16	0	717.36	168	--	0	0	38	0
		32	32	0	717.36	168	--	0	0	38	0
		64	0	0	717.36	304	--	0	0	70	0
		64	16	0	717.36	304	--	0	0	70	0
		64	32	0	717.36	304	--	0	0	70	0
	Double	32	0	0	717.36	204	--	0	0	38	0
		32	16	0	717.36	204	--	0	0	38	0
		32	32	0	717.36	204	--	0	0	38	0
		64	0	2	456	329	--	0	0	134	1
		64	16	2	456	329	--	0	0	134	1
		64	32	2	456	329	--	0	0	134	1
Arria 10 (10AX1 15H4F3)	Single	32	0	0	--	168	--	0	0	38	0

continued...

Family	Input Precision	Output Width	Output Fraction	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
4I3SP)		32	16	0	--	168	--	0	0	38	0
		32	32	0	--	168	--	0	0	38	0
		64	0	0	--	304	--	0	0	70	0
		64	16	0	--	304	--	0	0	70	0
		64	32	0	--	304	--	0	0	70	0
	Double	32	0	0	--	203	--	0	0	38	0
		32	16	0	--	203	--	0	0	38	0
		32	32	0	--	203	--	0	0	38	0
		64	0	2	407.33	328	--	0	0	134	0
		64	16	2	407.33	328	--	0	0	134	0
		64	32	2	407.33	328	--	0	0	134	0

Table 96. FXFPF

Family	Input Width	Input Fraction	Output Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
Arria V (5AGXF B3H4F4 0C5)	32	0	Single	6	283.61	154	0	--	0	195	14
	32	0	Double	5	328.19	165	0	--	0	180	17
	32	16	Single	6	283.61	154	0	--	0	195	14
	32	16	Double	5	328.19	165	0	--	0	180	17
	32	32	Single	6	293	152	0	--	0	193	13
	32	32	Double	5	336.59	159	0	--	0	180	16
	64	0	Single	7	282.01	217	0	--	0	297	16
	64	0	Double	7	256.48	330	0	--	0	451	18
	64	16	Single	7	282.01	217	0	--	0	297	16
	64	16	Double	7	256.48	330	0	--	0	451	18
	64	32	Single	7	282.01	217	0	--	0	297	16
	64	32	Double	7	256.48	330	0	--	0	451	18
Cyclone V (5CGXF C7D6F3 1C7)	32	0	Single	8	230.04	168	0	--	0	264	21
	32	0	Double	7	292.74	180	0	--	0	258	23
	32	16	Single	8	230.04	168	0	--	0	264	21
	32	16	Double	7	292.74	180	0	--	0	258	23
	32	32	Single	8	237.14	166	0	--	0	262	20
	32	32	Double	7	268.6	179	0	--	0	258	29
	64	0	Single	9	248.51	219	0	--	0	391	18

continued...

Family	Input Width	Input Fraction	Output Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
										Primary	Secondary
	64	0	Double	10	176.87	338	0	--	0	648	18
	64	16	Single	9	248.51	219	0	--	0	391	18
	64	16	Double	10	176.87	338	0	--	0	648	18
	64	32	Single	9	248.51	219	0	--	0	391	18
	64	32	Double	10	176.87	338	0	--	0	648	18
Stratix V (5SGXE A7K2F40C2)	32	0	Single	3	579.71	148	--	0	0	97	1
	32	0	Double	2	547.95	161	--	0	0	72	1
	32	16	Single	3	550.66	148	--	0	0	97	1
	32	16	Double	2	536.19	160	--	0	0	72	0
	32	32	Single	3	558.66	145	--	0	0	96	1
	32	32	Double	2	496.28	154	--	0	0	72	1
	64	0	Single	3	454.55	194	--	0	0	125	0
	64	0	Double	3	434.22	304	--	0	0	194	3
	64	16	Single	3	454.55	194	--	0	0	125	0
	64	16	Double	3	434.22	304	--	0	0	194	3
	64	32	Single	3	454.55	194	--	0	0	125	0
	64	32	Double	3	434.22	304	--	0	0	194	3
Arria 10 (10AX115H4F34I3SP)	32	0	Single	3	464.9	147	--	0	0	97	0
	32	0	Double	2	458.93	161	--	0	0	72	0
	32	16	Single	3	464.9	147	--	0	0	97	0
	32	16	Double	2	432.15	160	--	0	0	72	0
	32	32	Single	3	451.67	145	--	0	0	96	0
	32	32	Double	2	419.99	154	--	0	0	72	0
	64	0	Single	3	417.54	193	--	0	0	124	3
	64	0	Double	3	407.33	305	--	0	0	193	3
	64	16	Single	3	417.54	193	--	0	0	124	3
	64	16	Double	3	407.33	305	--	0	0	193	3
	64	32	Single	3	417.54	193	--	0	0	124	3
	64	32	Double	3	407.33	305	--	0	0	193	3


Family	Input Precision	Output Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
									Primary	Secondary
Arria V (5AGXFB3H4F40C5)	Single	Double	2	371.61	93	0	—	0	71	0
	Double	Single	2	370.64	127	0	—	0	74	1

continued...

Family	Input Precision	Output Precision	Latency	f _{MAX}	ALMs	M10K	M20K	DSP Blocks	Logic Registers	
									Primary	Secondary
Cyclone V (5CGXFC7D6F31C7)	Single	Double	2	346.14	93	0	—	0	72	1
	Double	Single	3	349.9	126	0	—	0	111	2
Stratix V (5SGXEA7K2F40C2)	Single	Double	0	—	76	—	0	0	0	0
	Double	Single	0	717.36	126	—	0	0	34	0
Arria 10 (10AX115H4F34I3SP)	Single	Double	0	—	75	—	0	0	0	0
	Double	Single	0	—	126	—	0	0	34	0

16.6. FP_FUNCTIONS Intel FPGA IP Signals

Figure 41. FP_FUNCTIONS Intel FPGA IP Signals



- 1) The floating point and fixed point data widths determine the port width of this port.
- 2) This port is not relevant for convert and square root functions.

Table 97. FP_FUNCTIONS Intel FPGA IP Input Signals

Port Name	Required	Description
clk	Yes	All input signals must be synchronous to this clock.
areset	Yes	Asynchronous active-high reset. Deassert this signal synchronously to the input clock to avoid metastability issues.
en	No	Optional port. Allow calculation to take place when asserted. When deasserted, no operation will take place and the outputs are unchanged.
a	Yes	Data input signal.

continued...

Port Name	Required	Description
b	Yes	Data input signal (where applicable).
s	Yes	Select port for Add/Sub function.
c	Yes	Data port for integer exponent port for LDExp function.

Table 98. FP_Functions Intel FPGA IP Output Signals

Port Name	Required	Description
q	Yes	Data output signal.

16.7. FP_FUNCTIONS Intel FPGA IP Parameters

These tables list the FP_FUNCTIONS Intel FPGA IP parameters.

Table 99. FP_FUNCTIONS Intel FPGA IP Parameters: Functionality Tab

Parameter	Values	Descriptions
Function		
Family	- All - Arithmetic - Comparisons - Conversions - Exp/Log/Pow - Roots - Trigonometry	Allows you to chose which functions to be displayed in the Function Name Parameter list. The default value is All.
Name	- Add - Subtract - Add/Sub - Multiply - Divide - Reciprocal - Absolute - Scalar Product - Multiply Accumulate - Accumulate - Multiply Add - Complex Multiply - Sin - Cos - Tan - Arcsin - Arccos - Arctan - Arctan2 - Exponent - Exponent base 2 - Exponent base 10 - Log - Log₂ - Log₁₀ - Log_(1+x)	Allows you to choose your desired function. <i>Note:</i> This parameter only displays the options you have selected from the Family Parameter.

continued...

Parameter	Values	Descriptions
	- Power - Square Root - Reciprocal Square Root - Cube Root 3D Hypotenuse - Minimum - Maximum - Less Than - Less Than or Equal - Equal - Not Equal - Greater Than - Greater Than or Equal - Fixed to Floating Point - Floating Point to Fixed - Floating to Floating Point	
Use Select Signal	—	Select this option to generate a Select signal. Use the Select signal to choose the option to use both addition and subtraction functions or either one of the functions.
Represent angle as multiple of Pi	—	Select this option to represent angles as multiple of Pi. <i>Note:</i> Not available for Arctan2 function.
Inputs are within range -2pi to +2pi	—	Select this option to disable range reduction. <i>Note:</i> Only available for Sin and Cos function.
Floating Point Data		
Format	Single Double Custom	Allows you to choose the floating point format of the data values. The default value is single.
Exponent	5 to 11	Allows you to specify the width of the exponent. This parameter is only available when the Format parameter is set to custom. The default value is 8.
Mantissa	10 to 52	Allows you to specify the width of the mantissa. This parameter is only available when the Format parameter is set to custom. The default value is 23.
Input Vector Dimension	Integer	Provide the desired the number of inputs to compute the vector dimension.
Input Format	Single Double Custom	Allows you to choose the floating point format of the input data values. The default value is single. <i>Note:</i> Only available for Floating to Floating Point function.
Input Exponent	Integer	Allows you to specify the width of the input exponent. This parameter is only available when the Format parameter is set to custom. The default value is 8. <i>Note:</i> Only available for Floating to Floating Point function.
Input Mantissa	Integer	Allows you to specify the width of the mantissa. This parameter is only available when the Format parameter is set to custom. The default value is 23. <i>Note:</i> Only available for Floating to Floating Point function.
<i>continued...</i>		

Parameter	Values	Descriptions
Output Format	Single Double Custom	Allows you to choose the floating point format of the output data values. The default value is single.
Output Exponent	Integer	Allows you to specify the width of the output exponent. This parameter is only available when the Format parameter is set to custom. The default value is 8. <i>Note:</i> Only available for Floating to Floating Point function.
Output Mantissa	Integer	Allows you to specify the width of the mantissa. This parameter is only available when the Format parameter is set to custom. The default value is 23. <i>Note:</i> Only available for Floating to Floating Point function.
Fixed Point Data		
Width	16 to 128	The bit width of the fixed point data port. This parameter is only available when the Name parameter is set to Fixed to Floating Point . The default value is 32.
Fraction	-128 to 128	The bit width of the fraction. This parameter is only available when the Name parameter is set to Fixed to Floating Point .
Sign	Signed , Unsigned	Choose if the fixed point data is signed or unsigned. This parameter is only available when the Name parameter is set to Convert. The default value is signed.
Rounding		
Mode	• nearest with tie breaking to even	The rounding mode.
Relax rounding to round up or down to reduce resource usage	—	Choose if the nearest rounding mode should be relaxed to faithful rounding, where the result may be rounded up or down, to reduce resource usage. Only available for arithmetic functions
Ports		
Generate Enable Port	—	Choose if the FP_FUNCTIONS Intel FPGA IP core should have an enable signal.

Table 100. FP_FUNCTIONS Intel FPGA IP Parameters: Performance Tab

Parameter	Values	Descriptions
Target		
Goal	- Frequency - Latency - Combined - Manually Specify DSP Registers	If the Goal is the frequency, then the Target is the desired frequency in MHz. This, together with the target device family, determines the amount of pipelining. If the Goal is Combined then two Targets are displayed, one is the desired frequency in MHz, one is the target latency in cycles. When you set the Goal parameter to frequency, the default value is 200 MHz. When you set the Goal parameter to latency, the default value is 2. If the Goal is Latency, then the Target is the desired latency. The report generates the achievable latency if it can't meet target latency.

continued...

Parameter	Values	Descriptions
		If the Goal is set to Manually Specify DSP Registers , you can manually select the register and function subblocks within the DSP IP core.
Target	Any Positive Integer	Specify your target frequency and latency.
Report		
Latency on Arria 10 is <x> cycles	—	This report shows the latency of the function.
Resource Estimates: • Multiplies • LUTs • Memory Bits • Memory Blocks	—	This report shows the number of multipliers, LUTs, memory bits, and memory blocks utilized by the IP core.
Check Performance	—	Click this to check if the design can achieve the target latency. <i>Note:</i> Only available when Goal is set to Latency .

17. Floating-Point IP Cores User Guide Document Archives

IP versions are the same as the Intel Quartus Prime Design Suite software versions up to v19.1. From Intel Quartus Prime Design Suite software version 19.2 or later, IP cores have a new IP versioning scheme.

If an IP core version is not listed, the user guide for the previous IP core version applies.

Intel Quartus Prime Version	IP Core Version	User Guide
16.1	16.1	Floating-Point IP Cores User Guide
15.0	15.0	Floating-Point IP Cores User Guide

18. Document Revision History for the Floating-Point IP Cores User Guide

Document Version	Intel Quartus Prime Version	Changes
2021.09.13	20.1	Updated <i>List of Functions Supported by FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP</i>
2021.03.12	20.1	Updated the supported device for FP_FUNCTIONS Intel FPGA IP to include Intel Arria 10 devices in Table: <i>IP Cores Available in Intel Quartus Prime Standard Edition Software</i> .
2020.06.22	20.1	- Added <i>IP Cores Available in Intel Quartus Prime Standard Edition Software</i> and <i>IP Cores Available in Intel Quartus Prime Pro Edition Software</i> tables. - Removed information about the ALTERA_FP_MATRIX_INV and ALTERA_FP_MATRIX_MULT_IP cores. These cores are no longer supported in versions 20.1 and later. - Rebranded the entire document as per Intel standards. - Renamed ALTERA_FP_ACC_CUSTOM IP Core to FP_ACC_CUSTOM Intel FPGA IP or Floating Point Custom Accumulator Intel FPGA IP. - Renamed ALTERA_FP_FUNCTIONS IP Core to FP_FUNCTIONS Intel FPGA IP or Floating Point Functions Intel FPGA IP. - Added <i>Floating Point Custom Accumulator Intel FPGA IP Release Information</i> and <i>Floating Point Functions Intel FPGA IP Release Information</i> tables. - Updated the <i>VHDL LIBRARY-USE Declaration</i> topic to correct the VHDL Library declaration example from USE altera_mf_altera_mf_components.all; to USE altera_mf.altera_mf_components.all;.

Date	Document Version	Changes Made
December 2016	2016.12.09	- Added simple description about the IP core on each introduction page. - Added descriptions for each function supported by ALTERA_FP_FUNCTIONS IP core. - Clarified that ALTERA_FP_FUNCTIONS IP core replaces all ALTFP IP cores in Arria 10 devices. - Added new parameters in ALTERA_FP_FUNCTIONS Parameter: Functionality Tab table. - Clarified the functionality of en signal for ALTERA_FP_FUNCTIONS.
July 2015	2015.07.30	- Updated link to Floating-Point IP Cores Design Examples. - Updated Memory Blocks numbers in ALTERA_FP_MATRIX_MULT Resource Utilization and Performance for the Arria 10 and Stratix V Devices table. - Added notes on default settings for WIDTH_INT and WIDTH_DATA.
December 2014	2014.12.19	- Remove all references to the complex mode in ALTFP_MATRIX_MULTIPLY. - Updated ALTERA_FP_MATRIX_MULT and ALTERA_FP_FUNCTIONS sections.
continued...		

Intel Corporation. All rights reserved. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

*Other names and brands may be claimed as the property of others.

Date	Document Version	Changes Made
November 2013	7.0	- Added "ALTERA_FP_FUNCTIONS" on page 3-1. - Added "ALTERA_FP_ACC_CUSTOM" on page 2-1. - Updated Table 1-1 on page 1-1 to list ALTERA_FP_FUNCTIONS and ALTERA_FP_ACC_CUSTOM. - Updated the "ALTFP_MATRIX_INV" on page 17-1 section to include 4 x 4 and 6 x 6 dimensions. - Updated "Rounding" on page 1-4 to clarify that the code for round-to-nearest-even mode is T0_NEAREST. - Removed Design Example section for "ALTFP_MATRIX_MULT" on page 18-1. - Removed device family support for HardCopy III, HardCopy IV, Stratix II, and Stratix II GX devices from "Device Family Support" on page 1-2.
November 2011	6.0	Updated "General Features" on page 1-2.
May 2011	5.0	Added "ALTFP_ATAN" on page 12-1.
January 2011	4.0	Added "ALTFP_SINCOS" on page 13-1.
July 2010	3.0	- Updated architecture information for the following sections: ALTFP_MATRIX_MULT ALTFP_MATRIX_INV. - Added specification information in all sections.
November 2009	2.0	- Updated resource utilization information for the following sections: ALTFP_ADD_SUB ALTFP_DIV ALTFP_MULT ALTFP_SQRT ALTFP_EXP ALTFP_INV ALTFP_INV_SQRT ALTFP_LOG ALTFP_COMPARE ALTFP_CONVERT ALTFP_MATRIX_MULT - Added the ALTFP_MATRIX_INV section. - Updated the Ports and Parameters section for all floating-point megafunctions.
March 2009	1.0	Initial release.