Designer Resource. Accelerate. Simplify.

Issue 3 Summer 2008
Document Number: BR8BITBYNDBITS Rev 2

Beyond Bits

Find the power in going green

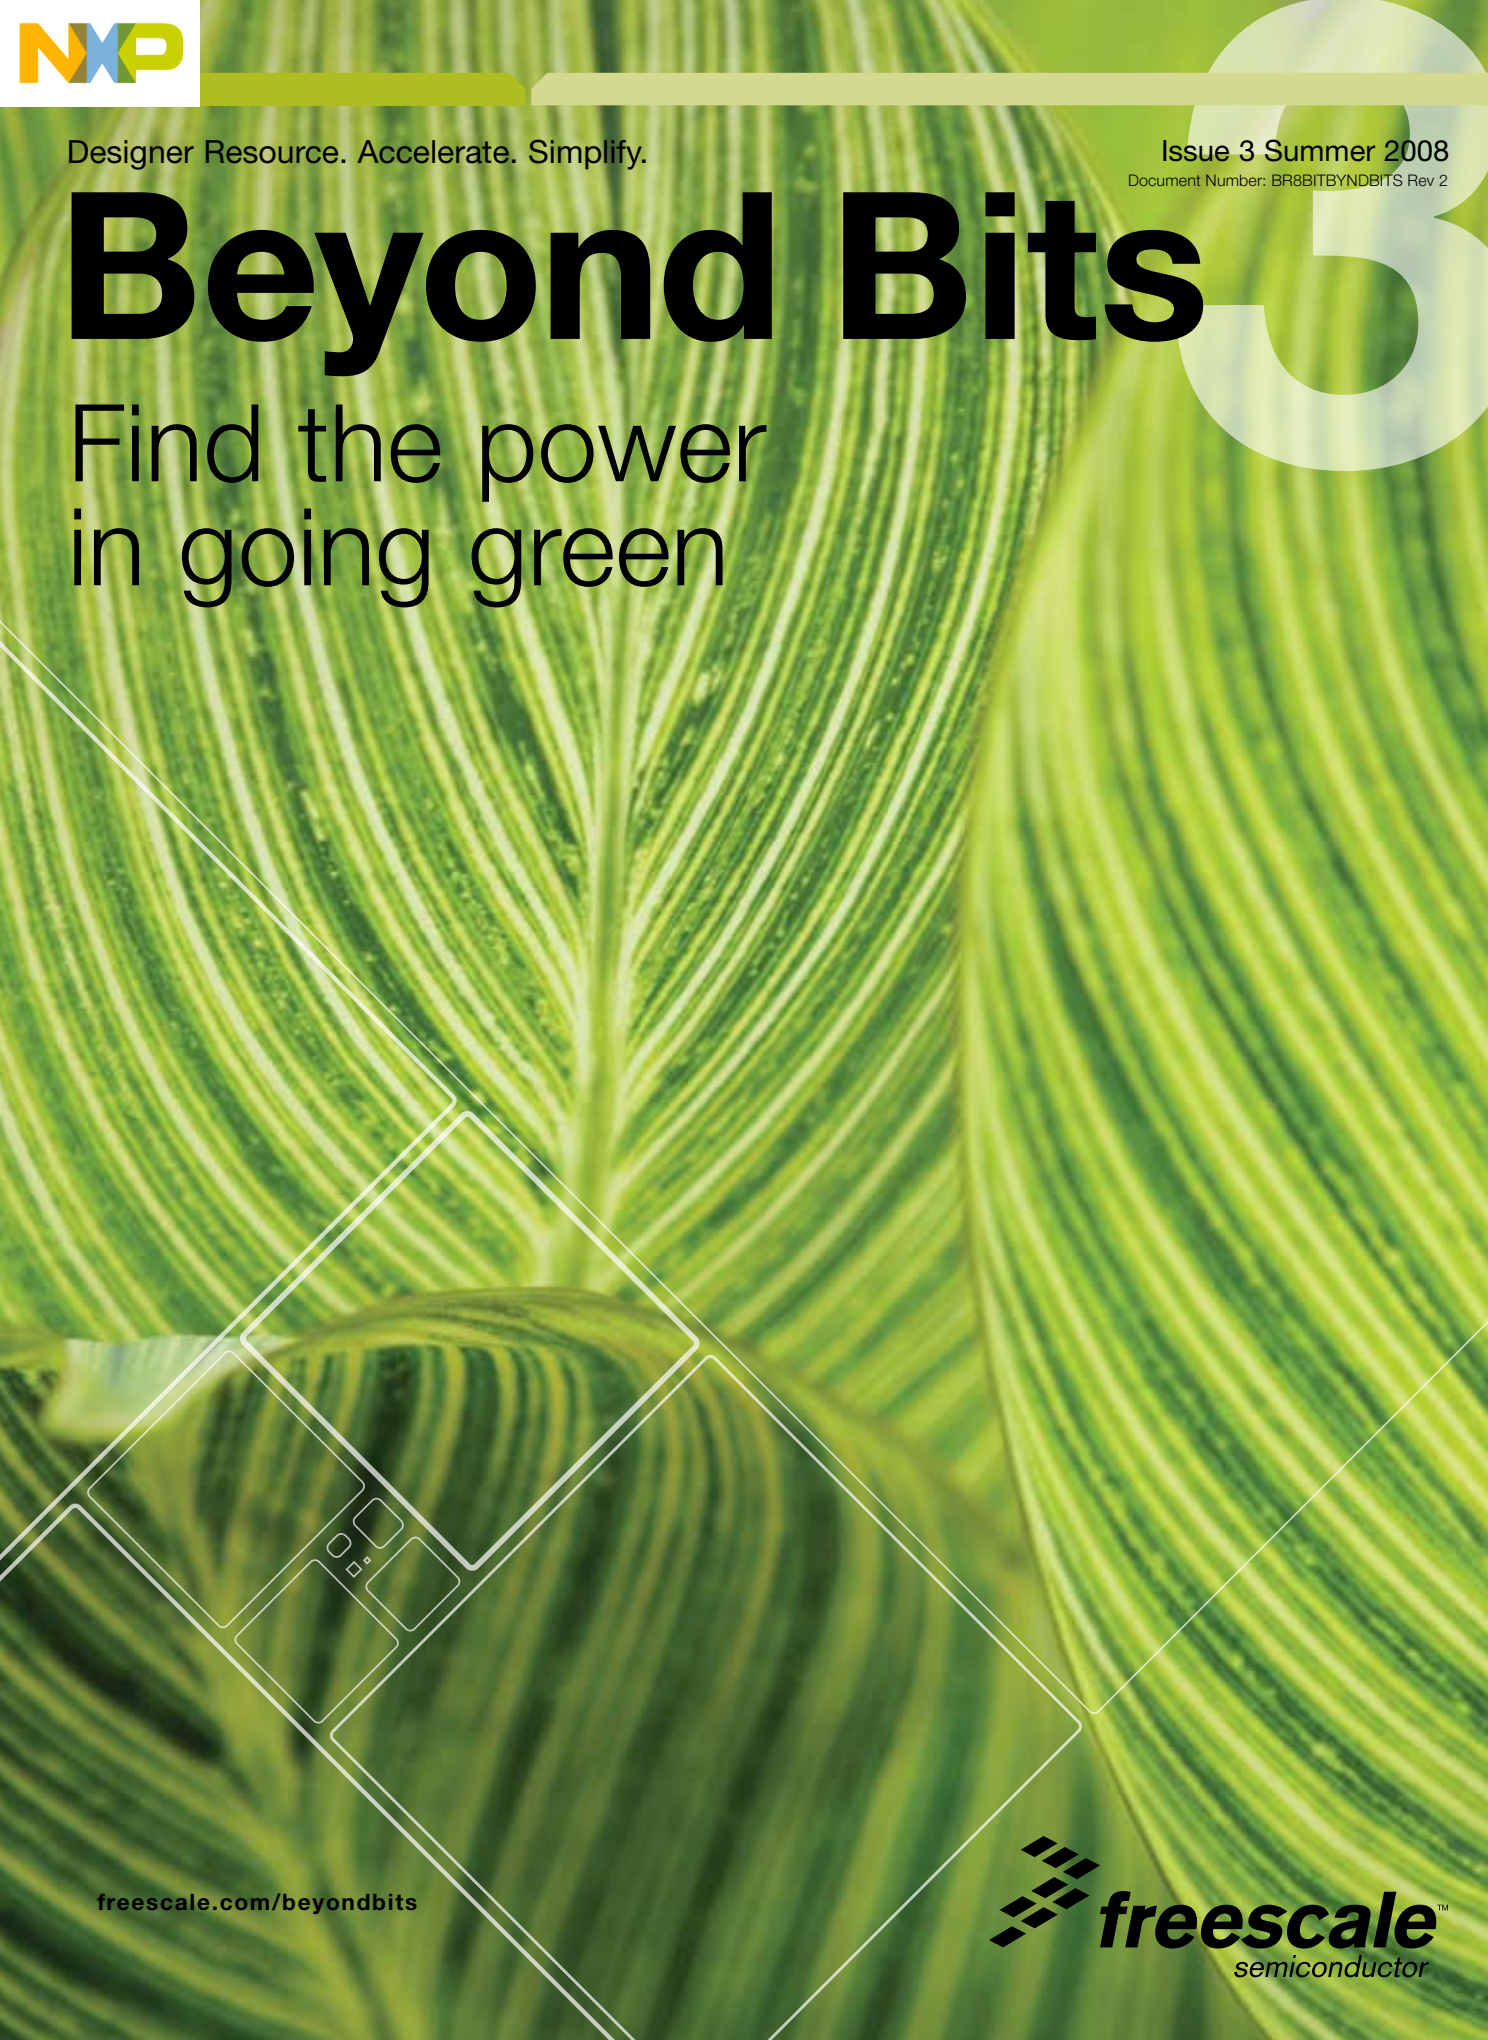

freescale.com/beyondbits

freescale semiconductor

It's a family that's ready for any design scenario.

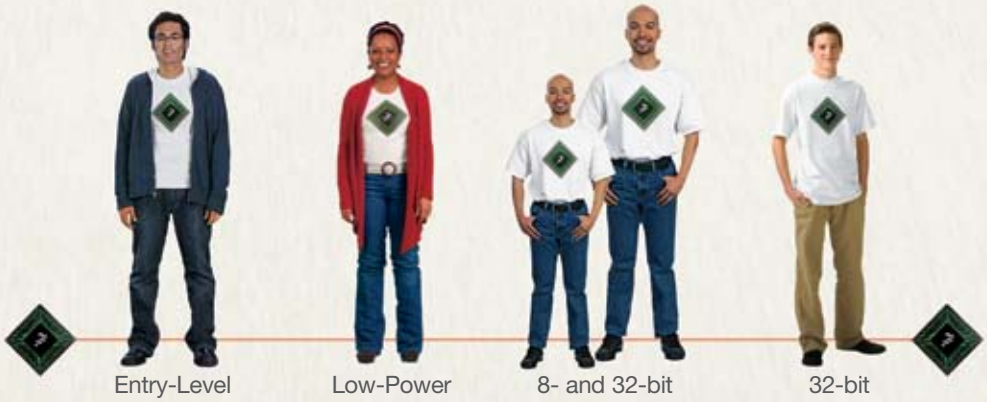

The Controller Continuum Family

Meet the family of MCUs that can help you solve any design scenario, no matter what the challenges. You'll find true ease-of-use in our entry-level 8-bit solutions. Industry-leading power efficiency in our low-power solutions. Unprecedented 8- and 32-bit compatibility in our Flexis™ series of MCUs. And one of the industry's best-connected 32-bit architectures in our ColdFire® series. The entire family shares an extensive library of reference designs, world-class enablement tools like our CodeWarrior® Software and 24/7 support. It's a family you can count on to help you start and finish your next design.

Get details and technical specs at freescale.com/continuumfamily

freescale™ semiconductor

Issue 3

Beyond Bits
Freescale Goes Green

Welcome to this third edition of *Beyond Bits*.

In this issue, we highlight green product development and the broad Freescale product portfolio that enables you to build energy-efficient, environmentally conscious designs.

Energy conservation and carbon awareness are not only becoming business mandates, they are at the heart of the technology industry's social responsibility for earth-friendly living. Consequently, improving energy efficiency is the primary challenge facing the engineering community. This publication features Freescale products in the context of real solutions for low power and energy efficiency.

Freescale technologies, such as 8-bit, 16-bit and 32-bit microcontrollers, capacitive proximity sensors, ZigBee®/IEEE® 802.15.4 transceivers, digital signal controllers and embedded microprocessors enable your new, environmentally friendly designs while continuing to reduce costs and improve time to market. Examples to inspire and guide you are available in this booklet—from advanced motor control to green building technology to fire detection.

This publication has evolved since its inaugural issue—in large part due to the great suggestions we have received from you. Please enjoy *Beyond Bits III*, and continue to provide us your feedback. And, most of all, thank you for considering Freescale as your green partner for energy-efficient design.

Henri Richard
Senior Vice President, Chief Sales and Marketing Officer

We want to hear from you!
Send your design ideas, tips or questions to **freescale.com/beyondbits**.

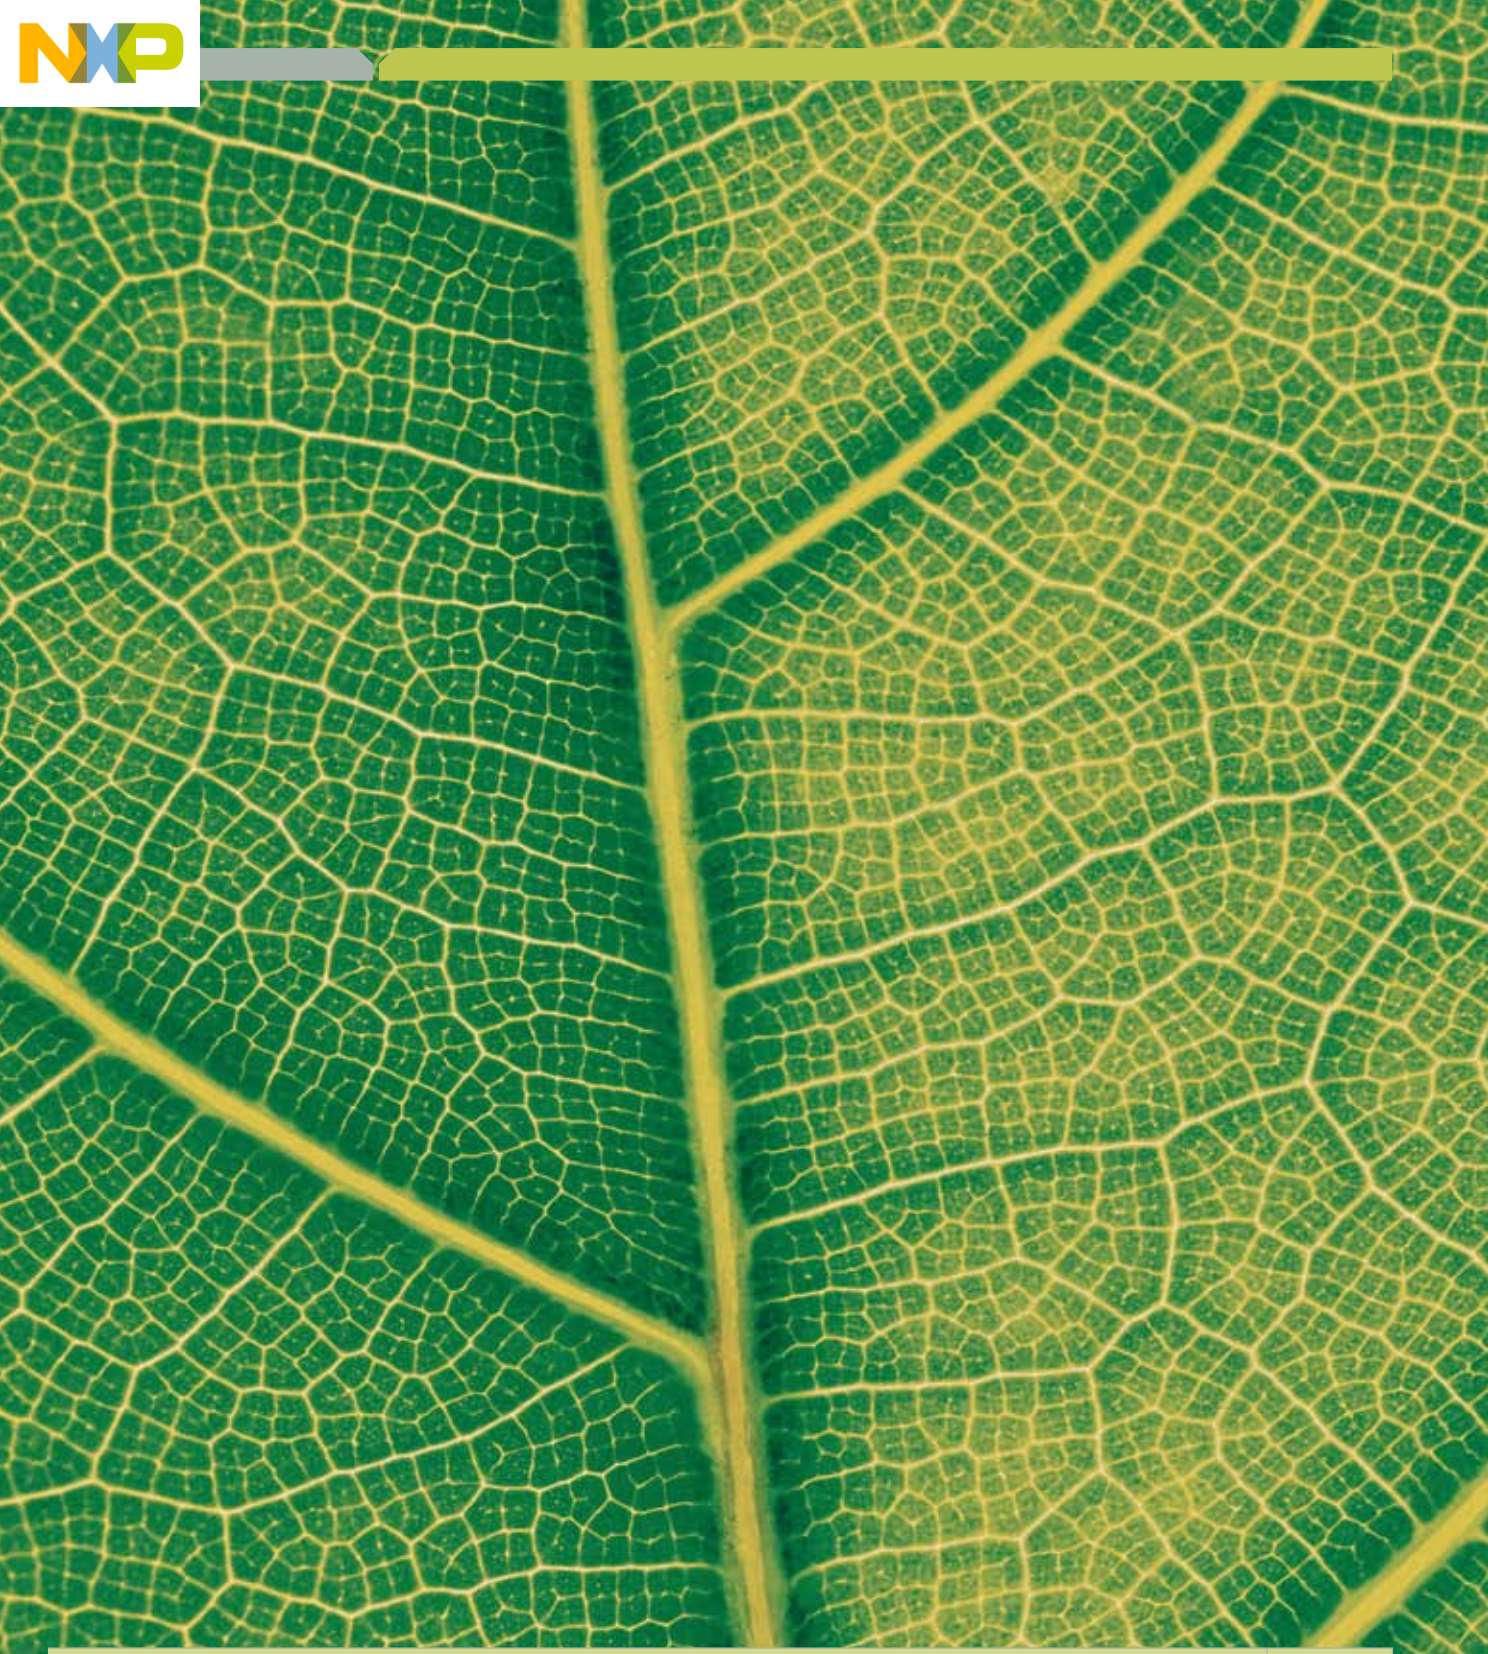

Freescale's Dedication to Sustainability

The following resources were saved in the production of this magazine.

Trees Preserved	Wastewater Flow Saved	Energy Not Consumed	Solid Waste Not Generated	Greenhouse Gases Prevented
75 Fully Grown	32,071 Gallons	54 Million BTU	3,586 Pounds	6,995 Pounds

Calculations based on research by Environmental Defense and other members of the Paper Task Force.

FPO postconsumer waste fiber

FPO © FSC

Beyond Bits
Table of Contents

Articles

Design Challenge

Product Summaries

We want to hear from you!
Send your design ideas, tips or questions to **freescale.com/beyondbits**.

Gonzalo Delgado Huitrón

Preventing and Controlling Forest Fires with ZigBee®/IEEE® 802.15.4

Detecting Forest Fires with ZigBee®/IEEE® 802.15.4

In 2007, more than 800 square miles of California forests were consumed by fire, resulting in tremendous damage to the environment and the loss of a significant portion of the biosphere. Forests are one of Earth's primary air filters, scrubbing the atmosphere of CO^2, a greenhouse gas and contributor to global warming. Fire, whether started by natural causes, cigarettes, bonfires, agricultural techniques or any other reason, is contributing to massive losses of forest land, and it is time to use technology to help Mother Nature preserve those lands.

Today modern technologies, such as infrared scanners and satellite photos, are being used to detect fires, but they have their drawbacks. Satellites are complex and very expensive while infrared scanners have problems with thick vegetation and blocked visibility. The old "watch tower" technique is still being used in many places around the world, but it requires a significant manpower investment that must always be on alert.

Taking advantage of new, low-cost technology, a low-power wireless sensor network can detect a fire nearly the moment it starts. Time is crucial when combating forest fires because the wind can quickly push a small fire into a major conflagration. A sensor network can provide the early warning necessary to allow firefighting resources to converge on the source of the fire as soon as possible.

By definition, a sensor network is a set of nodes that periodically reports data to a coordinator for processing. In a forest, hundreds of low-cost, low-power smoke detecting nodes could be distributed within wireless transmission range of the central, or coordinator, node (see Figure 2).

The coordinator node would gather information from the nearby sensor nodes through a ZigBee compliant wireless network,

Forest Fires per Year[1]

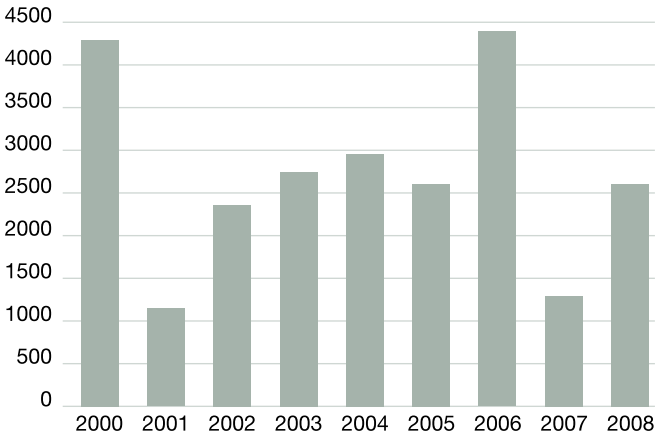

Figure 1

Wireless Sensor Network

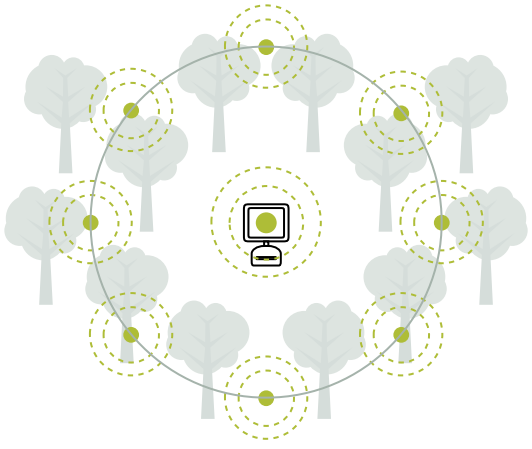

Figure 2

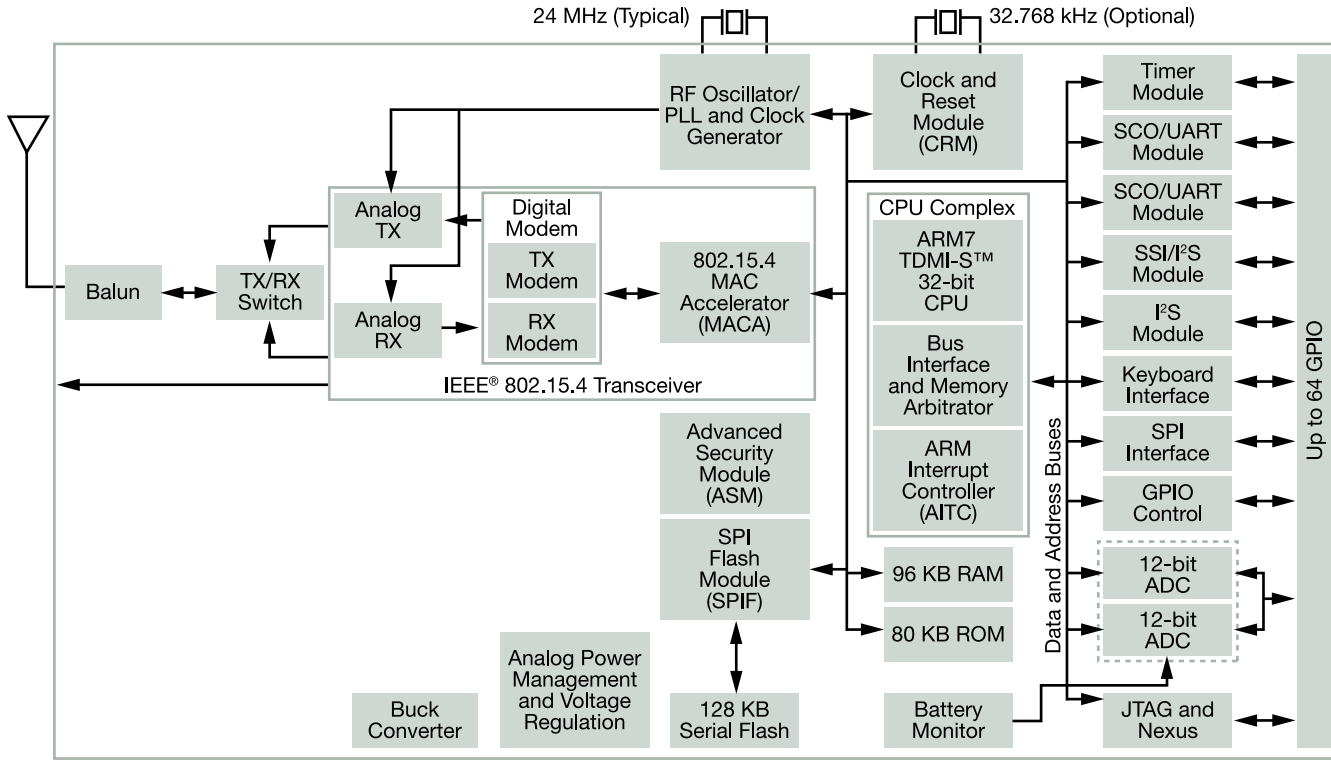

Figure 3

process the data and, if any smoke is detected, send out an alert using a larger wireless network, such as 3G or GPRS. This alert could be sent directly to the firefighting station with not only fire location information but also the time of the detection. Similar solutions have been implemented in Italy with great success.

A feasibility study must be made to determine where each node should be placed for optimal coverage. Since the ZigBee/IEEE 802.15.4 protocol defines a MAC address for each node, this could be correlated to a fixed GPS point to determine the location of the affected zone. One great advantage of using the ZigBee protocol is that a mesh network can be implemented to extend the network's range, using routers instead of coordinators.

Battery life is not a problem, since ZigBee devices can operate up to 20 years between battery replacement. If additional power is required for the external components, such as power amplifiers, smoke sensors or emergency lights, a battery charger could be implemented with solar cells.

With a low power IEEE 802.15.4 system-on-chip solution, such as the MC1322X (see Figure 3), a wireless network could be easily implemented. It contains an ARM7™ core and consumes only about 20mA in transmission and reception.

The MC1322X ZigBee Platform in Package™ (PiP) includes a dual 8-channel, 12-bit analog-to-digital converter that can be interfaced with several sensors to gather different data. The network system would only need an added 24 MHz crystal and antenna since other elements, such as baluns and switches, are already included in the package. The versatility of the highly integrated MC1322X PiP could also provide the "forest watch" nodes with additional functionality to assist the firefighters by gathering meteorological information, such as humidity, temperature and precipitation. As an added benefit, forest watch nodes could even function as emergency stations, in case a person gets lost in the woods.

According to the paper "Estimates of CO^2 from fires in the United States: implications for carbon management," U.S. fires release about 290 million metric tons of CO^2 a year, the equivalent of 4 to 6 percent of the nation's carbon dioxide emissions from fossil fuel burning. Much of this could be avoided by using simple wireless sensor networks, such as that described in this article[2].

1 http://www.nifc.gov/fire_info/nfn.htm, "Fire Information – National Fire News", National Interagency Fire Center, February 22nd 2008

2 Wiedinmyer, C., and J.C. Neff. 2007. Estimates of CO2 from Fires in the United States: Implications for Carbon Management. Carbon Balance and Management

Gonzalo Delgado Huitrón is part of the Wireless Connectivity Operation in Mexico. He is a software engineer with expertise in SMAC, BeeStack™, Synkro and 8/32-bit microcontrollers.

Gonzalo Delgado Huitrón

Reducing Home Power Consumption with Wireless Controllers

Home Automation and Power Efficiency

Did you know that a TV can consume up to 16 watts per hour when it is turned off? Is your water heater working at 4 a.m., when nobody is taking a shower? There are many ways to reduce energy consumption at home, and with the help of wireless network standards, this task can be achieved easily and efficiently.

Each day we hear that energy reserves are getting lower, shortages and blackouts are present and people are looking for new ways to save energy. So, beyond these inconveniences and the rising cost of energy use, why should we care about energy conservation? The answer is very simple—numbers show that almost 87 percent (refer to Figure 1) of worldwide energy is generated using carbon-based fuels, which, when burned, release massive quantities of pollutants into our environment.

One solution to our non eco-friendly energy production is to use green technologies, such as wind, geothermal, solar and hydroelectric power. The bad news is that this change cannot come about overnight (it takes several years, and new technologies are coming everyday). This is why energy conservation now is so important. Another factor to consider is that energy costs have gone up four times[2] in the U.S. in the last 30 years and more drastic increases are expected as shortages continue.

Most residential power consumed is in heating and cooling, followed by lighting and appliances (refer to Figure 2). This is where our energy efficiency efforts should be focused.

There is a very simple law of electronics to follow: I = V/R (Ohm's Law). This means that in order to have a current (I), we need voltage (V) and resistance (R). If we remove one of those items from the equation, then no current will flow. In other words, "If it is not being used, turn it off." It's as easy as that.

Another problem that might not look very alarming is the standby mode. During this mode, the devices are not completely on or off, but they do consume some power (refer to Figure 3) to be ready to start up when needed. The total annual standby power consumption in the U.S. can range from 5 to 26 percent of the total annual bill[4].

In 15 European countries in 2000, the total household energy lost to standby was estimated at 94 billion kWh[6]. According to the Energy Information Administration (EIA), the average standby power consumed each day per house was 60 W[7]. That translates to 525.6 kWh per year (0.060 kW x 8760 hrs). There are many initiatives to reduce standby power to 1 W daily, but disconnecting every single device that is not in use (then reconnecting when needed) would be tedious and time consuming. We need other solutions, and this is where technology comes in to help us.

IEEE 802.15.4/ZigBee® networking standard can be the solution we're looking for. Applications focused on monitoring and

World Energy Supply in TeraWatts[1]

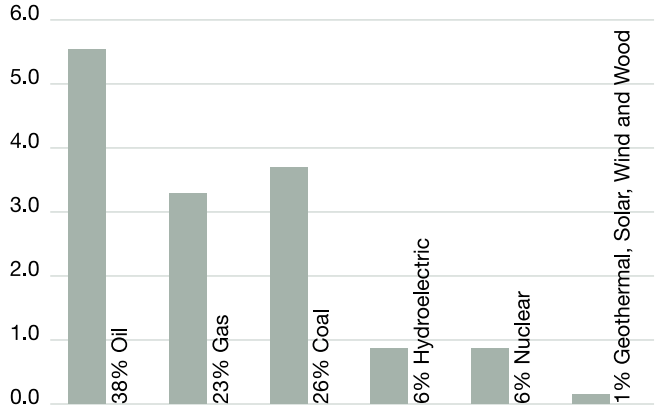

Figure 1

Residential Energy Consumption[3]	
Heating/Cooling	42%
Lighting/Appliances	36%
Water Heating	14%
Other	8%

Figure 2

controlling energy usage are primary markets for IEEE 802.15.4/ZigBee, and this is where Freescale takes the lead. Some key features that make it a good choice are: battery life (devices can last for years), complexity, range (up to 300m), security (128 AES) and cost (refer to Figure 4).

Imagine each house having its own remote controller, not only to turn on/off devices but also to record if they are being used or not and how much power they are using. By using fewer devices inside the home, interior heat could be reduced and air conditioners won't have to work so hard.

Alternatively, using a multi-socket power bar can make it easier to turn off several appliances at once. By adding a wireless controller, you would have the convenience of being able to turn off all devices in a certain section of your house (living room, kitchen, entertainment room) remotely.

Water heating consumes a lot of energy, whether it uses gas or electricity, but a remote controller for the gas/electric heaters could be implemented using Freescale's MC1321X family of ZigBee® ICs. Most of the time people take a shower in the mornings or in the night. All other times the water heater could be turned off remotely to conserve energy.

Automated meter reading (AMR) is one of most promising ZigBee applications (the profile is still being defined by the ZigBee

Power Consumption in Standby Mode[5]	
Appliance	**Watts: Standby**
LCD TV	1.5 (average)
Plasma TV	1.5 (average)
RPTV	3 (average)
DVD	2
Hi-Fi system	1
Computer	1–5
ADSL modem	2–10
Microwave	3
Phone	2
Digital Clock	1

Figure 3

Alliance). Linked to a "smart house," this could have instant benefits. If too much power is being used, instructions through the wireless link could direct the house to enter into low power mode and automatically turn off devices that are not used.

Using Freescale's IEEE 802.15.4/ZigBee products, such as the MC1321X family (with an 8-bit MCU) and the MC1322X family (with a 32-bit MCU), these hardware implementations could be done easily. From the software side, Freescale offers several stacks for the software implementation (SMAC, IEEE 802.15.4 MAC and BeeStack™). Using our BeeKit™ configuration tool, new projects can be easily created. With just a few mouse clicks, you can create an on/off light switch controlled through the ZigBee protocol without writing a single line of code.

Saving energy is not complicated. You just have to focus on the times when energy is really needed and when it is not. With a little help from Freescale's transceivers, systems-on-chip and system-in-package solutions, this can be accomplished.

Comparing Wireless Technologies			
Feature(s)	**IEEE 801.11b Wi-Fi**	**IEEE 802.15.3 Bluetooth**	**IEEE 802.15.4/ZigBee**
Battery Life	Hours	Days	Years
Complexity	Very Complex	Complex	Simple
Nodes/Master	32	7	65540
Latency	Enumeration up to 3 seconds	Enumeration up to 10 seconds	Enumeration 30 ms
Range	100 m	10 m	70 m–300 m
Extendability	Roaming possible	No	Yes
Base Data Rate	11 Mbps	1 Mbps	250 Kbps
Effective Throughput	5–7 Mbps	700 Kbps	100 Kbps
Security	Authentication Service Set ID (SSID)	64-bit, 128-bit	128-bit AES and Application Layer
Application	Computer Networking	File Transfer	Monitoring and Control

Figure 4

1 World Energy Intensity: Total Primary Energy Consumption per Dollar of Gross Domestic Product using Purchasing Power Parities, 1980–2004 (XLS). Energy Information Administration, U.S. Department of Energy (August 23, 2006). Retrieved on 2007-04-03.

2 ftp://ftp.eia.doe.gov/pub/consumption/residential/2001ce_tables/enduse_consump2001.pdf, Energy Information Administration

3 ftp://ftp.eia.doe.gov/pub/consumption/residential/2001ce_tables/enduse_consump2001.pdf, Energy Information Administration

4 http://eetd.lbl.gov/EA/Reports/45967.pdf, "Whole-House Measurements of Standby Power Consumption" J.P. Ross and Alan Meier,

5 http://eetd.lbl.gov/EA/Reports/45967.pdf, "Whole-House Measurements of Standby Power Consumption" J.P. Ross and Alan Meier,

6 http://www.greenpeace.org/international/campaigns/climate-change/take_action/your-energy, Greenpeace International

7 http://www.esru.strath.ac.uk/EandE/Web_sites/01-02/RE_info/Standby.htm, Energy Systems Unit at University of Strathclyde, in Glasgow, Scotland, UK.

Gonzalo Delgado Huitrón is part of the Wireless Connectivity Operation in Mexico. He is a software engineer with expertise in SMAC, BeeStack™, Synkro and 8/32-bit microcontrollers.

Raghavan Sampath

Green Lighting
Solar-based HBLED Lighting Solutions

Introduction

In the last fifteen years, energy demand has grown ten times and the cost of energy has increased four times. Blackout and brownout conditions have occurred in United States and other countries and could continue as energy demand increases and energy generation fails to keep up. According to a recent report from the North American Electric Reliability Corp., electricity demand in the U.S. alone is expected to grow by 141,000 megawatts in the next decade while only 57,000 megawatts of new resources have been identified. This leaves a shortfall of 84,000 megawatts, an amount equivalent to 160 large power plants.

Making electricity (much of it power lighting applications) creates 37% of all greenhouse gases, according to the United States Energy Information Administration. Hence, we should start looking at alternative energy resources, such as solar power, which is a green technology and does not cost as much as hydroelectric, geothermal or nuclear energy generation.

In the 21st century, there are as many as 300 million homes in under-developed countries deprived of adequate lighting. The reasons can vary, but may include affordability (the cost is too high), lack of infrastructure and supply and demand mismatches. Many homes use homemade kerosene lamps or candles for lighting. These dim, yellow, smoky light sources can be hazardous and are not eco-friendly. In the long run, they are more expensive light sources than bright, white, solar-powered high brightness light emitting diodes (HBLEDs).

Why HBLED for Lighting

Longer life

LEDs last longer than any other light source, in excess of ten years in many applications. They contain solid state technology similar to that used in the latest microprocessors. These solid state devices have no moving parts, no fragile glass environments, no mercury, no toxic gasses and no filament. There is nothing to break, rupture, shatter, leak or contaminate. Unlike typical conventional light sources, LEDs are not subject to sudden failure or burnout. There is no point in time at which the light source ceases to function. Instead, LEDs gradually degrade in performance over time. Most LED's are predicted to deliver an average of 70 percent of initial intensity after 50,000 hours of operation. In an application where the light source

Light Source Comparative Study

Lamp Type	Homemade Kerosene	Incandescent	Compact Flourescent	WLED
Efficiency (Lumens/watt)	0.03	5–18	30–79	25–50
Rated Life (Hours)	Supply of Kerosene	1,000	6,500–15,000	50,000
Durability	Fragile and Dangerous	Very Fragile	Very Fragile	Durable
Power Consumption	0.04–0.06 liters/hour	5W	4W	1W
CCT °K	~ 1,800°	2,652°	4,200°	5,000°
CRI	~ 80	98	62	82
$ After 50,000 Hours	**1251**	**175**	**75**	**20**

Figure 1 Source: Light up the World Foundation

LED Efficiency (Lumen/watt)

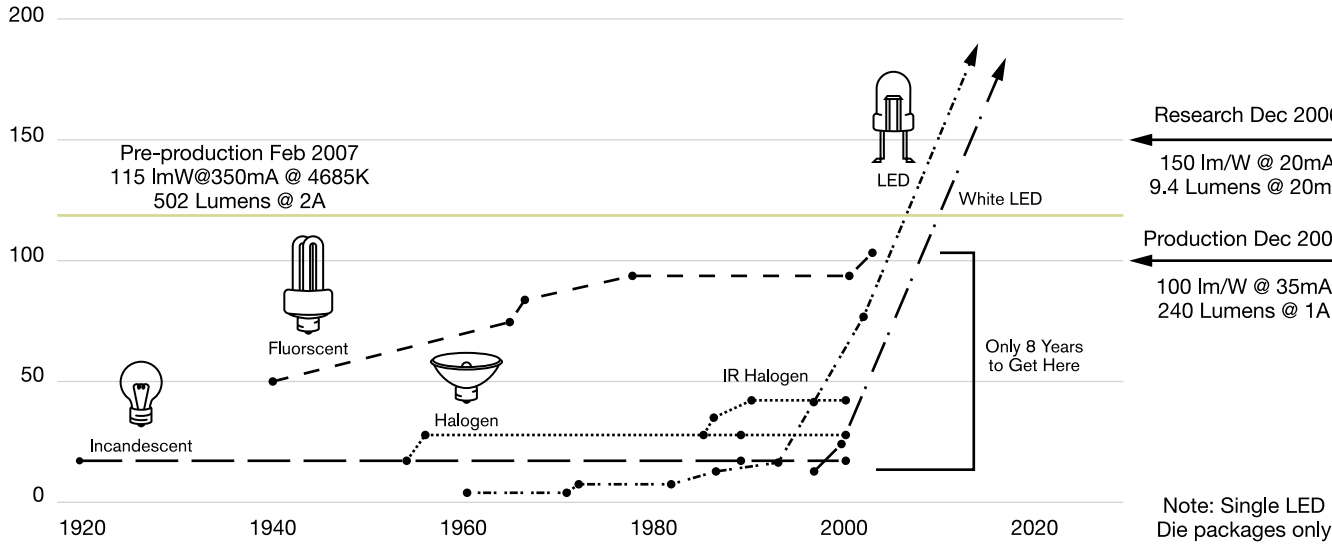

Figure 2 Source: Philips-Lumileds, Cree, Seoul Semi

would be used for 12 hours per day, 365 days per year, this could result in a system lifetime of over eleven years with only 30 percent degradation (70 percent lumen maintenance) from initial luminous output.

Reduced maintenance costs

Since LED-based light sources last up to ten times longer than a normal light source, there is no need for frequent replacement, reducing or even eliminating ongoing maintenance costs and periodic re-lamping expenses. This can be particularly important in critical, regulated lighting applications, such as buoys, beacon lights, emergency exit lighting, back up lighting and security lighting that would normally require scheduled, periodic bulb replacement.

More energy efficient

LED light sources are more efficient than incandescent and most halogen light sources. White LEDs today deliver more than 20 lumens per watt and are predicted to achieve greater than 50 lumens per watt in the future. When choosing solid state lighting as an alternative, it is important to consider the total system level benefits. For example, with superior lumens per watt performance, LEDs used in a building lighting system consume less energy per hour than competing lighting sources, making them more eco-friendly and cost-effective.

LED benefits

- Reliable (100K hours)—reduced maintenance costs
- More energy efficient—green, cost-effective solution
- Instant on and fully dimmable without color variation—pulse width modulation (PWM) control
- No mercury—conforms with environmental regulations
- Low-voltage DC operation—eliminates high-voltage connections

How Solar Powered LED Lighting Works

Figure 3 is a block diagram that illustrates the solar based LED lighting implementation.

A solar panel converts solar energy to electric voltage, which is stored in a battery. Freescale's 8-bit HCS08QG4 /8 microcontroller (MCU) with a two-channel sixteen bit timer is used for battery charging and monitoring and to drive the LED's. PWM is used for battery charging and an analog-to-digital

Solar-based LED Lighting

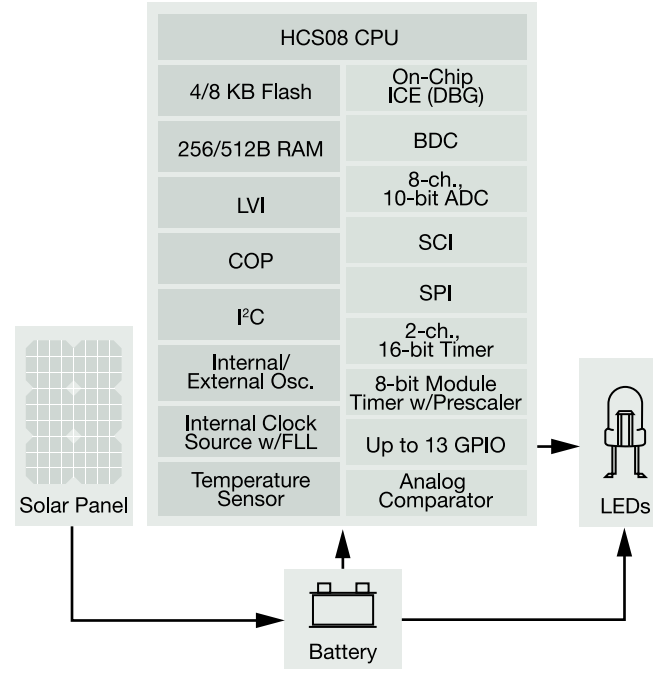

Figure 3

converter (ADC) is used for monitoring the battery voltage. If that voltage falls below 50 percent of a full charge, the LED's brightness is automatically reduced to 50 percent by varying the duty cycle of the second PWM channel. The idea is to provide light for longer durations, even at lesser brightness. If the battery voltage drops to ten percent, then the MCU turns the LED off to ensure the battery is not completely drained.

Choosing the Right Microcontroller

Multiple alternatives exist in selecting an MCU for the solar powered LED application; the HCS08QG4/8 device offers an excellent combination of the features needed at a very competitive cost.

The MC9S08QG8/4 extends the advantages of Freescale's HCS08 core to low pin count, small-package options. QG devices are low voltage, with on-chip in-circuit flash memory programmable down to 1.8V. They include the standard features of all HCS08 MCUs, including wait mode and multiple stop modes, strong analog capabilities, a complete set of serial modules, a temperature sensor and robust memory options.

MC9S08QG8/4 Features	
Features	**S08QG**
Core	HCS08
Flash	8K/4K
RAM (byte)	512/256
Bus Frequency	10 MHz
ADC	Up to 8 channels (10-bits)
Analog Comparator	Yes
Keyboard Interrupt	Up to 8-pins
Timers (up to)	One 16-bit timer (2 channels) One 8-bit timer
SCI	1
SPI	1
I^2C	1
Operational Voltage	1.8—3.6V
Package	8 DIP/SOIC/QFN 16-pin DIP/TSSOP/QFN

Figure 4

Applications

Solar-powered HBLED lighting can be used for street lights, home lighting,emergency lights and rural lighting.

Traffic lights are operational for the entire day in most cities and street lights for much of the day. Substituting HBLED technology for the traditional halogen or compact fluorescent lamp (CFL) solutions can provide enormous savings in terms of reduced energy consumption and lower maintenance costs.

The International Finance Corporation (IFC), the private sector investment arm of the World Bank Group's "Lighting the Bottom of the Pyramid" project, plans to sell solar powered LED lighting systems to 1.6 billion people around the world that are not connected to the electrical grid. Lacking access to electrical lighting, many of these people and their businesses instead rely on carbon fuels, such as kerosene, for their lighting needs. By using solar energy and LED lighting in these locations, we can reduce the health hazards and greenhouse gas emissions associated with the burning of fossil fuels.

A Last Note on Energy Savings

In 2006, China was the world's second largest electricity consumer, generating 2,475 billion kWh. Twelve percent of that went to lighting. By 2020, if LED efficiency reaches 150 lm/W, and penetrates just 30 percent of the Chinese market, the energy savings could reach 200 B kWh per year.

Conclusion

LED lighting technology is a significant, quantifiable, energy and cost reducing lighting alternative that provides low maintenance solutions for a wide variety of lighting applications. By using significantly less electricity than more traditional lighting methods, LEDs also help the environment by cutting electric power generation and its associated greenhouse gas emissions.

Raghavan Sampath graduated from Bangalore University in 2000 with a Bachelor of Engineering in Electronics and Communication. He has been with the company for two years and is a Field Application Engineer at Freescale Bangalore.

Virginia N. MacDonald

Evaporator Superheat Control System
Refrigeration Cycle Optimization

Overview

With global energy prices soaring to record high levels and no end in sight, home appliance manufacturers struggle to find innovative and effective ways to improve energy efficiency. Among the most energy-consuming and expensive of all home appliances are our compressor-based heating and cooling systems. These include living space heating, cooling and comfort control equipment, as well as refrigerators and freezers which are present in millions of households throughout the world.

This article describes a Freescale control system which can be used in compressor-based home appliances to optimize refrigeration cycle efficiency. Specifically, an important parameter called “evaporator superheat” is accurately monitored and controlled to provide near-perfect evaporator heat absorption across a wide range of thermal load conditions. The end result is significant utility cost savings to the homeowner as well as prolonged equipment life.

Functional Description

A typical refrigeration cycle schematic is shown in Figure 1. This cycle is utilized in one form or another in all common refrigeration, air conditioning and heat pump appliances. The cycle uses a suitable refrigerant compound that is formulated to change phase from liquid to gas and vice versa, at suitable temperatures and pressures for a particular application. We will use a home air conditioning application to describe how the cycle works.

The cycle starts at the system compressor. The refrigerant at the compressor input is a cool, low pressure gas. The compressor physically compresses the refrigerant to a hot, high pressure gas, which then enters the condenser coil. Here the refrigerant changes phase from a hot, high pressure gas to a hot, high pressure liquid. The key point is that the refrigerant phase change that occurs in the condenser causes a very large heat transfer from the condenser coil to its surroundings (in this case, the outdoor air).

Hot, high pressure liquid refrigerant emerging from the condenser coil is then passed through some type of fixed nozzle or metering device, exiting as a cold, low pressure liquid. Finally, this liquid is passed through the evaporator coil, where it changes phase to a cold, low pressure gas. Again, a very large amount of heat is transferred, but in this case the heat is absorbed by the evaporator coil from its surroundings (in this case, the indoor air). The entire cycle is then repeated until the indoors is sufficiently cool and comfortable. In short, heat is absorbed from the indoors and ejected to the outdoors via this special arrangement of coils, refrigerant pressures and temperatures.

Typical Refrigeration Cycle

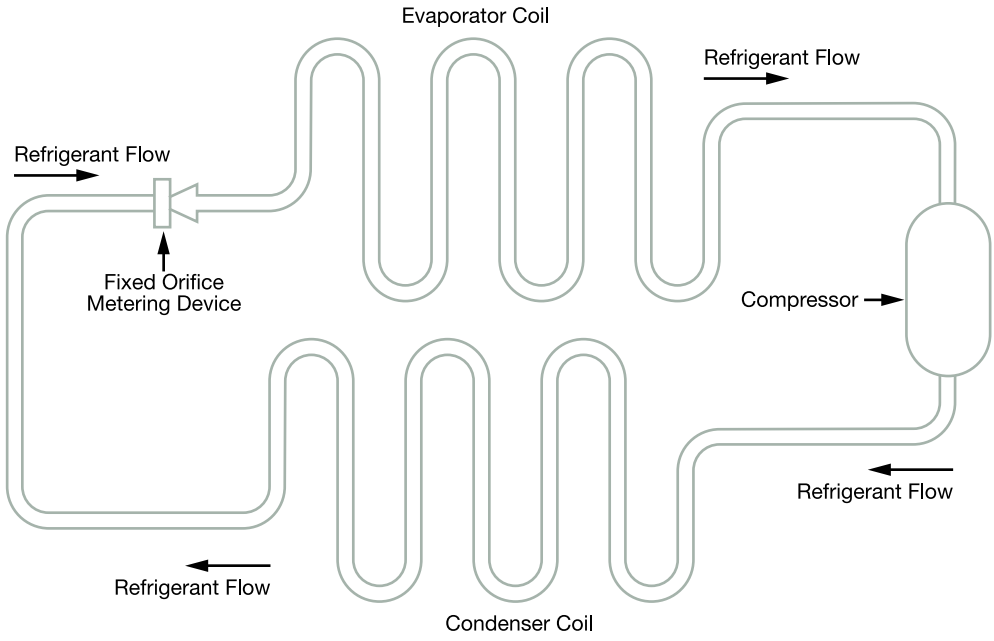

Figure 1

The fixed orifice metering device is a major contributor to degraded system performance. In the vast majority of air conditioning systems, the metering device is sized for "nominal" design conditions (a certain quantity of heat energy transferred per unit of time). Unfortunately, heat loads can vary significantly in real-world applications, and system performance degrades significantly as actual load conditions vary from nominal. Factors such as outdoor temperature excursions, open windows and doors, sunlight and the operation of other home appliances all contribute to significant heat load variations within the home.

When the heat load is higher than nominal, the liquid refrigerant in the evaporator coil starts to boil sooner. That is, when under nominal load the last bit of evaporation occurs at the coil outlet, but under high load conditions this last bit of evaporation occurs upstream from the coil outlet. This condition is equivalent to reducing the effective size of the evaporator coil, resulting in reduced heat absorption and lower efficiency. Moreover, in the presence of extreme heat loads the system compressor can overheat, which impacts its reliability.

Conversely, when the heat load is lower than nominal, the last bit of liquid refrigerant boils downstream of the evaporator coil outlet. Evaporation downstream of the evaporator coil absorbs heat from some source other than the conditioned space, reducing the air conditioner's cooling efficiency. Worse yet, this phenomenon can result in liquid refrigerant reaching the compressor. This can cause permanent damage to the compressor and incur very high service costs to both the equipment manufacturer and the homeowner.

Evaporator Superheat Control System

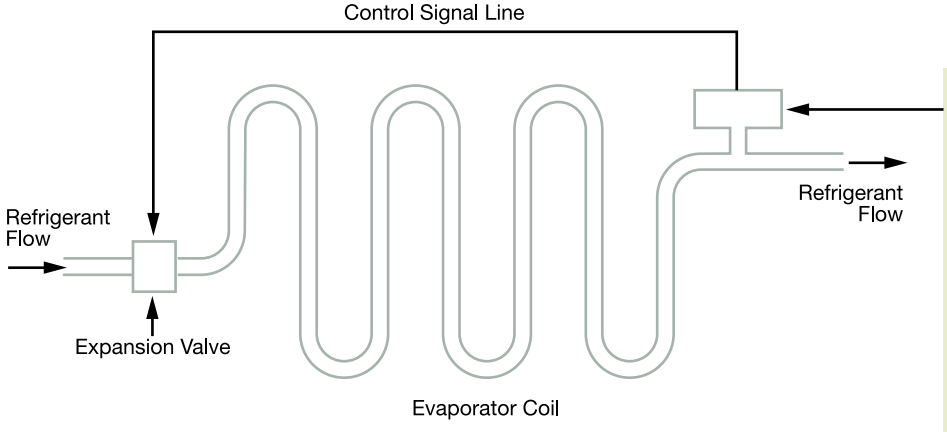

Superheat Control Module contains:

- Freescale industrial pressure sensor
- Freescale MC9S08GB60 microcontroller
- Temperature sensor
- Operating software
- Power transistor (optional)
- ZigBee® transceiver (optional)

Figure 2

Control System Logic Flowchart

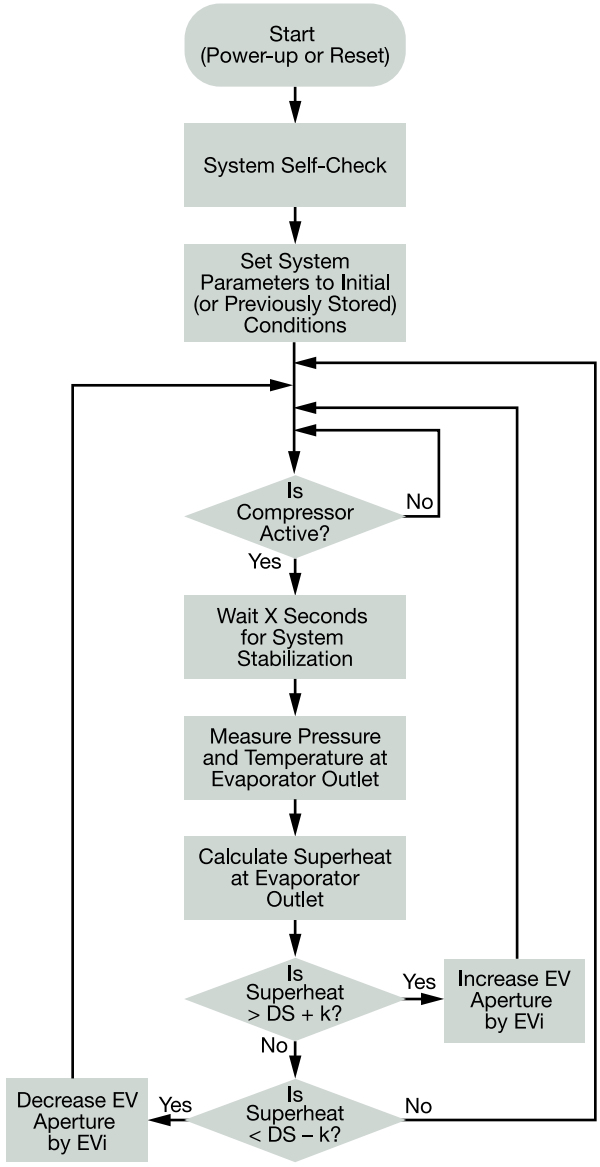

Key:
EV=Expansion valve
X=Constant, defines stabilization time from state change
DS=Desired superheat in degrees
k=Constant, defines "deadband" around DS in degrees
EVi=Expansion valve increment in mm

Figure 3

The Freescale evaporator superheat control system eliminates these anomalies by monitoring and controlling an important refrigerant parameter called superheat. Superheat is measured at the evaporator outlet and represents the difference between the actual refrigerant temperature and its boiling point temperature at the measured pressure.

The control system is comprised of a Freescale high-performance industrial pressure sensor (currently under development), a Freescale MC9S08GB60 microcontroller, a temperature sensor, operating software and an expansion valve. A power transistor may also be needed, depending on the type of valve used in the application. An optional ZigBee transceiver may be included to provide a wireless link to another device, such as a utility meter or service technician's computer. A schematic of the control system is shown in Figure 2.

As the heat load increases from nominal, the evaporator superheat rises. The controller senses this condition and adjusts the expansion valve aperture to allow more refrigerant into the evaporator, thus causing the superheat to fall.

Conversely, as the heat load decreases from nominal, the evaporator superheat falls. The controller senses this condition and adjusts the expansion valve aperture to allow less refrigerant into the evaporator, causing the superheat to rise. The result is optimized evaporation within the evaporator coil and maximum heat absorption from the indoor air. A logic flowchart for the control system is shown in Figure 3.

Recall that the above description is specific to home air conditioning. However, the same control methodology can be applied to refrigerators, freezers, heat pumps or any other home appliance that uses a compressor-based refrigeration cycle. The methodology is also directly applicable to commercial and industrial refrigeration systems.

Note: The Freescale evaporator superheat control system is 100 percent compatible with all commonly used refrigerants, including the environmentally-friendly R410A and R134A types.

Virginia MacDonald is currently a Product Definer from the Freescale Sensors and Actuators Solutions Division. She earned a BSEE from Clarkson College in 1980, and has over 27 years of engineering experience in the electronics, sensors, machine vision and HVAC industries. She has been awarded 5 U.S. patents and currently resides in the Phoenix, Arizona area.

Peter Balazovic

Sensorless AC Motor Control in Home Appliances

Introduction

World-wide interest in water and energy conservation as well as the overall environmental friendliness of new products and services definitely impacts the home appliance market. The actual regulatory standards and recommendations, either nationwide or multination (EU standards), enforce basic demands for a new generation of home appliances, particularly washing machines, dryers, dish washers and refrigerators. To meet these demands while still reducing system costs, enhanced microcontrollers equipped with appropriate software must be employed.

For the home appliance industry, a digital signal controller (DSC) based solution is proposed, combining the processing power of a digital signal processor with the functionality and ease of use of a microcontroller on a single chip. A flexible set of peripherals allows the designers to realize numerous functions, such as standard motor drive algorithms, advanced control algorithms, sophisticated feedback signal sensing, power factor correction schemes and communication with external environments.

Employing variable speed motor drives gives the designer an opportunity to use more sophisticated control programs. This enhances appliance performance and increases overall energy efficiency. Thus, home appliances equipped with variable speed drives and intelligent controls outperform those with uncontrolled, fixed-speed motor drives.

The majority of electric motor drives in home appliances are controlled very simply by either fixing the motor speed at pre-defined levels or by running it directly from the AC main supply without additional control electronics. The alternating current mains-supplied single-phase induction motor (ACIM) is extensively used because it is low cost, robust and reliable. However, ACIM solutions also have significant drawbacks, such as lower efficiency and reduced effective speed control, which hinder control improvements designed to enhance product flexibility and meet regulatory guidelines as well as customer expectations. In contrast to single–phase ACIM, the variable speed drives address energy efficiency requirements by maintaining precise torque control, thereby improving overall appliance efficiency.

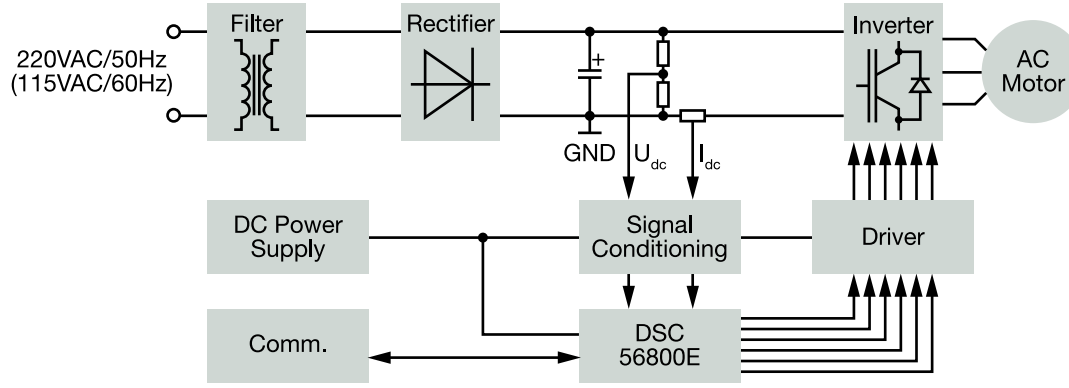

Figure 1

AC Motor Choice

There are different motor categories that can drive home appliances, such as ACIM, permanent magnet (PM) motors or switched reluctance (SR) motors. The product of the electronic drive efficiency and the electric motor efficiency determines the overall energy efficiency of the system.

In general, to achieve variable speed with an ACIM, variable frequency and voltage need to be supplied to the ACIM. This is referred to as constant volt per Hz control. The ACIM speed drive efficiency can be further improved by using vector control. This requires accurate velocity information sensed by a speed or position sensor attached to the rotor. However, the additional sensor, connector and associated wiring increase the cost of the motor drive.

PM motors use permanent magnets to establish the flux instead of creating it from the stator winding. Replacing the electromagnetic excitation with permanent magnets has several advantages. Most obvious is the absence of excitation losses, which means PM motors have higher power density than comparable DC motors. PM motors have no mechanical commutator, and power density exceeds the AC induction motor because there is no flux oriented current. Overall efficiency approaches 90 percent while single-phase ACIM efficiency reaches only about 70 percent. Permanently excited synchronous motors are very attractive solutions for home appliances, however, they cannot be operated by applying AC main supply to the stator winding.

Variable Speed Motor Drives

Experts say that the total electric power consumption can be reduced up to 30 percent if electric motors are controlled to optimize the supplied energy. The controller for a variable-speed drive actually controls the rotational speed of an AC electric motor by regulating the frequency of the electrical power supplied to the motor. Electronic variable speed drives allow electric motors to continuously operate over their full speed range.

As Figure 1 indicates, the complete system for a variable speed drive includes the EMI filter, the input rectifier, the on-board DC power supply, a DSC, the signal conditioning circuits, the power inverter and gate driver.

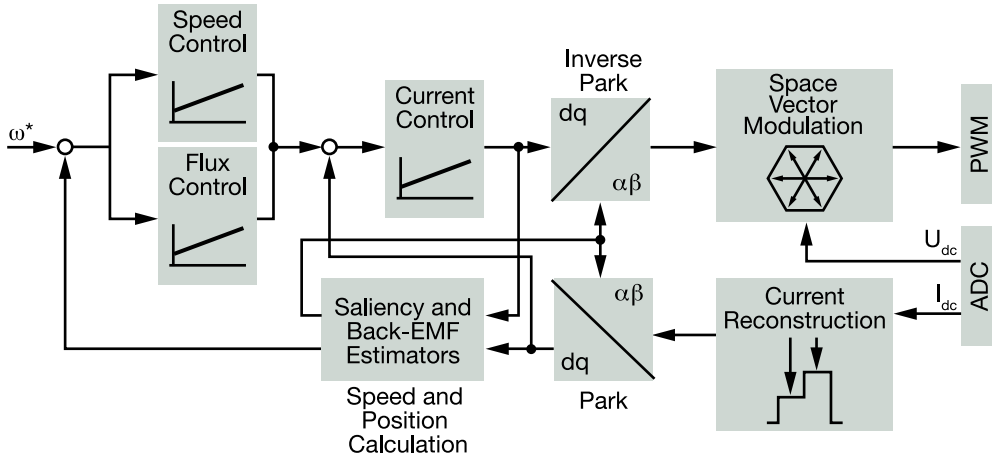

Figure 2

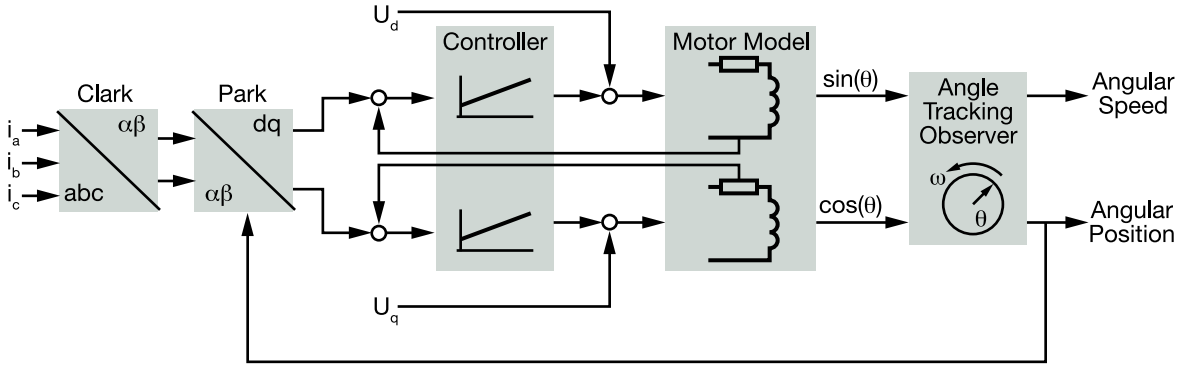

Figure 3

Motor Control Strategy

The open loop scalar control represents the most popular control strategy for variable speed drives predominantly used with ACIM. As depicted in Figure 2, the scalar control is based on the variation of supply voltage frequency. Voltage magnitude is proportional to the voltage frequency and changes according to the frequency variation. This technique requires modest computational power that can be handled by an 8-bit microcontroller.

The great advantage of this simple method is that it is a sensorless mode, and the control algorithm does not need information about the angular speed or actual rotor position. Speed dependence on the external load torque, however, is a big disadvantage, which results in reduced dynamic performance. Because of this, electric motors under scalar control must be oversized in order to deliver the required torque during load transients. In addition, scalar techniques result in an inefficient system, power factor degradation in the utility's network and noisy operation. In this control approach the energy efficiency might degrade down to the 50 percent range from theoretical maximum. Motor model-based method estimates an electromotive force in which the electrical position information of machine is encoded. The estimator depicted in Figure 3 consists of the stator current observer with knowledge of RL motor parameters, fed by applied motor voltage and measured current.

Yet, market demands require the highest possible dynamic performance and operating speed range. Vector control (field oriented control) of AC machines, as a novel approach in electrical drives, provides very good performance in comparison to scalar control. Vector control eliminates most of the disadvantages of constant volt-per-hertz control.

Within a vector control system the synchronous frame current regulators have become the industry standard for inverter current regulation. The inverse Park ($\alpha\beta$/dq) reference-frame transformation function calculates the stator currents and voltages in a reference frame synchronized to the rotating rotor field. All the electrical variables have DC steady state values when viewed in a rotating reference frame, enabling a simple PI regulator to provide zero steady state error. Additionally, it is possible to set up the coordinate system to decompose the current vector into a magnetic field generator and a torque generator. The structure of the motor controller is shown in Figure 3.

The inner current loop calculates the direct and quadrature stator voltages required to create the desired torque and flux currents. The Park (dq/$\alpha\beta$) functions transform these voltages into three-phase AC stator voltage demands in the stationary reference frame. The motor currents are sinusoidal, thus producing smooth torque, which minimizes acoustic noise and mechanical vibration. The outer velocity loop adjusts the applied torque magnitude, which is directly proportional to quadrature torque-current and enables maintaining the required angular velocity. In order to effectively extend the operating speed range above base speed, an additional flux-weakening loop is added, which manipulates directional stator flux-current.

Sensorless Control

To run an AC motor in vector control mode, it is important to synchronize the frequency of the applied voltage with the position of the rotor flux coming from permanent-magnet in the rotor. This results in a sensorless mode of operation where the speed and position calculation algorithm replaces the sensor. The sensorless control approach to AC PM motors is an innovative motor drive feature that improves reliability and maintains high performance levels without additional cost. The system should be more accurate and efficient, smaller, lighter, less noisy and should have more advanced functionality at a lower cost. Sensorless algorithms can be broadly divided into two major groups—those that utilize magnetic rotor saliency for tracking rotor position and those that estimate rotor position from calculated motor models.

In a refrigerator, conventional control turns the compressor on and off to maintain the temperature in a predefined range. With sensorless PM motor control, the controller can accelerate the compressor to its target speed within a few seconds and can regulate speed to within 1 percent of its target. The smooth running of the compressor reduces audible noise, and the lower operating speed helps minimize the temperature cycles in the refrigeration compartment, improving overall efficiency.

Normally, a washing machine operates in two basic cycles—tumble-wash and spin-dry. During the tumble-wash cycle, the motor drive operates at low speed with high torque. During the spin-dry cycle, the motor drive works at high speed for short time periods. In newer models, the estimated speed ripples and calculated load torque provide valuable information on the washing load distribution. The speed ripple is used to estimate the load imbalance before starting the spin cycle. A variable speed motor can be used to regulate the clothes tub rotational speed and direction to redistribute the clothes to correct the imbalance.

MC56F80xx Family of DSCs

Freescale's 56800E core is an ideal solution for this particular application. It processes all motor control functions, including space vector modulation, vector control, current, flux and velocity loop control. Digital control improves drive reliability by reducing the number of discrete components found in early designs and facilitates advanced algorithms for optimal motor performance. Sensorless speed vector control discussed above is implemented on the MC56F8025 DSC with all control routines using C-callable, optimal assembler language with fractional numerical representations.

56F802x3x Family Block Diagram

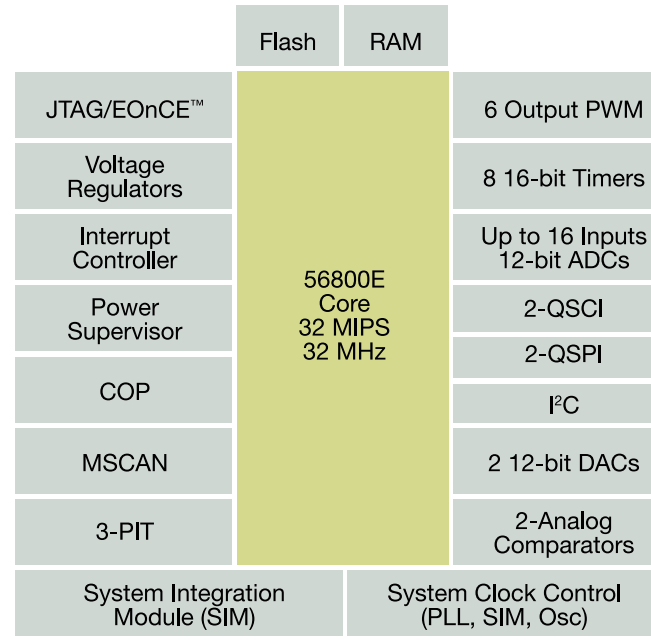

Packages: 32 LQFP, 44 QFP, 48 LQFP, 64 LQFP

Figure 4

The MC56F80xx family members provide these peripheral blocks:

- One pulse width modulation (PWM) module with PWM outputs, fault inputs and fault-tolerant design with dead-time insertion, supporting both center-aligned and edge-aligned modes
- 12-bit analog-to-digital converters (ADCs), supporting two simultaneous conversions. ADC and PWM modules can be synchronized
- One dedicated 16-bit general-purpose quad timer module
- One serial peripheral interface (SPI)
- One serial communications interface (SCI) with LIN slave functions
- One inter-integrated circuit (I²C) port
- On-board 3.3 V to 2.5 V voltage regulator for powering internal logic and memories
- Integrated power-on reset and low-voltage interrupt module
- All pins multiplexed with general-purpose input/output (GPIO) pins
- Computer operating properly (COP) watchdog timer
- External reset input pin for hardware reset
- JTAG/On-Chip Emulation (OnCE™) module for unobtrusive, processor speed-independent debugging
- Phase-locked loop (PLL) based frequency synthesizer for the hybrid controller core clock, with on-chip relaxation oscillator

The application uses the ADC periphery to digitize analog feedback signals (voltage, current) before being processed. The ADC trigger is synchronized to the PWM reload flag, but the sampling time instances vary as a function of the actual PWM pattern. This configuration allows multiple conversions of the required analog values for the DC-bus current and voltage within the one PWM cycle.

The PWM module is capable of generating asymmetric PWM duty cycles in the center-aligned configuration. This feature enables three-phase current reconstruction in critical switching patterns.

The quad timer is an extremely flexible module, providing all required services relating to time events. The application requires two of the quad timer's channels:

- One channel for PWM-to-ADC synchronization
- One channel for system velocity control loop (1 ms period)

The application allows communicating with the front-panel master control over an isolated serial link. This allows speed profile information to be downloaded to the controller and motor speed and torque information to be uploaded to the home appliance's master controller.

Conclusion

It is widely acknowledged that home appliance demands drive technological advances and, in turn, the production of next-generation products. In response to this trend, vector control with sensorless functionality is a technique used for this purpose, and since it eliminates mechanical sensors, it enables a cost-effective solution. Even such demanding control techniques can be realized using the low-cost DSC based on the Freescale 56800E core, which combines both MCU and DSP capabilities.

The results also indicate satisfactory performance for other various motor drive applications, such as pumps, fans, blowers, dryers, etc. MC56F80xx devices designed with a focus on motor control applications are being integrated into many home appliance product designs, particularly where national or multinational regulations on power consumption are strict.

Freescale offers a comprehensive portfolio of embedded chips, application software and development tools that are optimized for electronic motor systems. These embedded solutions are designed to allow customers to extract a high level of functionality and performance out of highly reliable, cost-efficient motors. A digitally controlled motor in a washing machine is designed to provide a more efficient agitation cycle so that less water is needed and may allow a short-high speed spin cycle, resulting in drier clothes and less energy consumption. The chips incorporate motor-specific functions which are designed to enable the appliances and other motor-driven products to operate more efficiently and reliably.

Peter Balazovic is a system application engineer with the motor control and power efficiency group for the Consumer and Industrial markets. He works closely with the DSC/MCU product line for motor control applications and designs motor control algorithms at the Roznov Czech System Center. He joined Freescale Semiconductor in 2001.

Akshat Mittal

Way Forward to Greener Buildings
New Power Reduction Technique for Elevators

Introduction

A new buzz of the 21st century is environment friendliness. Societies, governments and corporations have lately realized the need of an environment friendly society. The change in attitude is a consequence of many facts, shortage of conventional sources of energy due to wasteful lifestyles, problems associated with disposal of waste like polythene, circuit boards, etc. and climate change threats in the form of global warming. Many corporations have started taking action to counter these problems. These include use of renewable sources of energy (solar panels), use of non-toxic material in product manufacturing processes, spreading awareness and designing environmentally friendly products (RoHS Compliance is an example). Focusing on energy efficiency at each level of product design has become the norm.

A substantial amount of energy savings can be done by optimizing the use of power in industries, commercial establishments and offices. This has been made possible by the presence of microprocessors and microcontrollers.

A respectable share of power consumed in buildings is claimed by the elevators. The type of drive, capacity and the total full load mass of the elevator, number of floors served, elevator system efficiency and the traffic pattern in the building are some of the variables affecting the energy consumption of elevator systems. In the past, many innovations have been done in the mechanical controls (Mitsubishi's VVVF Inverter Control Technology), manufacturing process (reduction in the amount of lead used) and electronic controls (standby modes in off-peak hours) to reduce power consumption in elevators.

Another aspect of power waste which we often forget is the usage pattern of elevator in peak hours. The focus of this article is to analyze, discuss and propose a control algorithm to reduce the overall power consumption of elevators and provide a working solution on a suitable microcontroller; keeping in mind the computational complexity, motor driving capability and cost. The proposed algorithm also reduces wait time and frustration levels along with power consumption during the up-peak and down-peak hours of an office building.

Power Reduction Challenges

Elevators have faced a constant aggressive push for more energy efficiency since the beginning. We have seen the drive system shift from AC two speed controls to primary voltage control, and then to inverter drive control in 1984. This has had a particularly big influence on the energy saving achieved by elevators. Standard elevators installed in buildings with less than 12 floors have seen a reduction in the energy consumption level of approximately 32 percent from the 1970s to the 1990s [1]. With the aim of further energy saving, work has also been done on the traction machine (traction motor). We have seen progress from worm gear methods to helical gear methods, in the 1990s, and a recent shift to gearless methods. Not to forget the replacement of MG (Motor Generator) drivers by the energy efficient VVVF (Variable Voltage Variable Frequency) drivers [2].

Control algorithms for elevator scheduling have evolved with time to reduce the passenger wait time, total elevator runs and enhance transport efficiency for a single or group of elevators [3] [4]. Advanced artificial intelligence neural networks, complex heuristic algorithms have been developed to achieve this and can already be found in modern elevators. The required computing power has increased since this type of controls have been implemented.

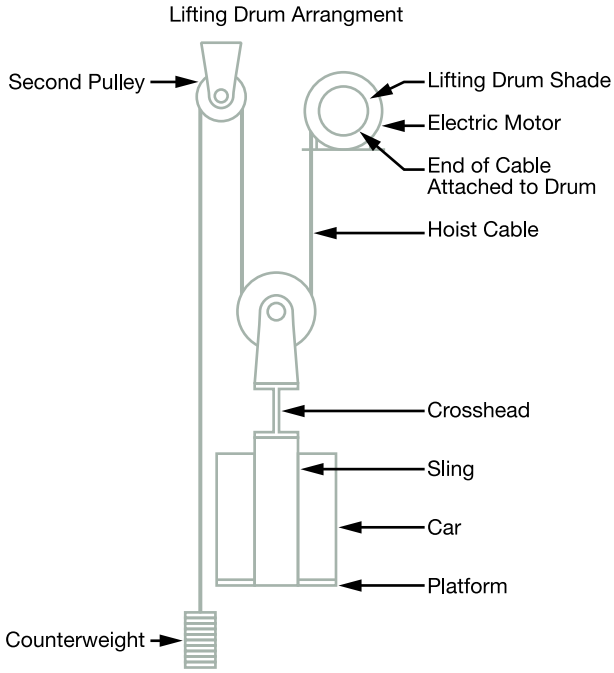

Figure 1

Elevator Function and Motor Theory

Commonly used in commercial buildings, geared traction elevators are driven by AC/DC electric motors. A typical traction elevator system is shown in Figure 1. Geared machines use worm gears to control mechanical movement of elevator cars by rolling steel hoist ropes over a drive sheave which is attached to a gearbox driven by a high speed motor. A brake is mounted between the motor and the gearbox to hold the elevator stationary at a floor. This brake is usually an external drum type and is actuated by spring force and held open electrically. Cables are attached to a hitch plate on top of the cab and then looped over the drive sheave to a counterweight attached to the opposite end of the cables which reduces the amount of power needed to move the cab. Both, the car and the counterweight ride separate rail systems. Both move in opposite directions. The traction motor is controlled by a microcontroller that directs the starting, acceleration, deceleration and the stopping of the elevator cab.

The control panel inside the elevator contains call buttons to choose a floor, door control buttons, alarm and hold buttons. Some elevators also have an overload indicator. The external control buttons request up and/or down travel.

A 3-phase permanent magnet synchronous motor (PMSM) is normally used to drive the traction elevator. A three-phase PMSM is a permanently excited motor. Its very high-power density, very high efficiency and high response, makes it suitable for most sophisticated applications in mechanical engineering. It also has a high overload capability. A PMSM is largely maintenance-free, which ensures the most efficient operation. The power-to-weight ratio of a PMSM is also higher than induction machines. Progress in the field of power electronics and microelectronics enables the application of PMSMs for high-performance drives.

Power Reduction Technique

A simple algorithm to control a single elevator is given below.

Pseudo Algorithm

1. Remain in IDLE ready to serve state at floor zero until interrupted by an interrupt.
2. Move to the level from where the request came and wait for floor button activation.
3. Stop on the level to serve an external request of same direction.
4. Stop at levels in descending/ascending order to complete all internal requests for upward/downward movements respectively.

Now imagine a crowded ground floor of an office establishment in the morning, heavy downward elevator traffic in the evening and mixed traffic during lunch time. Very often, the elevator is fully loaded at the starting floor (lower or top floor).

Working on the above mentioned control algorithm, an elevator stops at each level to serve requests in the same downward direction during down-peak hours but cannot accommodate more because it is already full. Such behavior increases the wait time for people at each level, causes frustration to the people inside the cab and most importantly causes a lot of wasted power due to acceleration and deceleration of the cab and extra door movements of the elevator.

A maximum-duty cycle of 35 percent is often assumed for traction elevator systems. This means the car can be in motion and drawing full-load currents for a maximum of 35 percent of a given time period. It has also been seen that peak currents are drawn by the motor while the elevator is accelerating or decelerating. These peaks normally persist for three to five seconds, and the amplitude of current is typically 2–2.5 times the full-load current [5].

A change in the above algorithm will counter all the above problems. As mentioned earlier, that elevators these days come with an overload indicator. The indicator can be used to ignore the external interrupts when the elevator is fully loaded. So step 3 in the above algorithm will become “Stop on the level to serve an external request of same direction IF elevator is 90 percent full; (because 100 percent full elevator should not move).”

Experiments done on a 10 story building with a Mitsubishi elevator show that time taken by the elevator to go from the top floor to the ground floor is double when the elevator stops at every floor (85 seconds) then when elevator do not stop on

Power Reduction Technique State Diagram

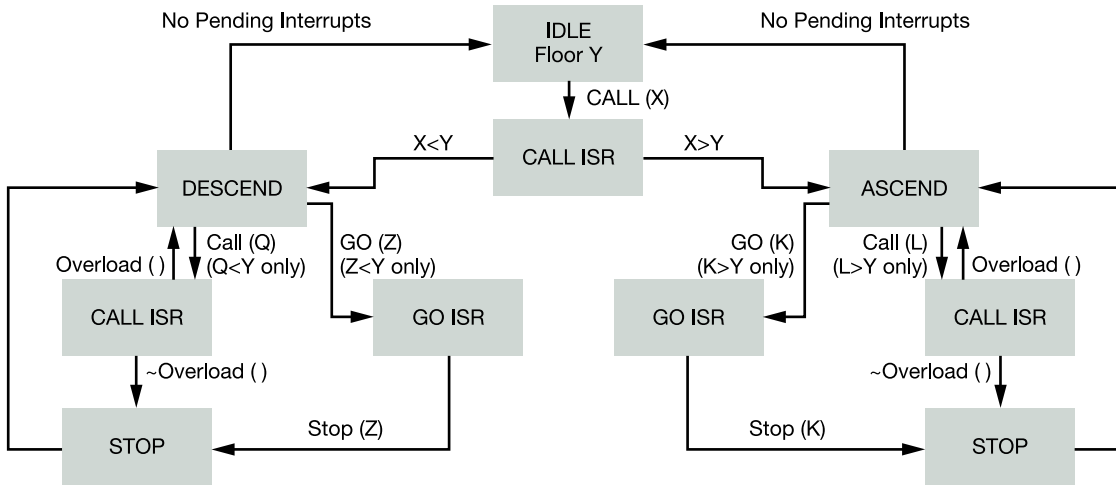

Figure 2

any floor in between (40 seconds). This shows that an elevator travels at constant speed for only half the duration of the commute time, and the other half is taken for accelerating and/or decelerating (assuming door open/close time to be negligible, which actually also consumes extra power). Thus there is a potential for 40–50 percent savings in energy through the proposed method during peak hours.

Technical specifications of Mitsubishi elevators show an estimated yearly power consumption of its elevators using permanent magnet synchronous traction motor in a typical office complex to be around 3000 kWh. The above calculation thus puts absolute power saving achieved by the technique to a remarkable figure of 1200–1500 kWh.

Working Solution

A state machine to explain the implementation of the proposed scheme is described in Figure 2.

Interrupts

CALL(x)=External Interrupt from Floor ‘x’ to call the elevator

GO(x)=Internal Interrupt from inside the elevator; occurs when one of the floor buttons is pressed to go to a particular floor

Functions

Ascend( )=function to control the PMSM motor to rotate the sheave in a direction for UP travel

Descend( )=function to control the PMSM motor to rotate the sheave in the counter direction for DOWN travel

Stop(x)=Stop the elevator at floor ‘x’

Overload ( )=Returns a 1 if elevator is loaded 90% of maximum load. (For an elevator with a maximum capacity of 1000 Kg)

State Description

IDLE=Elevator is stationary positioned at level Y

ASCEND=Elevator is traveling UP, Ascend( ) function is executed in this state

DESCEND=Elevator is traveling DOWN, Descend( ) function is executed in this state

CALL_ISR=Handles the processing if external call interrupts are pending

GO_ISR=Handles internal interrupts from floor buttons

STOP=Execute Stop(x) function

Motor control circuit

A PMSM drive circuit to implement the Ascend/Descend functions is shown in Figure 3. It contains an input diode bridge that rectifies ac-line voltage and a bank of capacitors, which then filters it. Together, these passive components form a simple ac-to-dc converter. The right side of the conversion circuit is a “full three-phase bridge” and converts DC-to-AC. High-side transistors apply positive voltage to the motor phases, while low-side transistors apply negative voltage. By controlling bridge transistors on and off states, the drive causes current to flow in or out any of the three motor phases.

Digital PMSM drivers control all transistors through pulse width modulation (PWM). PWM voltage regulates phase current, which in turn produces motor torque. Analog feedback voltage and current is immediately digitized before being processed. Software processes all motor control functions, including modulation, field orientation, current, and velocity loop control. Digital control improves drive reliability by reducing the number of discrete components found in early designs and facilitates advanced algorithms for optimal motor performance.

To implement the above proposed technique, a microcontroller with following capabilities is needed:

- Pulse width modulator modules
- High priority Interrupt controllers to handle overload sensors, brakes, fuse switches
- Fast ADC modules
- Serial port for the display panel
- High processing prowess for precise velocity control, door and cab positioning
- Sufficient RAM and on chip memory, to store the collected data from sensors

Traction Machine Control Circuit

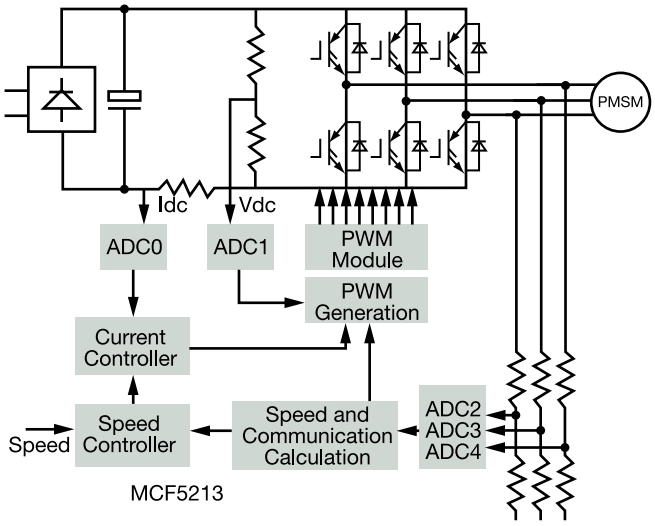

Figure 3

The MCF5213 is a member of the ColdFire® family of reduced instruction set computing (RISC) microprocessors. This 32-bit device is based on the Version 2 ColdFire core operating at a frequency up to 80 MHz, offering high performance and low power consumption. On-chip memories connected tightly to the processor core include up to 256 Kbytes of flash memory and 32 Kbytes of static random access memory (SRAM). On-chip modules include:

- V2 ColdFire core delivering 76 MIPS (Dhrystone 2.1) at 80 MHz running from internal flash memory with Multiply Accumulate (MAC) Unit and hardware divider
- Three universal asynchronous/synchronous receiver/transmitters (UARTs)
- Inter-integrated circuit (I^2C) bus controller
- Queued serial peripheral interface (QSPI) module
- Eight-channel 12-bit fast analog-to-digital converter (ADC)
- Four-channel direct memory access (DMA) controller
- Four 32-bit input capture/output compare timers with DMA support (DTIM)
- Four-channel general-purpose timer (GPT) capable of input capture/output compare, pulse width modulation (PWM), and pulse accumulation
- Eight-channel/Four-channel, 8-bit/16-bit pulse width modulation timer
- Two 16-bit periodic interrupt timers (PITs)
- Programmable software watchdog timer
- Interrupt controller capable of handling 57 sources
- Clock module with 8 MHz on-chip relaxation oscillator and integrated phase-locked loop (PLL)
- Test access/debug port (JTAG, BDM)

Conclusion

In this article we discussed the need for environmentally friendly products and the impact of power needs posed by rapid urbanization. We justified the need for reducing power consumption in every aspect of life. In an attempt to make the buildings greener, we proposed a way to reduce elevator power during peak hours.

A working solution to implement the scheme was discussed using MCF5213 from Freescale's ColdFire portfolio, because of its motor driving capability, high processing power and large number of peripherals.

References

1. Development of Energy-Saving Elevator Using Regenerated Power Storage System. Shinji Tominaga", Ikuro Suga, Hiroshi Araki, Hiroyuki Ikejima, Makoto Kusuma, and Kazuyuki Kobayashi; Industrial Electronics & Systems Laboratory; Mitsubishi Electric Co.
2. Energy Consumption Analysis for Geared Elevator Modernization: Upgrade from DC Ward Leonard System to AC Vector Controlled Drive, Ashok B. Kulkarni Thyssen Dover Elevators.
3. Dynamic Optimized Dispatching System for Elevator Group Based on Artificial Intelligent Theory, Cao Liting, Zhang Zhaoli, Hou Jue; Beijing Union Univ., Beijing
4. Elevator Group Control Scheduling Approach Based on Multi-Agent Coordination; Qun Zong Liqian Dou, Weijia Wang; College of Electric and Automation Engineering, University of Tianjin, City of Tianjin, China
5. Elevator World Magazine, January Issue, 1996

Akshat Mittal is an ASIC Design Engineer at Freescale Semiconductor India, and has been with the company for almost two years. He holds a Bachelors Degree in Information and Communication Technology. He specializes in digital IP design.

Matt Maupin

A Green Bee
ZigBee is Going Green

ZigBee® Technology is Green Technology

Going green is a major trend that is impacting all of our businesses. In the last year, we have seen more companies jumping on the green bandwagon then ever before. However, since day one, ZigBee technology has been focused on green issues. Among other things, ZigBee technology is ideal for sensing and controlling energy consuming devices, such as lighting and heating/ventilation/air-conditioning (HVAC), in both home and commercial buildings. Now that the world has become more concerned with rising energy costs, which have gone up 4X in the last 30 years[1], it is clear that ZigBee technology is increasingly viewed as a leading technology to help reduce energy cost in home and commercial buildings through Smart Energy products.

Energy Usage Outpaces Energy Generating Investments

The U.S. is a major consumer of electricity, and demand has increased 25 percent since 1990 while annual investment in generating electricity has decreased 30 percent in the same time period[2]. While some may say we need to increase supply through the building of new generation plants, or even through alternative energy sources (wind and solar power) that takes time and costs billions of dollars. An alternative is to try to control the amount of energy used. Energy can be 20 percent[3] of a building's operating cost, with lighting and heating/cooling using most of it in both residential and commercial buildings. They are prime targets for energy reduction through smart energy usage with lower energy components.

Commercial Building Energy Usage

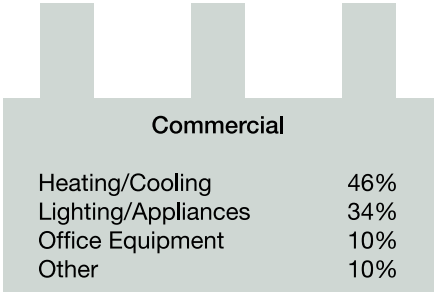

Figure 1 Source: Energy Information Administration

A study by the U.S. Department of Energy's Pacific Northwest National Laboratory showed that households with digital tools controlling temperature and price preferences saved on average 10 percent on utility bills while cutting peak loads on utility grids by up to 15 percent. This translates into $70 billion dollars saved over a 20 year period on new power plants and infrastructure expenditures. Utility companies are looking for smart ways to control rising energy usage while limiting capital expenditures. Utilities are looking to implement both load control and variable pricing to help manage load during peak times. Through load control, utilities companies can turn off customer loads, such as HVAC, for sort periods of time (a few minutes) during peak loads. In addition, utility companies can employ demand response, which allows them to communicate changing utility rates to the home to encourage voluntary energy conservation. These actions will help prevent rolling brownouts by reducing peak demand.

ZigBee technology helps address both these needs with an open and interoperable standard for communication from the utility to the home area network (HAN) for purposes of load control and demand response. This allows ZigBee-enabled meters to communicate and control ZigBee enabled devices in the home, such as the HVAC system. The home owner will also have the option of taking voluntary actions during demand/response periods to reduce their personal consumption by turning off appliances, lights, etc. This not only reduces the peak load on the utility grid but also helps home owners make smart decisions on their energy use. ZigBee technology can also convey intelligent communication between the HAN and individual appliances to facilitate load control and demand response. For example, a dishwasher can delay its start to a time of day that has lower utility rates.

Regulation Helps Drive Requirements

Regulations and requirements that are being enacted to reduce energy consumption and ensure a more reliable energy grid are helping push the adoption of green technologies. These regulations and requirements are being driven at the federal, regional and state level. For instance, the Energy Policy act of 2005 and the Energy Independence and Security Act of 2007 mandate standards for items such as smart metering, grid reliability and energy efficiency in both buildings and appliances. Regional drivers, such as the Electric Reliability Council of Texas, have incorporated requirements for smart grid functionality as well as smart meters. The council has recommended using a HAN for rate recovery with a standards-based approach, "…such as ZigBee."

We can see even more stringent state-level drivers, with California leading the way with its California Load Control Proceeding (formally Title 24). Utility companies, such as Southern California Edison, are reacting by rolling out five million ZigBee-enabled meters and HANs over the next several years[4].

The Savings Add Up

Much of the recent activity targets residential customers. And while it may seem that this is benefiting the utility company more, it also benefits the home owner with reduced consumption and operating expense. In addition, by enabling a more intelligent home, consumers can begin automating such things as lighting control. Lights that are not needed or are accidentally left on can be turned off. Answering a demand/response notification, lights can be automatically dimmed and starting other appliances, such as a dishwasher, oven, coffeemaker, etc. can be delayed. The below example looks at a simple lighting automation for the HAN that shows significant savings.

- 2,000 Square foot home with 4,500 watts in lighting
  - 3 bedrooms x 440 watts
  - 2 bathrooms x 460 watts
  - 1 kitchen x 460 watts
  - 1 dining room x 240 watts
  - 2 living rooms use 340 watts
  - Other has 880 watts (closets, garage, half bath, outdoor lights, etc.)
- Typical lights are bathroom, bedroom, closet, or outdoor lighting
- $246 annual cost
  - 450 watts x 10 hours/1,000 (kilowatts) x .15 (cost per kWh) x 365 (days a year)

Residential Smart Energy Example

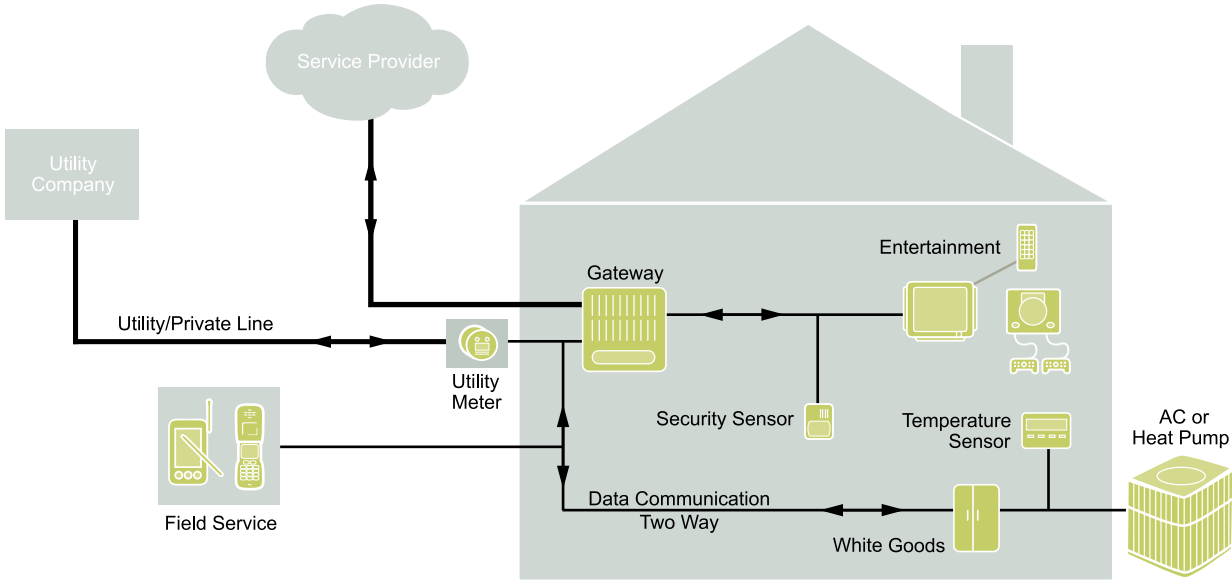

Figure 2

- If these 2-3 lights are controlled by a ZigBee-enabled HAN in a Smart Energy home, these lights might only be on a few minutes at a time, saving a considerable amount of the original $246 cost.

Products are Hitting the Shelves

While ZigBee-enabled products for home automation have been on the market for several years, they have been expensive and remained a niche market. However, with the introduction of the ZigBee Smart Energy Profile and the rollout of Smart Energy enabled meters from the utility companies, we will start seeing more products on the market. Initially the rollout will include a ZigBee-enabled meter, thermostat and premise device to communicate pricing information. But soon we will start to see products such as lighting and appliances that can benefit from the Smart Meters and demand response information sent by the utility. Eventually, we could expect to see the home with multiple ZigBee-enabled devices, such as lighting, HVAC, appliances and even security; all communicating with each other and the utility meter.

Freescale Enables Smart Energy and ZigBee Technology

Freescale has an extensive portfolio of ZigBee-enabled ICs and low power MCUs that make up an ideal platform for ZigBee and Smart Energy products. Freescale's MC1322x ZigBee Platform in a Package (PiP) is the latest in a long line of low-power platforms for ZigBee devices. With its best-in-class power consumption and highly integrated packaging, it is becoming the metering companies' preferred choice to enable Smart Meters with ZigBee technology.

For customers that need additional functionality or memory, Freescale offers a number of low-power MCUs as well that can be paired with stand-alone transceivers. The Flexis™ family of MCUs offers low power and an upgrade path from 8-bit to 32-bit processing power. In addition to ICs, Freescale has taken a platform approach that includes software, development tools and reference designs to help simplify development. Freescale's BeeStack™ ZigBee compliant stack with BeeKit™ Wireless Toolkit provides a simple software environment to configure network parameters. This tool is unique to Freescale, allowing customers to use a wizard and drop down menus to help configure the ZigBee network parameters. Our competitors force the user to wade through lines of code to edit the network parameters.

Finally, while the MC1322x PiP simplifies RF design, many customers don't have the expertise to ensure robust and optimized designs. Freescale has a number of reference designs, from IC to antenna, that include the design details for the development hardware in the development kits as well as designs that are more "form factor." They can take the bills of materials, gerbers and schematics and simply copy our design or integrate it into theirs. The complete platform approach helps you reduce development time and speed time to market.

Freescale development kits provide the ideal development platform for the ZigBee protocol. The MC1321x and MC1322x are specifically targeted for ZigBee-enabled device development, providing the necessary hardware, software, tools and demo applications to streamline the development process. These kits come with development boards programmed to allow developers to have a ZigBee network up and running in just fifteen minutes. The demonstration is based on Freescale's ZeD (ZigBee enabled Demo) using the Home Automation profile. The kit also comes with the BeeKit Wireless Toolkit and additional demo applications. In addition, for customers running the ZigBee protocol that requires a different low-power MCU, they can use the MC13202 RF transceiver and the Flexis QE128 MCU. The MC1320x-QE128-DSK provides a simple and low-cost development platform and includes a 2-node example using the ZeD HA lighting application.

Summary

With all this activity, it appears that ZigBee technology has found a home in the Smart Energy market. While this is just an example of how ZigBee technology is going green in the home, we expect ZigBee to move beyond lighting and HVAC to other appliances and to migrate to commercial buildings. With Freescale's leadership in ZigBee technology development, our focus on Smart Energy and our relationship with customers, we are positioned for success in this rapidly growing market.

Matt Maupin has over 14 years of experience in the high-tech industry focusing on wireless connectivity, including WiFi, Bluetooth and ZigBee. Matt has been active in numerous industry groups for wireless communication and is currently active in the ZigBee Alliance. Matt joined Freescale in 2001.

Pavel Sustek

Cost-Effective Vector Control Drive
Using Single Shunt Current Sensing

Introduction

Persistent demand for energy-saving industrial and home appliances has recently escalated due to energy efficiency, environmental issues and the necessity to comply with new energy consumption regulations. These regulations force the development of energy-efficient motors for appliances, such as washing machines, air conditioner compressor systems and fans. It's estimated that electric motors consume over 70 percent of all electricity in industrial applications, and studies by the Electric Power Research Institute say that over 60 percent of industrial motors are operating under their rated load capacity.

Alternating current AC induction motors (ACIM) are popular in industry and consumer electronics for a number of reasons (See Figure 1). They are very simple (they have no brushes) with favorable manufacturing costs, robust and require minimum maintenance. They have also been produced for years, so their construction is extremely optimized. Traditionally, these motors have been run without speed control, and are started and stopped frequently in order to achieve the desired result. About 50 percent of the electricity used during such a process is wasted.

Many new methods of reducing electricity in ACIMs are being considered, including new electric motor efficiency technologies. System costs and power consumption can be drastically reduced utilizing digital control of an analog motor circuit. This article describes a 3-phase AC induction motor vector drive solution based on Freescale's MC56F8013/23 digital signal controller (DSC) that takes advantage of a cost-efficient solution for consumer and industrial motor drives.

Motor Overview

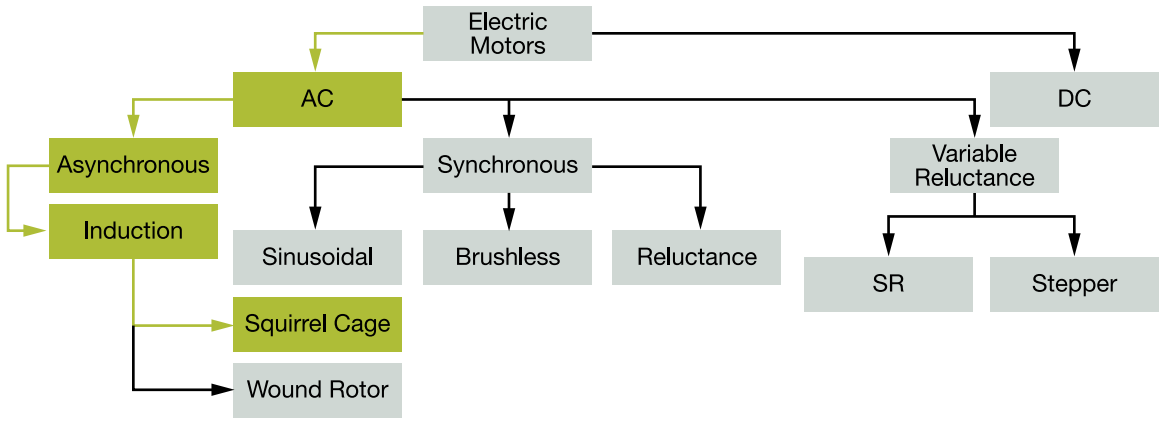

Figure 1

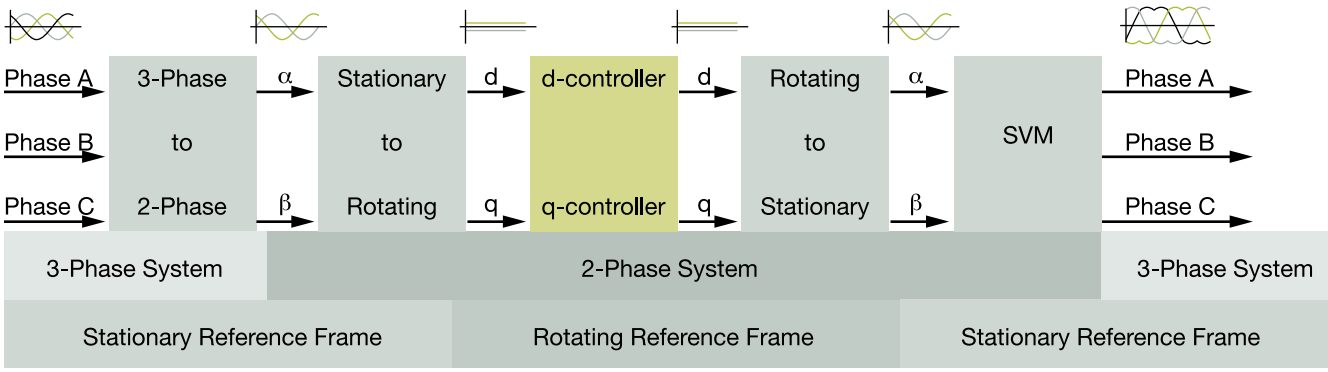

Figure 2

Three-Phase AC Induction Motor

The ACIM is a rotating electric machine designed to operate from a 3-phase alternating voltage source. Slots in the inner periphery of the stator each accommodate a 3-phase winding. The turns in each winding are distributed so that a current in a stator winding produces an approximately sinusoidally-distributed flux density around the periphery of the air gap. When three currents that are sinusoidally varying in time, but displaced in phase by 120 degrees, flow through the three symmetrically-placed windings, a radially-directed air gap flux density is produced that is also sinusoidally distributed around the gap and rotates at an angular velocity equal to the angular frequency of the stator currents.

The most common type of induction motor has a squirrel cage rotor in which aluminum conductors, or bars, are cast into slots in the outer periphery of the rotor. These are shorted together at both ends of the rotor by cast aluminum end rings, which also can be shaped to act as fans. As the sinusoidally-distributed flux density wave produced by the stator magnetizing currents sweeps past the rotor conductors, it generates a voltage in them. The result is a sinusoidally-distributed set of currents in the short-circuited rotor bars. Because of the low resistance of these shorted bars, only a small relative angular velocity between the angular velocity of the flux wave and the mechanical angular velocity of the two-pole rotor is required to produce the necessary rotor current. The relative angular velocity is called the slip velocity. The interaction of the sinusoidally-distributed air gap flux density and induced rotor currents produces a torque on the rotor.

Vector Control of AC Induction Motor

In order to achieve variable speed operations in a three-phase AC induction motor, variable voltage and variable frequency need to be supplied to the motor. Modern three-phase variable speed drives (VSD) are supplied with digitally controlled switching inverters, which can considerably reduce overall system power consumption. Using a variable speed drive motor can provide electricity savings of up to 60 percent, a three to four times increase of resources and can enable functional possibilities unattainable before. The power range of a variable speed drive can recover 0.2–0.4 kW in refrigerator compressors, 0.8–1 kW in washing machines, 2–3 kW in air conditioners and 3–100 kW in electric drives for housing and communal services (for example, pumps for cold and hot water in many-storied houses, cold water pipelines in trunks, etc.).

The control algorithms can be sorted into two general groups. The first group is scalar control. The constant volt per hertz control is a very popular control technique. The other group is called vector or field oriented control (FOC). The vector oriented techniques provide improved drive performance over scalar control. FOC advantages include higher efficiency, full torque control, decoupled control of flux and torque, improved dynamics and more.

The basic idea of the FOC algorithm is to decompose a stator current into flux and torque producing components. Both components can be controlled separately after decomposition. The structure of the motor controller is then as simple as that for a separately excited DC motor. Figure 2 shows the basic structure of the vector control algorithm for the AC induction motor.

To perform vector control, it is necessary to follow these steps:

- Measure the motor quantities (phase voltages and currents)
- Transform them into the 2-phase system (α,β) using a Clarke transformation

- Calculate the rotor flux space-vector magnitude and position angle
- Transform stator currents into the d-q reference frame using a Park transformation
- The stator current torque (isq) and flux (isd) producing components are separately controlled
- The output stator voltage space vector is calculated using the decoupling block
- The stator voltage space vector is transformed by an inverse Park transformation back from the d-q reference frame into the 2-phase system fixed with the stator
- Using space vector modulation (SVM), the output 3-phase voltage is generated

To be able to decompose currents into torque and flux producing components (isd, isq), we need to know the position of the motor magnetizing flux. This requires accurate velocity information sensed by a speed or position sensor attached to the rotor. Incremental encoders or resolvers are used as position transducers for vector control drives. In cost sensitive applications, such as washing machines, tachogenerators are widely used. In some applications, however, the use of speed/position sensors is not desirable. Then, the aim is not to measure the speed/position directly, but to instead employ some indirect techniques to estimate the rotor position. Algorithms which do not employ speed sensors are called "sensorless control."

Description of Vector Control Algorithm

The overview block diagram of the implemented control algorithm is illustrated in Figure 3. As with other vector control oriented techniques, it is able to control the excitation and torque of the induction motor separately. The aim of control is to regulate the motor speed, and the speed command value is set by high level control. The algorithm is executed in two control loops. The fast inner control loop is executed with a 125 µs period. The slow outer control loop is executed with a period of one millisecond.

To achieve induction motor control, the algorithm utilizes a set of feedback signals. The essential feedback signals are DC-bus voltage, three-phase stator current reconstructed from the DC-bus current and motor speed. For correct operation, the presented control structure requires a speed sensor on the motor shaft. In the case of the presented algorithm, an incremental encoder is used.

The fast control loop executes two independent current control loops. They are the direct and quadrature-axis current (isd, isq) PI controllers. The direct-axis current (isd) is used to control rotor magnetizing flux. The quadrature-axis current (isq)

Digital Signal Controller 56F8xxx

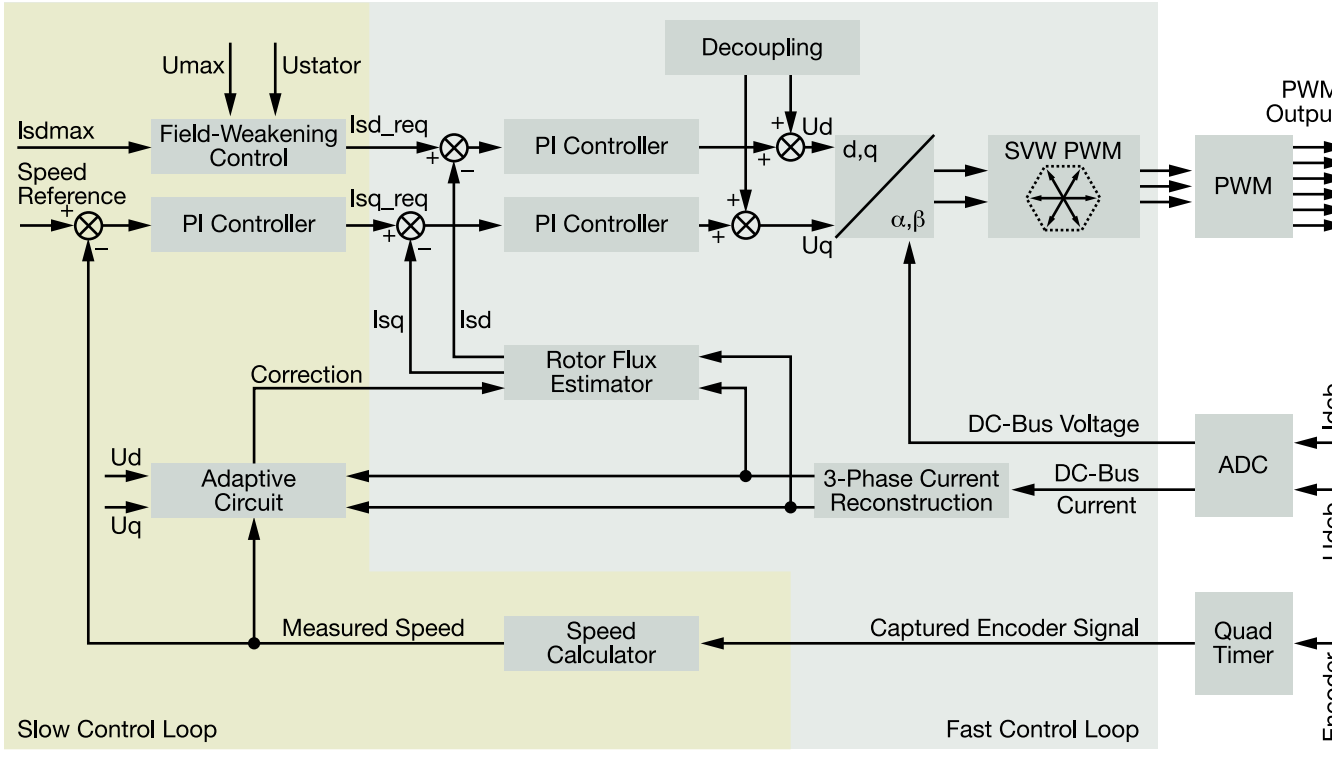

Figure 3

corresponds to the motor torque. The current PI controllers' outputs are summed with the corresponding d and q axis components of the decoupling stator voltage. Thus we obtain the desired space-vector for the stator voltage, which is applied to the motor. The fast control loop executes all the necessary tasks to enable an independent control of the stator current components. These include:

- Three-phase current reconstruction
- Forward Clark transformation
- Forward and backward Park transformations
- Rotor magnetizing flux position evaluation
- DC-bus voltage ripple elimination
- Space vector modulation (SVM)

The slow control loop executes speed and field-weakening controllers and lower priority control tasks. The PI speed controller output sets a reference for the torque producing quadrature axis component of the stator current (isq). The reference for the flux producing direct axis component of the stator current (isd) is set by the field-weakening controller. The adaptive circuit performs correction on the rotor time constant to minimize the error of the rotor flux position estimation.

System Concept

The Freescale MC56F80xx family is well suited for digital motor control, combining the DSP's calculation capability with the MCU's controller features on a single chip.

The MC56F80xx family members provide these peripheral blocks:

- One pulse width modulation (PWM) module with PWM outputs, fault inputs, fault-tolerant design with dead-time insertion, supporting both center-aligned and edge-aligned modes
- 12-bit analog-to-digital converters (ADCs), supporting two simultaneous conversions; ADC and PWM modules can be synchronized
- One dedicated 16-bit general-purpose quad timer module
- One serial peripheral interface (SPI)
- One serial communications interface (SCI) with Local Interconnect Network (LIN) slave functions
- One inter-integrated circuit (I^2C) port
- On-board 3.3 V to 2.5 V voltage regulator for powering internal logic and memories
- Integrated power-on reset and low-voltage interrupt module
- All signal pins multiplexed with general-purpose input/output (GPIO) pins
- Computer operating properly (COP) watchdog timer

System Concept: Three-Phase Control Board

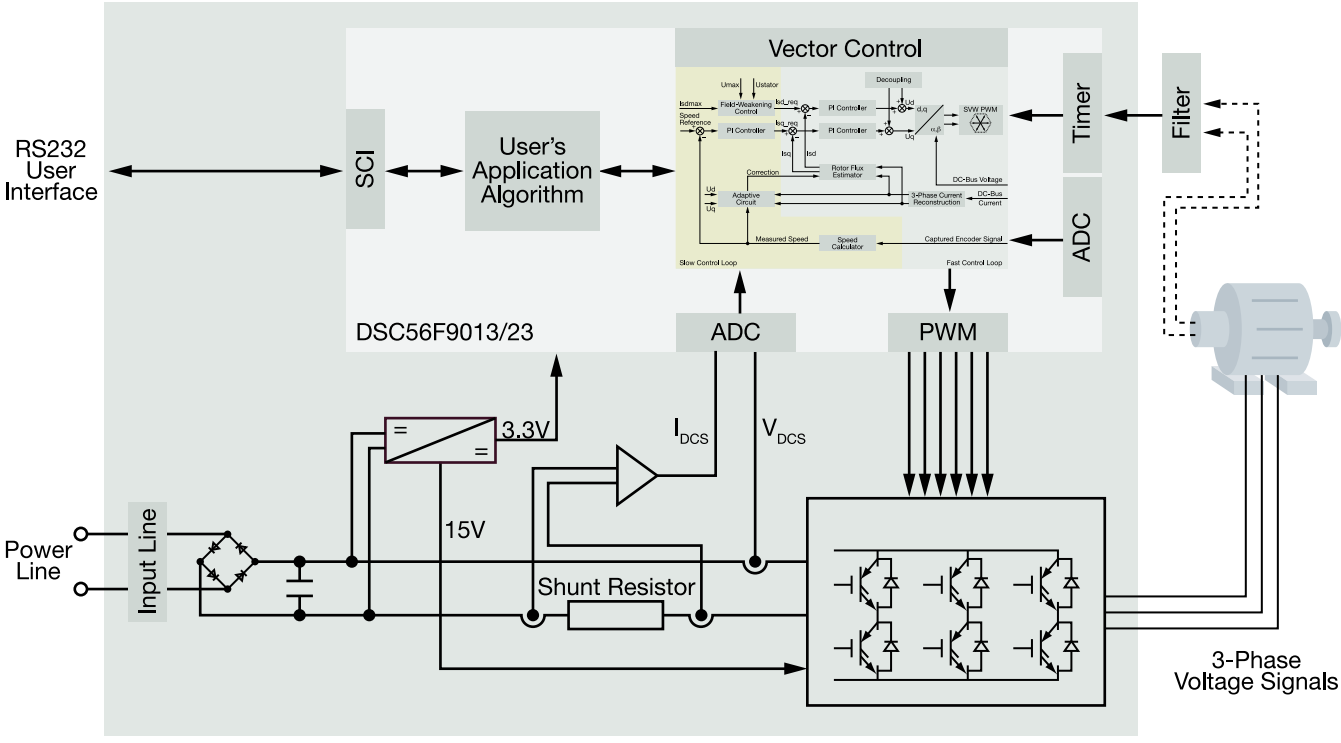

Figure 4

- External reset input pin for hardware reset (may also be assigned as GPIO)
- JTAG/On-Chip Emulation (OnCE™) module for unobtrusive, processor-speed-independent debugging
- Phase-locked loop (PLL) based frequency synthesizer for the hybrid controller core clock, with on-chip relaxation oscillator

The three-phase ACIM vector control with single shunt sensor greatly benefits from the flexible PWM module, fast ADC and quad timer module. The PWM offers flexibility in its configuration, enabling efficient three-phase motor control. The PWM module is capable of generating asymmetric PWM duty cycles in center-aligned configuration. We can benefit from this feature to achieve a reconstruction of three-phase currents in critical switching patterns. The PWM reload SYNC signal is generated to provide synchronization with other modules (Quadtimers, ADC). The application uses the ADC block in simultaneous mode scan and is synchronized to the PWM pulses. This configuration allows the simultaneous conversion of the required analogue values for the DC-bus current and voltage within the required time. The ADC conversions are triggered directly by the PWM without need for the DSC core to relay the event, resulting in predictable and constant relative timing.

The quad timer is an extremely flexible module, providing all required services relating to time events. The application uses four channels:

- One channel for PWM-to-ADC synchronization
- Two channels for reading quadrature encoder signals (one channel in case the tachogenerator is used instead of a quadrature encoder)
- One channel for system base of slow control loop (1 ms period)

An adaptive closed loop rotor flux estimator enhances control performance and increases the overall robustness of the system. Parameter drift sensitivity can be considerably minimized in this way. Minimizing system cost, the algorithm implements single shunt current sensing, which reduces three current sensors to one.

DC-Link Current Sensor

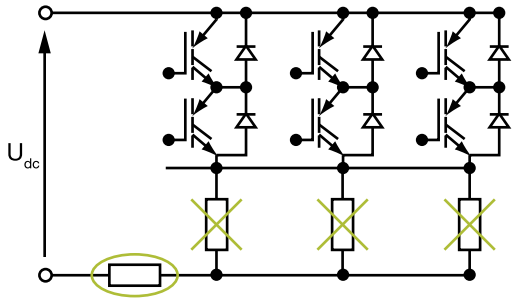

Figure 5

Another advantage is the high range of motor operating speeds, up to 20,000 rpm. Washing machines, for instance, require this high speed. The ratio between motor and drum speed is about ten to one in horizontal washers. Thus, to achieve a drum speed of 2,000 rpm, the motor has to run at 20,000 rpm. Three-phase induction motors for washers are designed for nominal speed that is much lower than 20,000 rpm (usually 6,000 rpm). The higher speed is achieved utilizing a field-weakening algorithm, which allows the motor to exceed the nominal speed while the magnetizing flux remains at the nominal motor voltage. Using a motor designed for lower nominal speed but can run up to 20,000 rpm with a field-weakening algorithm can provide significant cost and energy savings.

Three-Phase Current Reconstruction

The vector control algorithm requires the sensing of the three motor phase currents. A standard approach is to sense the phase currents directly through current transformers, or Hall effect sensors, directly coupled to the motor phase lines that carry the current between the switches and the motor. To reduce the number of current sensors and overall cost of the design, the three-phase stator currents are measured by means of a single DC-link current shunt sensor. (See Figure 5)

The DC-link current pulses are sampled at exactly timed intervals. A voltage drop on the shunt resistor is amplified by an operational amplifier inside the 3-phase driver and shifted up by 1.65V. The resultant voltage is converted by the ADC.

The stator's three-phase currents are reconstructed based on the actual combination of switches. The ADC measures the DC-link current during the active vectors of the PWM cycle. When the voltage vector V1 is applied, current flows from the positive rail into the phase A winding and returns to the negative rail through the B and C phase windings. When the voltage vector V2 is applied, the DC-link current returning to the negative rail equals the T phase current. Therefore, in each sector, two phase current measurements are available. The calculation of the third phase current value is possible because the three winding currents sum is zero. The voltage vector combination and corresponding reconstructed motor phase currents are shown in Table 1.

Measured Current	
Voltage Vector	**DC-Link Current**
V1(100)	+ia
V2(110)	-ic
V3(010)	+ib
V4(011)	-ia
V5(001)	+ic
V6(101)	-ib
V7(111)	0
V0(000)	0

Table 1

An Example of Estimating Fan Energy Cost

Pumps and fan systems account for almost 40 percent of all motor applications in the industry. The input power of a fan is proportional to cube of the flow. For example, if 100 percent flow requires full power, 75 percent flow theoretically requires $(0.75)^3 = 42$ percent of full power. Although this is the theoretical saving under zero static head conditions, even in practical applications, a substantial energy saving can be achieved. VSDs are most often used to save energy in centrifugal fans and pumps. When estimating the savings from such applications, it is useful to use the Fan Laws, which relate air (fluid) flow, input power and motor speed.

The equation forms for Fan Laws

$$Q_2 / Q_1 = N_2 / N_1$$
$$P_2 / P_1 = (N_2 / N_1)^2$$
$$HP_2 / HP_1 = (N_2 / N_1)^3$$

Where:

Q_1, Q_2—initial and new volumetric flow rate (l/s)
P_1, P_2—initial and new pressure or head (kPa)
N_1, N_2—initial and new fan speed (rpm)
HP_1, HP_2—initial and new input fan power (kW)

A centralized plant supplying a chilled water system with two flow requirements is a good example of how Fan Laws may be applied. The system is controlled by an electronic actuator that throttles the water flow based on the system requirements. With the throttle fully open, the flow rate is 80 liters per second (l/s) while power consumption has been measured at 20 kW. When the flow rate is reduced to 65 l/s using a throttling valve the power consumption is measured at 18 kW. The system operates for 8,760 hours per year divided into 60 percent and 40 percent for the throttled and un-throttled conditions, respectively. By installing a VSD, savings can be calculated as follows.

$$HP_2 / HP_1 = (N_2 / N_1)^3$$
$$\text{As } HP_2 / HP_1 = (Q_2 / Q_1)^3$$
$$HP_2 / 20 \text{ [kW]} = (65 \text{ [l/s]} / 80 \text{ [l/s]})^3$$
$$HP_2 = (65 \text{ [l/s]} / 80 \text{ [l/s]})^3 * 20 \text{ [kW]}$$
$$HP_2 = 10.77 \text{ kW}$$

This shows that for the 60 percent of the time that the system is throttled, the savings would be in the order of:

8760 hours/year * .60 * (18 kW - 10.77 kW) = **38,001 kWh**

VSD should be used when:

- Production volume fluctuates
- Multi-speed motors are currently used
- Dampers, control valves or recycle loops are used to control flow
- Very accurate speed and torque control is required
- Speed control systems do not perform satisfactorily
- Systems operate more than 80 hours per week

Conclusion

The solution presented based on of Freescale's MC56F80xx DSC shows cost-effective design for a wide range of both industrial and consumer motor control application. DSCs completely control the AC induction motor so it always functions at optimum efficiency. While the capacity of the motor corresponds to the process needs, variable speed drives can provide major savings compared to the wasteful practice of running the motor at full speed. This reduces stress and strain on the motor and extends motor life, both electrically and mechanically. With energy prices on the rise, Freescale DSC's introduce a smart option for energy savings on electric motors.

Pavel Sustek is an application engineer on the motor control team at Freescale and has been with the company for almost four years. He has a master's degree in electrical engineering.

Petr Frgal

Digital Power Factor Solution from Freescale
DSC Capable to Control PFC and Motor

Digital PFC and Motor Control Based on the 56F8xxx DSC Family as One Solution

When standard IEC31000-3-2 became mandatory in 2001, many companies began considering power factor correction (PFC) in their designs, including products such as lighting equipment, portable tools, all electronic equipment, consumer products, appliances and industrial equipment. This standard deals with the limitations of harmonic currents injected into the public mains supply system. It is applicable to electrical and electronic equipment having an input current up to and including 16A per phase that is intended to be connected to public low-voltage distribution systems.

Without using PFC, a typical switched mode power supply would have a power factor of around 0.6; therefore having considerable odd-order harmonic distortion (sometimes with the third harmonic as large as the fundamental). Having a power factor of less than one along with harmonics from peaky loads reduces the real power available to run the device. In order to operate a device with these inefficiencies, the power company must supply additional power to make up for the loss. This increase in power causes the power companies to use heavier supply lines or suffer burnout in the neutral line conductor. PFC use is increasing every year, especially in highly competitive enviroments with many companies sharing the market.

There are two main PFC classifications—active and passive. Passive PFC is less expensive but is hard to design for variable input voltages and variable loads. A passive PFC circuit is simple, with fewer components than active PFC. It is good for low-power applications below 200W and with one input voltage, for which inductors and capacitors are small and inexpensive.

For higher power applications, when inductors and capacitors are bigger and more expensive, active PFC is a better and more cost-effective approach. Passive PFC is suitable for low-power, fixed-voltage and fixed-load applications. Active PFC is used in medium and high power demand applications, such as PC power supplies, UPS, telecom equipments and plasma displays where passive PFC cannot meet the system requirements (SMPS, HF ballasts, converters, battery chargers, etc.).

Two other approaches are analog and digital PFC. Conventional analog PFC controllers offer many control algorithms and fewer additional components. Digital PFC can offer comparable results, but the more dynamic PFC can offer better performance and easier modification.

This article describes average current mode control for PFC on Freescale's MC56F8013 digital signal controller (DSC). There are various reasons to use PFC other than saving power and meeting current regulations. PFC can mitigate harmonic distortion, which can cause increased operating temperatures in power generating equipment. Higher temperatures can reduce the operating life of such equipment as rotating machines, cables, transformers, capacitors, fuses, switching contacts and surge suppressors. Problems are caused by the harmonics creating additional dielectric stresses in capacitors and cables, increasing currents in machinery windings and transformers and noise emissions in many products. They also can cause a skin effect, which creates problems in cables, transformers and rotating machinery. All these factors influence the reliability, performance and aging of the electrical equipment.

Although the application described in this paper is for the MC56F8013 DSC, it can be ported into the other members of the MC56F80xx family according to the application

requirements. The presented implementation represents a fully digital solution. Fast current and slow voltage loops are implemented digitally using the DSC, and the PFC power switch is controlled directly by the DSC PWM output. Therefore, it is called direct PFC. The direct PFC algorithm works in an average current control continuous conduction mode (CCM).

Using the direct control approach requires more DSC resources than the indirect solution, where the PWM is generated by external hardware circuitry. On the other hand, with direct control it is possible to generate a pure current sine wave drawn from the line and to obtain the ideal ohmic load character at the input. Another advantage of direct PFC is its constant transistor switching frequency which reduces noise.

You can achieve better system dynamics using direct PFC because the control algorithm is simple and fast. Also, you do not need a synchronization signal from the line voltage and fewer passive components are needed compared to indirect PFC. The solution is cost-effective and is suitable for medium (200–600W) and higher power (above 600W) applications. The DSC's high performance enables the concurrent operation of the PFC and motor control applications. This paper describes a PFC implementation into a 3-phase AC induction vector control drive with single shunt current sensing.

A boost converter is widely used as an active power factor correction pre-regulator. The control structure is divided into two loops: an inner current control loop and an outer voltage control loop, as shown in Figure 1. The outer voltage control loop is implemented via software in the DSC and keeps a constant voltage on the DC bus. The voltage control loop utilizes a propotional-integral (PI) controller, and the output defines the amplitude required for the PFC current. The PFC control algorithm provides a sinusoidal input current without phase shift to the input voltage through dedicated PFC hardware controlled by the DSC. The hardware incorporates an input bridge rectifier DB, PFC inductance L, PFC diode D and PFC switch Q. These analog quantities are sensed-rectified input voltage, input current and DC-bus voltage. The input current is controlled using the PFC switch to achieve the desired input current and the desired level of the DC-bus voltage (UREQ).

The inner current loop is implemented via software as is the outer loop, and it employs the PI controller to maintain the sinusoidal input current by directly controlling the PFC transistor. The input to the PI controller is the difference between the current reference, IREQ, and the actual current, IL. The sinusoidal waveform of IREQ is derived from the shape of the input voltage UDC RECT, as shown in Figure 1. The final current reference, IREQ, is acquired by multiplying a rectified input voltage waveform by the output of the voltage controller. The current PI controller's output generates a signal, D, corresponding to the duty cycle of a boost converter in an open loop. The bandwidth of the current PI controller has to be set

above 8 kHz to get a sufficient response. Therefore, the current PI controller algorithm has to be executed at least once every 60 µs, which puts a lower-limit requirement on the performance of the DSC. The DSC performance requirement for the voltage control loop is low, as the bandwidth of the voltage control loop is set below 20 Hz. Therefore, the DSC performance is not a limiting factor in this part of the PFC algorithm.

The dedicated PFC hardware is designed as part of the entire system. The PFC board together with the power stage and controller board form one compact system to drive a 3-phase AC/BLDC motor, including PFC control.

The application meets the following performance specifications:

- Hardware
  - MC56F8013/23 controller board
  - PFC board
  - 3-phase AC/BLDC high voltage power stage board
- Control technique
  - Inner current loop
  - Outer voltage loop
  - Current reference generation
  - RMS input voltage calculation
- FreeMASTER software monitor
- Fault protection
  - DC-bus over-voltage and under-voltage
  - Over-current protection
  - Input voltage over-voltage and under-voltage

The power factor correction application provides the sinusoidal input current by controlling the PFC switch. In the control loop, the actual DC-bus voltage is compared with the desired voltage. The control error is processed by the PI controller, which generates the amplitude of the reference current. Input rectified voltage is multiplied by the input rectified RMS voltage and by the output of the voltage controller. This multiplied value is the reference current, which is compared to the actual current sensed on the shunt resistor. The difference between them is

System Concept Block Diagram

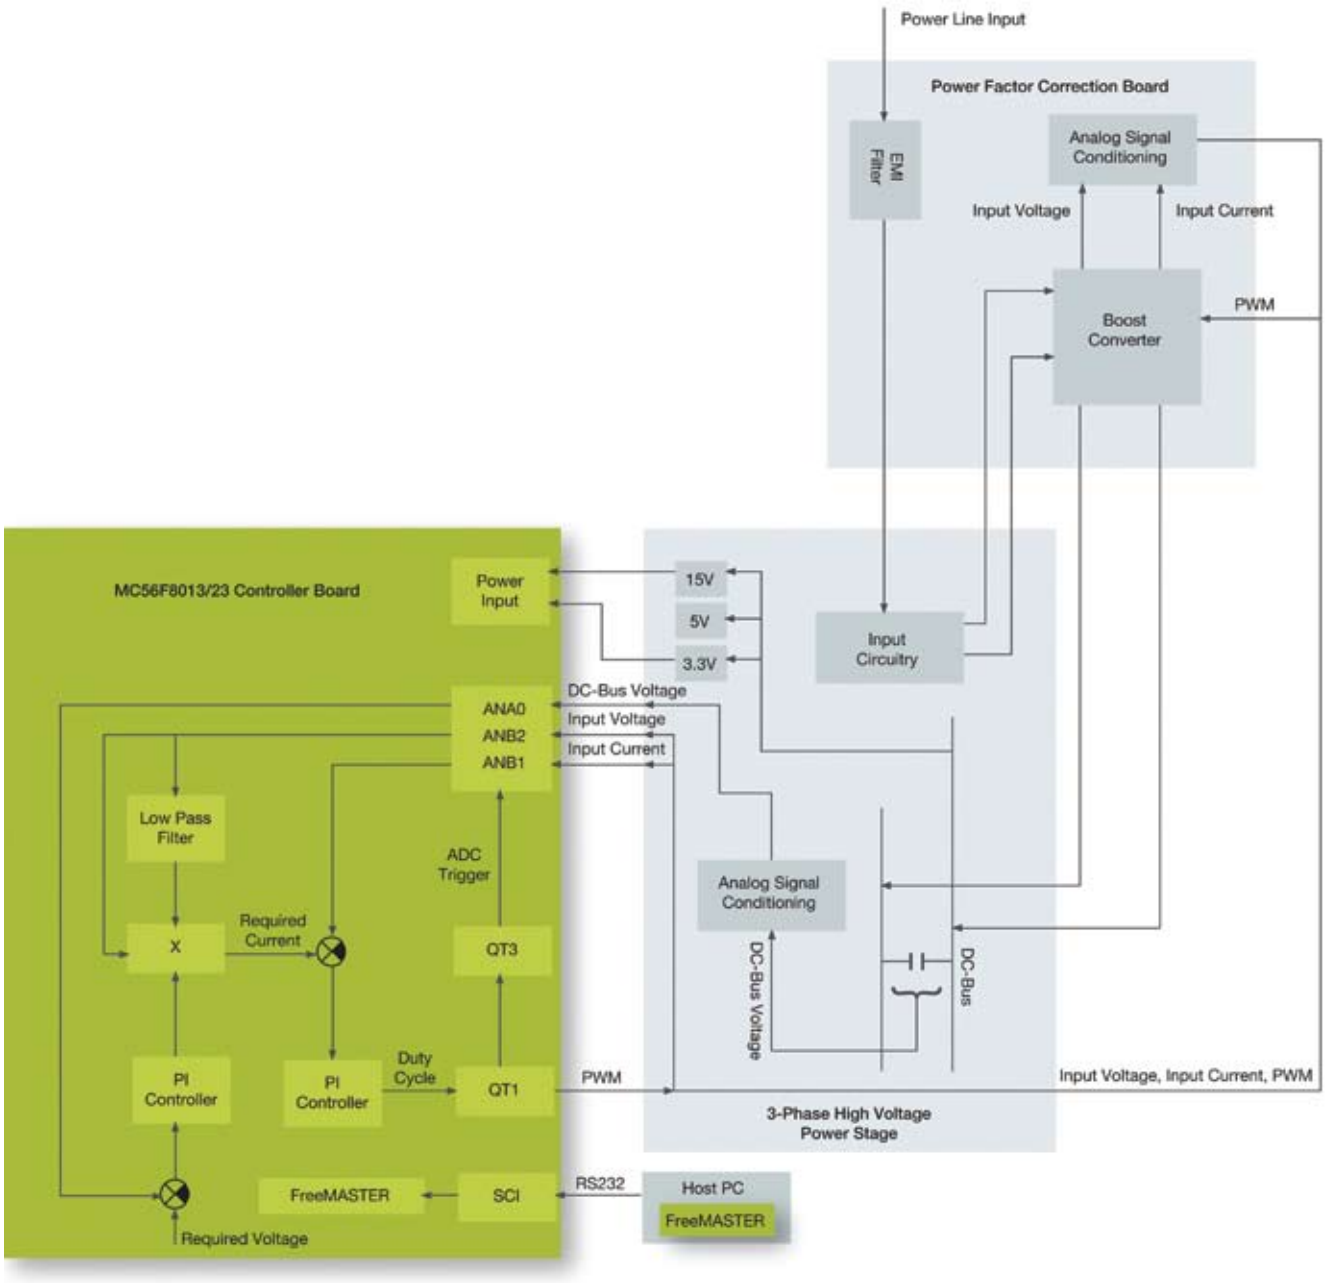

Figure 1

then processed in the PI current controller. The output from this controller is the PWM signal for quad timer 1, which directly switches the PFC transistor.

The whole application is controlled by the MC56F8013 DSC. This low-cost digital signal controller offers many important features and peripherals for this kind of application. The quad timer and analog-to-digital converter (ADC) are the most important peripherals used in this application. The ADC is used for sensing analogue quantities and the quad timer for timing the control algorithm, ADC sampling synchronization and control signal generation.

The entire PFC algorithm is implemented in one interrupt routine generated from the quadrature timer (QT1). This routine is called every 31.25 µs, corresponding to 32 kHz. Such a frequency is high enough to generate the proper current shape but doesn't load the DSC core more than necessary. The current loop is executed on every interrupt. The current controller utilizes a recursive algorithm for fast execution time. Quad timer, channel 2, is used as a synchronization signal between A/D converter B and the PWM signal generated by QT1. The A/D conveter B synchronization signal (scan start) is executed at the midpoint of the on-time PWM signal to measure an average inductor current. The ADC reads the input current and input voltage in one sequence using A/D converter B, and in another sequence,

converts the output voltage using A/D converter A. Input current is sensed every QT1 interrupt, 31.25 µs, and input voltage and output voltage are sensed every fourth QT1 interrupt, 125 µs.

All quantities are sensed at the midpoint of the duty. RMS input voltage is gained by filtering the input voltage by a 10 Hz low-pass filter. The switching frequency of the PFC transistor is set to 32 kHz. This constant switching frequency of the PFC transistor simplifies the design of the input filter. The result of the current controller defines the duty cycle of the PFC transistor.

PFC can be incorporated as a part of every system supplied from AC power. The system concept for implementing a PFC into a 3-phase AC induction vector control drive with single shunt current sensing is illustrated in Figure 3. The motor control application uses an AC induction motor with a tachogenerator. If you would like to use a motor with an encoder sensor, the MC56F8013 device does not have enough timer channels for encoder processing. You would need to use a device with more timers, such as the MC56F8037 DSC. Porting of the presented MC56F8013 software to the MC56F8037 is straightforward.

The system consists of three boards: power factor correction board, 3-phase AC/BLDC high voltage power stage board and the MC56F8013/23 controller board. Hardware and software implementations are described below. There is no need for any additional hardware to implement control of the ACIM. Configuration is the same as for standalone PFC purposes. A suitable motor must be connected.

The application software is interrupt-driven running in real time. There are three periodic interrupt service routines executing the major motor control and PFC tasks.

- The Timer 3 interrupt service routine performs a fast current control loop and PFC tasks. It is executed when a third DC-bus current sample is read, with a 125µs period.
- The PWM reload interrupt service routine performs a fast current control loop and PFC tasks. It is executed every PWM half-reload, with a 31.25 µs period.
- The ADC channel A End of Scan interrupt service routine reads the DC-bus current samples. It is executed for three consecutive sample readings within one PWM cycle.

There is also the non-periodical interrupt service routine. The PWM fault interrupt service routine is executed on an over-current event to manage an over-current fault condition. It is executed only if the fault condition occurs.

The background loop is executed in the application power line. It manages non-critical time tasks, such as the application state machine and FreeMASTER communication polling.

The PWM module is configured to run in center-aligned mode. The PWM_half_reload_sync signal is generated every PWM half-cycle with a 31.25 µs period. The PWM_half_reload_sync is connected to the timer module. An output from timer channel 3 is used to trigger the ADC channel A. A connection link between the PWM module, TMR module, and the ADC module enables the definition of the exact multiple time instants of ADC sampling, which are synchronized to the generated PWM signal.

ADC channel A is started for the third time after 31.25 µs and the third DC-bus current sample is read. Simultaneously, the Timer 3 interrupt is executed. After the third current sample has been read, the Timer 3 ISR is interrupted and the ADC channel A End of Scan ISR is executed. When this ISR has finished, the Timer 3 ISR continues processing.

The fast current control loop is executed in the PWM reload ISR, which is synchronized to the PWM_half_reload_sync signal. Prior to the PWM reload ISR being executed, three ADC samples of the DC-bus current are taken and processed by the ADC channel A End of Scan ISR.

A measurement was made for standalone PFC at full load, which met international regulation IEC61000-3-2 requirements. Total harmonic distortion (THD) is 4.5 percent and power factor (PF) 0.99.

PFC components are designed for the same rated power as the power stage board, which is 750W. The PFC was also tested together with the AC induction motor as a load. Current limits were also met, but THD has a higher value and PF has a lower value. This is due to the fact that PFC was designed for 750W. For better results, components should be redesigned for the rated power of the chosen motor.

In this particular implementation (3-phase AC induction vector control drive with single shunt current sensing) we have used a tachogenerator for speed sensing. Implementing a design with the MC56F8037 DSC will eliminate the need for a tachogenerator, thanks to the increased number of timers. It is then possible to use an encoder for speed sensing, which may be required in some cases.

A detailed description of the application is available as a design reference manual DRM098 on the Freescale web site **http://www.freescale.com/files/microcontrollers/doc/ref_manual/DRM098.pdf**.

Petr Frgal graduated from VSB Technical University of Ostrava with a BS in Electronic Engineering Technology. Since 2004, he has worked as an Application Engineer in Roznov, Czech Republic. He lives in the small city of Hranice, where he spends his free time with his family and participates in sports such as volleyball and squash.

Petr Cholasta

Remote Keyless Entry
In a Body Controller Unit Application

Many of us know this situation. When we leave the car, with a single click of a remote control we lock and secure it until we return. This is a significant safety and convenience improvement over manually locking your car (and sometimes locking the keys inside the car). The number of electronic systems within the car is increasing rapidly, offering engineering opportunities to rework and simplify common systems, develop new ones and make all of them smarter, safer, and more secure and eco-friendly than ever before. This article introduces a concept about how new technology can be used within the body controller unit (BCU) vehicle system.

First let's take a closer look at three new technology drivers:

- High-performance—enabling complex system control
- Lower power consumption—helping save the environment and resources for the next generation
- Ease of use—a high level of integration that enables engineers to build high-quality systems that users can quickly understand

Each bullet point above specifies just one key feature of a modern system. Although they seem separate, experience has proven that when building powerful systems, each of those bullet points becomes an essential part of the design.

Let's explore how Freescale automotive products fit these categories.

MPC5516 Overview

The MPC5516 is a member of the 32-bit MPC551x family of microcontrollers (MCUs) qualified for automotive applications. It provides a rich feature set balanced with smart development tools. Designers can benefit from dual-core computing power supported by a variety of embedded peripherals (DSPI, IIC, eSCI, FlexCAN, FlexRay™, timers, eQADC, etc.), 1 MB flash memory and 64 kBytes of RAM.

A number of advanced low-power techniques offer power-saving features to minimize system power consumption. The most impressive low-power mode, with RAM partly on, is called Sleep 2. During this mode, the MCU core consumes about 60 µA, maintaining the option to start running on either an internal timers overflow or an I/O pin level/edge. When the MCU wake-up procedure has finished, the user can choose where the program counter starts to count as defined in the reset recovery pointer register (i.e., which part of the application code will be processed immediately after wake-up).

The MPC5516 MCU is offered in LQFP and BGA packages.

MC33696 Overview

The MC33696 is a phase locked loop (PLL) tuned UHF transceiver designed for 304, 315, 426, 434, 868 and 915 MHz industrial, scientific and medical (ISM) bands.

The receiver module is based on superheterodyne architecture containing an Rx data manager with an incorporated module for Manchester encoded (known also as phase encoded, PE) data decoding. The receiver includes a signal strength measurement unit that is used, for instance, for LNA gain control. It also has a strobe oscillator that is extensively used when receiver power supply consumption needs to be reduced but the receiver needs to be kept ready.

The transmitter Tx data manager modulates coded data by ON/OFF or FSK keying. The frequency synthesis module consists of a PLL driven local oscillator and a crystal oscillator that provides the reference frequency.

The device is controlled via the SPI bus. During wireless message reception, the MC33696 behaves as an SPI bus master, however during transmission the SPI module is configured as an SPI slave. The SCLK and MOSI signals carry message data between the MCU and the MC33696.

MC33742 Overview

The MC33742 system basis chip (SBC) is a single-chip solution integrating the frequently used blocks of automotive systems in one package to reduce board space. The SBC incorporates two voltage regulators, VDD and V2, and a CAN transceiver, watchdog and low-power management module. The +5V VDD voltage regulator is fully protected and has been designed to deliver currents of up to 200 mA. The +5V V2 tracking regulator uses an external PNP transistor, thus the current capability is driven by transistor parameters.

The high-speed CAN transceiver can transmit data at a baud rate of up to 1 MBd. The CAN cell is fully protected with an option to wakeup the SBC out of the stop and sleep modes on a CAN message reception. An internal watchdog can be configured either to a window or to a time-out operation. The overflow period can be selected in four steps, from 10 ms to 350 ms. The watchdog can be disabled when running the device in the special debug mode.

The MC33742 can be configured to normal, standby, stop and sleep modes. An internally switched high-side output with four wake-up inputs gives the designer the opportunity to wakeup the system on an external event, such as one initiated through the keyboard. Eight control and status registers determine the SBC behavior. All of them are configurable by the SPI bus.

How Does It All Fit Together In the RKE Application?

All three Freescale devices are excellent solutions for remote keyless entry (RKE) applications. The MPC5516 MCU offers enough computing power for complete system control and message processing. It includes, among other features, various low power modes options and the FlexCAN peripheral for establishing CAN connectivity.

The MC33742 SBC maintains the system +5V power supply, and an integrated CAN transceiver carries messages through the CAN bus. The various SBC low-power modes are balanced by smart system wake-up features. Together, these features make the MC33742 device an essential part in the RKE application.

The UHF band wireless connection is established by applying the MC33696 transceiver, an integrated solution dedicated to wireless messages transmission.

Generally, there are two key system circuitry topologies. System 1 (Figure 1) achieves in low-power mode the lowest power consumption possible (see Table 1), however additional components are required. Conversely, System 2 (Figure 2) incorporates a lower number of essential components, enabling a system wake-up, but the power consumption in low-power mode is higher in comparison to System 1 (see Table 1).

System 1 and System 2 Initialization

When the system is powered up, the MC33742 +5V regulator VDD is turned on and the SBC enters normal request mode. The MPC5516 is now supplied and behaves as it would after a reset. The MCU configures the MC33742 watchdog via SPI 1. Next, the SBC enters normal mode and an initialization is processed. The SBC watchdog is now periodically triggered by the MCU. The MC33742 turns on an additional +5V regulator V2 to supply other system devices typical in the BCU area, such as the eXtreme switches, COSS and MSDI, and the MCU processes the BCU application configuration.

The system configuration is finished by placing the MC33696 in receive mode. Now the BCU is ready to perform the various tasks defined by the application designer.

Let's take a closer look to see how System 1 and System 2 behave when a driver either locks and secures a car or unlocks and unsecures a car.

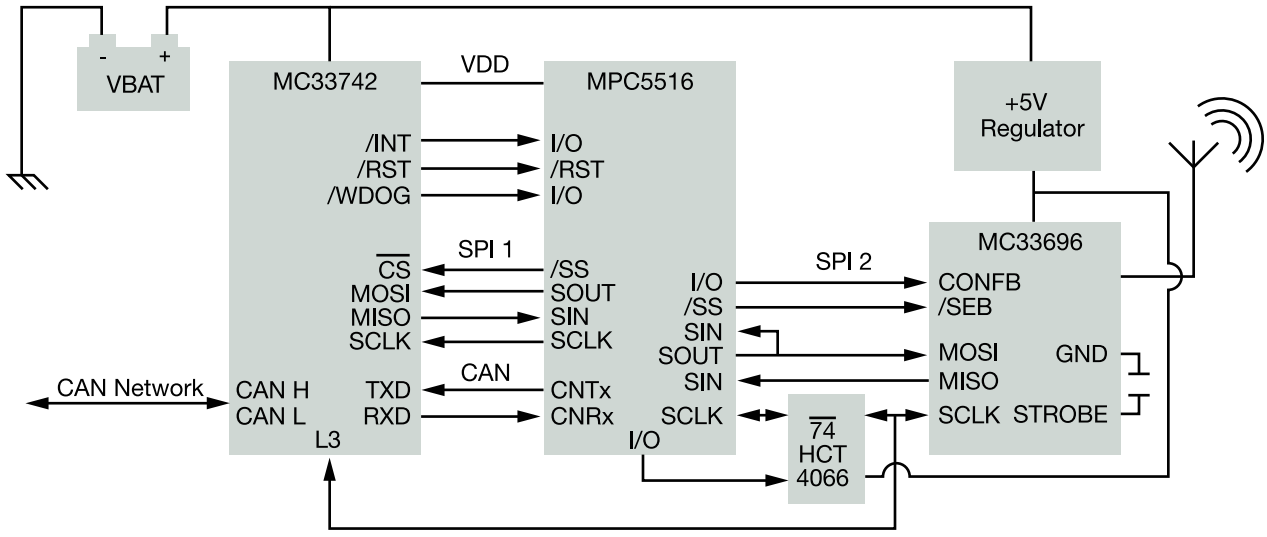

Figure 1

System 1—Securing the Car

The process can be separated into the following steps.

Receiving message

The remote control transmits the message to lock and secure the car. The MC33696 processes and decodes the Manchester encoded message. Now the SPI 2 bus carries the message data, during which the MC33696 behaves as an SPI 2 bus master and the MCU as an SPI 2 bus slave. The data transmission is managed by SCLK and MOSI signals, maintaining the configured baud rate. Once the message is accepted, the MCU verifies the data content to avoid misinterpretation.

Message verification

The verification message data is coded using the Manchester encoding algorithm. The MCU timer peripheral can be used to simplify the coding process. The MCU SPI module is now configured as the SPI 2 bus master and the MC33696 operation is changed from receive mode to transmit mode. The data transmission is started using the SCLK, MOSI and SS lines. The remote controller unit verifies the message data and, where successful, sends back the authentication code. The message verification process can be repeated as many times as needed, and the number of cycles is the system designer's choice.

System configuration

Once the authentication code is received, the MCU notifies the relevant application systems through a CAN bus that the car has to be locked and secured. Now the application systems can either enter low-power mode or process an action, depending on the desired function (e.g., activating the door lock). The BCU now waits for the systems' responses, configuring the MC33696 to Receive mode and disconnecting the SCLK line to offload the MC33696 SCLK pin buffer. The MCU configures the MC33742 L3 as a wake-up pin and sends the sleep command via the SPI 1 to switch off the VDD and V2 regulators. The module power consumption goes down from 50 mA to 100 µA (see Table 1), except for the MC33696 device, which needs the higher power to enable a system wake-up when a new wireless message is received.

System 1—Unsecuring the Car

Process a wake-up

The remote control message is received by the MC33696 transceiver. The MC33696 SCLK signal appears on the MC33742 L3 pin and wakes up the SBC. The SBC enters normal request mode and the VDD regulator is turned on. The MCU is supplied, and the MCU initialization is processed.

Trace wake-up source

The wake-up source is traced by reading the MC33742 wake-up register (WUR) contents via the SPI 1 bus. If the MC33696 has triggered a system wake-up, then the MCU configures the SPI 2 bus SPI module as a slave and interconnects the SCLK line. So far, a few bytes from the message were lost because the MCU was not able to receive the message. The MCU now waits for a new RKE message to verify message data and process authentication in the same manner as described above.

System 1 initialization is processed only in the case of a successful authentication. If the message data does not contain an unsecure command, the BCU enters the low-power mode again. The mechanism described is very useful when working in noisy environments. Sometimes the noise can be misinterpreted as remote control signals and cause errors in the system. Another source of system wake-up can be a CAN message received by the MC33742 CAN transceiver.

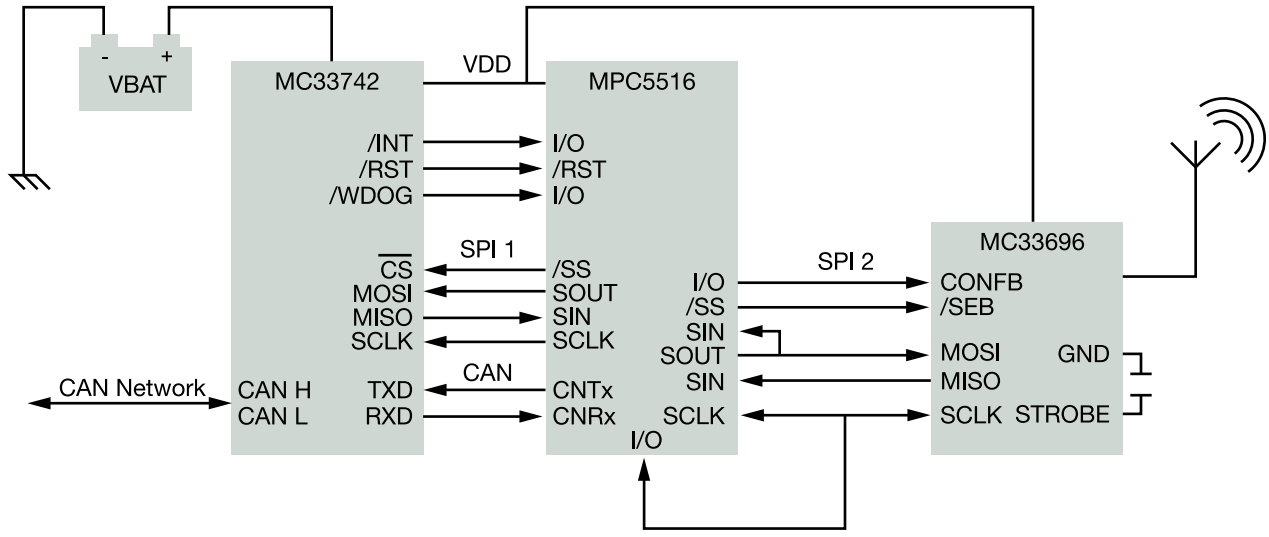

Figure 2

System 2—Securing the Car

The remote controller message processing and data verification is identical to that of System 1. System 2 configuration is also processed in the same manner, except the SBC and MCU enter different low-power modes. The SBC maintains the VDD power supply for the MCU and the MPC5516 enters Sleep 2 mode, enabling wake-up on an SPI message reception.

The MPC5516 offers wake-up sources when in Sleep 2 mode as follows:

- Real time clock (RTC)
- Autonomous periodical interrupt (API)
- I/O pins wake-up on negative, positive or either edge transitions

In the case of the System 2 solution depicted in Figure 2, the wake-up I/O pin connected to the SPI SCLK signal line is used to wake-up the MCU. During low-power mode, the MCU has to periodically update the SBC watchdog to prevent a reset.

System 2—Unsecuring the Car

Process wake-up

The remote control message is received and decoded by the MC33696 transceiver. The MPC5516 leaves Sleep 2 mode when the SCLK signal appears on the I/O pin. The MCU can now behave either as it would after a reset, or start at some place in program memory. This is enabled due to the reset recovery pointer register being loaded before the MCU enters the Sleep 2 mode. The reset recovery pointer register contains the program memory address, where the program counter starts when a wake-up recovery is processed.

The message authentication algorithm can be processed either out of RAM or out of flash memory, depending on the application designer's choice. Code running out of RAM has a faster system recovery, but it is limited by the RAM capacity and the necessity to copy an application program into RAM before entering Sleep 2 mode. Plus, MCU power consumption increases with the number of enabled RAM blocks. On the other hand, flash offers enough space for application code, but the trade-off is slower system recovery.

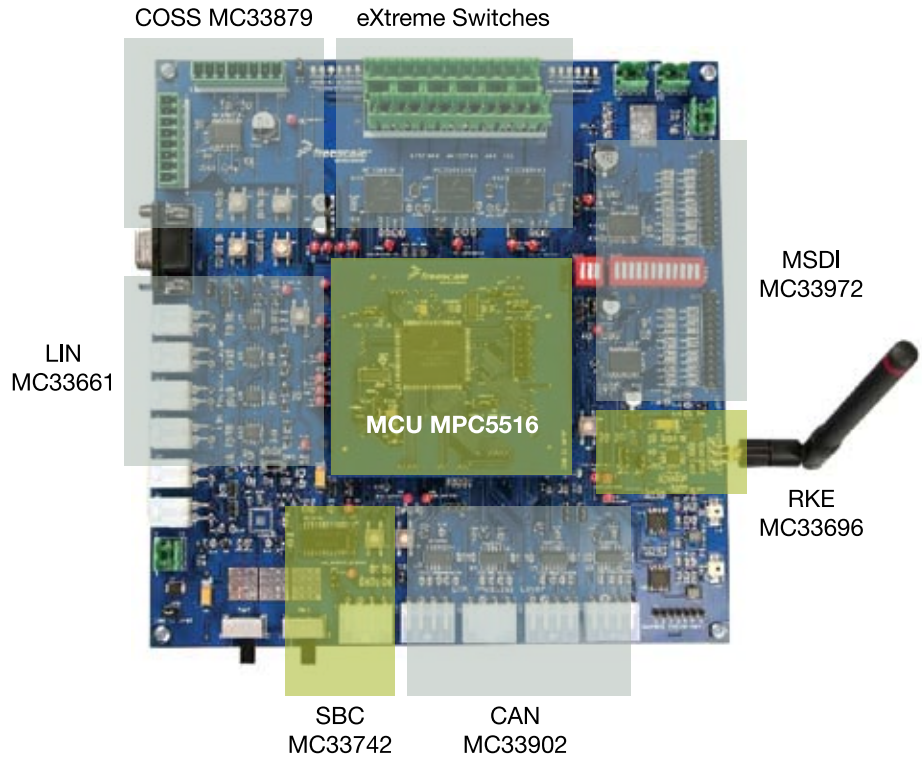

Figure 3

Test, Measurements and Evaluation

The tests and measurements were carried out on Freescale's BCU platform depicted in Figure 3. The BCU board contains devices to support:

- Connectivity—LIN (MC33661), CAN (MC33742, MC33902), FlexRay
- Load control—eXtreme Switches (MC15XS3400, MC35XS3400 MC10XS3412) and COSS (MC33879)
- Keyboard control—MSDI (MC33972)
- RKE—MC33696

For the remote control solution, the 8-bit MC9S08QG8 MCU and MC33696 UHF transceiver were chosen. The aim of the analysis is to see if both concepts can work in a real BCU platform. Testing was focused mainly on MCU, SBC and RKE interconnection verification and power supply consumption measurement.

The testing results are captured in Table 1 and Table 2. The measurement shows (Total—Strobe 1/10) that System 1 current consumption in low-power mode is roughly 50 times lower than in run mode, whereas System 2 low-power mode current consumption is only 10 times lower than in run mode.

In terms of reaching the lowest power consumption, the following rules can be extrapolated from Table 1:

- When the MCU power supply is maintained, disable all peripherals to keep the current consumption as low as possible
- The MC33696 transceiver requires the use of the strobe oscillator 1/10
- The more devices maintained by a power supply when the system is in low-power mode, the higher the power supply consumption (see Total (Strobe 1/10)—1.1 mA in comparison to 5 mA)

Reviewing Table 2, introducing the MPC5516 power consumption in Sleep 2 mode results in the MCU power supply current that is lower than 100 µA (VDDR, VDDA, VPP). The remaining 600 µA is drained by the devices connected to the MCU (e.g., eXtreme switches, COSS and so on). That current is not consumed by the MCU; it's just flowing through the ports to external devices. This means the supply current is also dependent on the MCU environment.

System 1 and System 2 Current Consumption			
Device	**System 1, System 2 Run Mode**	**System 1 Low Power Mode**	**System 2 Low Power Mode**
MPC5516	47mA	0.0mA	1mA
MC33742	3mA	0.1mA	3mA
MPC5516 + MC33742	**50mA**	**0.1mA**	**4mA**
MC33696 Strobe 1/10	1mA	1mA	1mA
MC33696 No Strobe	10mA	10mA	10mA
Total (Strobe 1/10)	**51mA**	**1.1mA**	**5mA**
Total (No Strobe)	60mA	10.1mA	14mA

Table 1

MPC5516 Sleep 2 Mode Current Consumption		
MPC5516 pin (pin number)	**Run Mode [mA]**	**Sleep 2 Mode [mA]**
VDDR (46)	40.5	0.06
VDDA (144)	6.6	0.014
VPP (78)	0.014	0.011
VDDE1 (96, 119)	0.880	0.330
VDDE2 (16, 33, 48)	0.64	0.001
VDDE3 (61)	0.07	0.308
Total	**48.7mA**	**0.724mA**

Table 2

Table 1 and Table 2 state the average current values.

The remote control measurement shows that an average current consumption equals 15 mA during message processing when in run mode and 2 µA when in low-power mode. The remote control is supplied by a popular 220 mAh CR2032 3V coin cell battery, which potentially enables it to continuously send five million messages per battery life.

The remote control is active only in case of BCU RKE message processing. Considering the fact that the battery leakage current also reduces the battery capacity, the battery life estimation becomes a challenging task. Experience shows that a remote control device can be used for a BCU RKE application demonstration for almost two years without battery failure.

Summary

Let's compare the systems based on the following categories:

- Hardware complexity
- Average power consumption
- Additional benefits

Comparing the system block diagrams, we can deduce that System 1 requires an additional +5V regulator for the MC33696 power supply and a switch for disconnecting the SPI SCLK line between the MCU and the MC33696 while in low-power mode. Those components are essential. When the application is in low-power mode, the SBC VDD and V2 regulators are turned off and the MCU is not powered. The current consumption is at the lowest possible level because the +5V supply VDD and V2 lines are powered down.

In System 2, the situation is slightly different. The +5V supply line is maintained, so devices consume energy. To keep the power consumption at a reasonable level, each device has to enter low-power mode. In comparison, System 1 has no need for such a capability.

On one hand System 2 is drawing more energy from the battery, but on the other hand it offers the user more flexibility due to the MCU having a maintained supply. Table 3 shows the differences between the systems' wake-up procedures. In System 1, wake-up is controlled by the SBC only. System 2 has the advantage of a supplied MCU, and thus additional wake-up sources can be utilized. The wake-up procedure is faster because there is no need to wait until the power supply is stabilized. An additional software control can tune System 1 and System 2 application performance to meet target power consumption and functionality.

Devices Wake-up Sources				
Device	**Source**	**System 1**	**System 2**	**Note**
MC33742	Lx pin	Yes	Yes	
	Internal timer	Yes	Yes	HS connected to Lx pin
	CAN bus	Yes	Yes	
MC33696	Message reception	Yes	Yes	SCLK connected to Lx pin
MPC5516	RTC	No	Yes	
	API	No	Yes	
	I/O pin	No	Yes	

Table 3

Lx pin can be one of MC33742 L0, L1, L2, L3 pins.

Petr Cholasta is an application engineer at Freescale Semiconductor and has been with company for almost five years.
He focuses on automotive application design using a wide range of Freescale products (8/16/32-bit MCUs, IDCs and APDs).
He holds a master's degree in electrical engineering.

Inga Harris

Waste Management
Eco-Friendly 32-bit Embedded Systems

Introduction

It is widely accepted that all consumer, industrial and automotive electronic system designers must consider, at every stage of the development and manufacturing process, the environmental impact of their products. Consumers are more aware of their carbon footprints while still demanding higher performance, and governments are penalizing companies that do not comply with more stringent environment regulations. This means that manufacturers must carefully balance performance and environmental factors as not to sacrifice one for the other.

The heart of most electronic and mechatronic systems is a microcontroller (MCU), and it is here that the win-win balance between efficiency and performance can be engineered. However, many MCU decision makers do not take the time to fully understand its key contributions and the datasheet numbers on which they base their decisions.

This article discusses some methods which can be used to achieve an effective balance between performance and efficiency for battery powered systems by challenging preconceived notions on power versus performance. Using a thermostat as an example, the article benchmarks power consumption and performance data among three very different cores from two separate manufacturers and shows what a massive impact one core can make. The data will prove that increased performance does not have to result in higher power, which translates to shorter battery life. It will also show that electronic and mechatronic systems can now afford to migrate from 8-bit MCUs to higher performing 32-bit MCUs while still achieving reduced power at a similar system cost.

Waste Management–Tips to Reduce Power

Whether the goal is longer battery life or lower power consumption from main supplies, many techniques are the same. The Battery Life Calculator application, available for download at **www.freescale.com/lowpower**, is a good place to begin a low-power project. It indicates the effect CPU speed, operating mode, voltage, temperature and the number of enabled modules can have on an MCU's power consumption.

While it is true that microcontrollers use more power at higher temperatures and higher voltages, system designers can minimize power consumption by carefully analyzing the MCU's supply voltage and the environment in which it operates. Where previously there were issues sourcing system components that would operate below 2V, the latest generation components, such as external EEPROMs and transceivers, now offer this voltage option.

Clock gating is a very effective and widely used strategy that reduces the power consumption while maintaining system performance and functionality. Silicon circuitry uses more power when it is clocked, even if it is disabled, than when the clock is turned off. Clocking non-required modules can consume as much as 40 percent of a chip's power. By shutting off (gating) the clocks in unused portions of the MCU, thus stopping the data toggling, sizable energy savings can be realized.

Peripheral modules can often be clocked at different speeds since modules often have several clock sources. Selecting the lowest frequency clock source and the lowest possible operating frequency to achieve the module's purpose will

provide important power saving. For example, the analog-to-digital converter (ADC) on the MCF51QE32 MCU has four clock sources:

- Bus Clock
- Bus Clock ÷ 2
- External reference is a clock that, if available, can be optionally enabled in stop modes by simply setting a register bit (EREFSTEN in ICSC2r)
- Asynchronous clock is an internally generated clock (high/low speed). This clock remains active while the MCU is in wait or stop3 mode and allows conversions in these modes for lower noise operation.

Since the ADC sampling frequency is governed by Nyquist sampling theory, the minimum clock frequency is determined by the signal being sampled, which may eliminate the use of the lowest frequency clock source. More clock source options mean more low-power flexibility in the product design.

More low-power modes are finding their way onto microcontroller features lists. These power modes are designed to offer maximum system performance only when required. To achieve the best energy efficiency over the profile of the application, power modes such as run, wait, stop and standby should be used wisely to manipulate power usage and to get the most efficient use of the power source. It is best to spend as much time as possible in the lowest power mode and to execute the main functions as quickly as possible, as shown in Figure 1.

The lowest useable power mode is usually governed by the wake up interrupt sources. The main function execution speed should be as fast as possible, but in practice this time must be weighed against how long it will take to re-enable the clocks and get to that speed, especially when clock source switching and FLL/PLL stabilization is a factor. Figure 1 shows this slower speed being used when the device is wakened by an

Best Practice Low Power Design

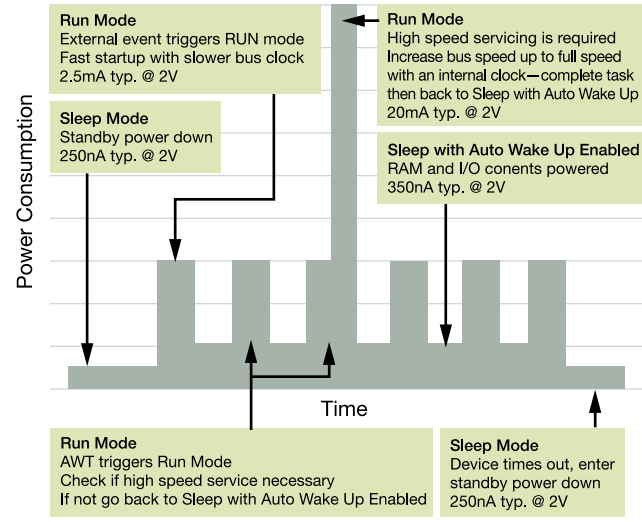

Figure 1

auto wake-up (AWT) function, e.g. a real time clock, but is not ramped up to full speed unless it actually needs to respond to a stimulus.

A key component of longer battery life is at what voltage level the applications fail and how much notice is given. The microcontrollers referred to as 3V devices are designed to run off two fully charged AA batteries, up to 3.6V, and to continue operation down to 1.8V, when the batteries are considered spent. At this voltage, microcontrollers have restrictions on speed and functionality and this should be understood before choosing a target microcontroller for the system.

Applications where the flash memory is used to store data can only function down to the lowest flash write voltage, which is frequently above the minimum microcontroller operating voltage as illustrated in Figure 2. By selecting a microcontroller with a

Microcontroller Operating Range

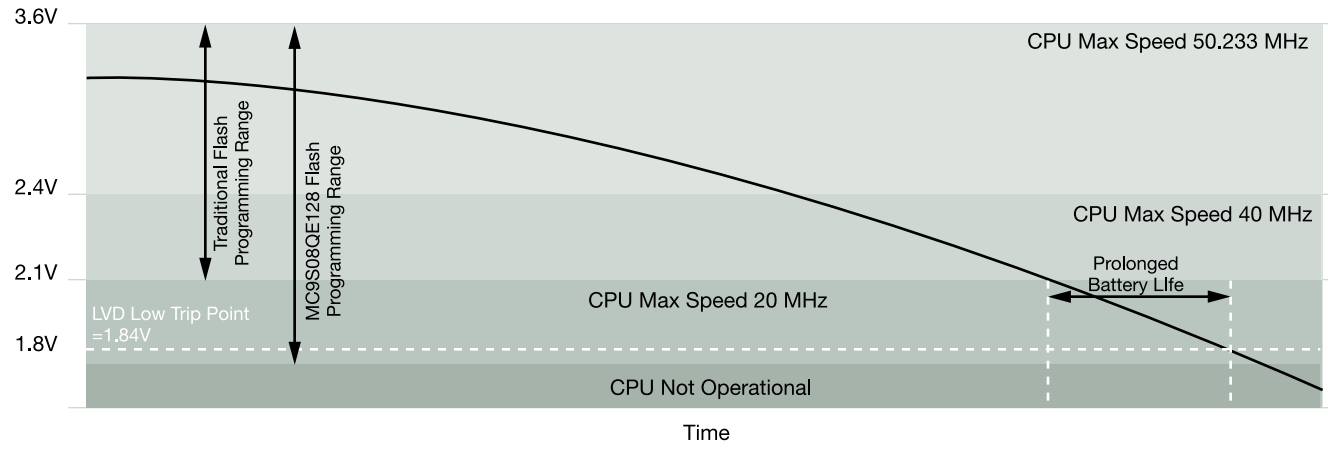

Figure 2

larger flash write operating range, battery life can be extended by over 10 percent, depending upon the load. (The smaller the load the longer the battery will last with a 1.8V dead level, compared with a 2.1V dead level.) In this situation, it is also important to consider what warning level of the approaching end of battery life will be given in order to send change battery messages and store key system data. The low voltage circuitry, which is embedded in most microcontrollers, offers this function when low voltage warnings (non-reset) flags are available. These are normally interrupt-based flags that can be a few hundred millivolts higher than the reset level.

Proper termination of unused pins is another obvious, yet often overlooked, method to reduce the systems power consumption. Floating pins use current. In the interest of system robustness, unused pins should be grounded and then set as inputs (or outputs pulled up internally) as per the device's reference manual's recommendation. In the case where some pins are not bonded out on smaller package options, it is also important to remember to terminate them for robustness and power reasons.

Power versus Performance

Contrary to popular understanding, the power-performance relationship is not linear. The interaction is different for all microcontroller architectures, primarily based on the peripheral modules and component design rather than the microcontroller core.

Process technology is one aspect of the component design that affects the power-performance relationship. The process technology is the basic building block of any semiconductor product, and the characteristics of the process determine the power consumption of all circuits built on that process technology. Regulator design, temperature compensation, dynamic voltage and frequency compensation and state retention circuit techniques are also contributing factors to component design. Since all these techniques are manufacturer independent, the resultant power characteristics manifest themselves in the electrical characteristics of the device.

As discussed previously, it is best if the device spends as much time as possible in the lowest power mode and executes the main functions as quickly as possible. Running the main function as fast as possible limits the power consumed, (as the current and frequency relationship is not necessarily linear) and there may be varying degree slopes and step functions as different DCO multipliers and silicon areas are used.

- Once in run mode, execute as simply and quickly as possible and then return to the low-power mode, minimizing subroutine calls.
- Use the PLL/FLL to run at top speed when the CPU is active, and bypass the PLL/FLL to drop to the lowest possible bus frequency when required, without unlocking the FLL/PLL (as re-lock times can be long).
- It may be beneficial to use a wait mode, with an immediate wake up time, if a peripheral will be clocked and the system is waiting for an imminent interrupt instead of continuously polling a flag.
- Interrupts which occur seconds later are best handled using sleep/stop modes.

These recommendations will enable the system to maintain a healthy balance between power and performance. The following case study highlights how one such application, a domestic thermostat, running on either an 8-bit, a 16-bit or a 32-bit microcontroller, performs. By migrating to a carefully selected 32-bit microcontroller and using all the applicable techniques, a reduction in power usage can be demonstrated.

Upgrading from 8-bit to 16-bit and 32-bit Microcontrollers without Compromising Efficiency

Domestic thermostat example: A domestic thermostat is a common low-power, battery-powered application that is understood globally. It usually features a real-time clock (RTC), a 12-bit ADC, a serial peripheral interface (SPI), general purpose I/O pins, timer and simple pushbuttons for user inputs to change the display information and to adjust settings. The main code loop runs once every second, prompted by an interrupt from the RTC module, to update the clock. One second is a long time in low-power applications, so between interrupts, the device is in a sleep mode, resulting in very low-power operation.

- An RTC interrupt flag is set every second, and the clock function is executed immediately to ensure that it is always completed and no seconds are lost. To minimize time spent in the main software loop, the maximum bus frequency possible with a 32 kHz crystal is used.
- The ADC takes a voltage reading from the application's thermistor every 30 seconds and compares it to a target value set by the user. The CPU then decides if the room temperature needs adjusting, indicated by an LED.
- The SPI transmits three parameters (one at a time): current room temperature, temperature setpoint and time, controlled by the state variable.
- Using the SET button while in set state, the user can adjust the target temperature. When the device is in the set state, pressing the SET button repeatedly cycles the temperature range from 10 to 40 degrees Celsius. Once the correct temperature is reached, the mode is exited via the STATE button.
- The time is updated by the SET button when the device is in the time state. Each button press will add one minute to the time (24-hour clock). Time state is only enabled for 10 seconds, after which it automatically moves to temp state.

More details on the application block diagram, state machine and code can be found in the application note AN3506/D, titled Migration from TI MSP430 to 9S08QE128 or MCF51QE128 Flexis™ Microcontrollers—Enhancing Low-Power Performance, available as a download from **www.freescale.com**.

This same application was run on Freescale's Flexis™ MC9S08QE128 (8-bit), MCF51QE128 (32-bit) microcontrollers and Texas Instruments MSP430FG4619 (16-bit) microcontroller, taking advantage of all low power features available. The main loop uses run mode and active modes. The button press debounce function uses low-power wait mode and LPM3 mode. The sleep-every-second mode uses stop3 and LPM3.

Results: Figure 3 shows the normalized energy used by the 8-bit, 16-bit and 32-bit devices, with the 8-bit MC9S08QE128 used as the reference in the three main sections of the application code, along with the total energy used by each device.

The QE128 devices have a significant advantage in the main loop due to their higher bus speed capability, enabling a better mA per MHz ratio. Although the current draw is higher, the amount of time it draws is shorter. Between the two QE128 devices, the MC9S08QE128 has the advantage, predominantly due to physical silicon size.

The MSP430FG4619 device's use of LPM3 mode, instead of the QE128 LPW, during the button press function gives this device the advantage. A key point is that the LPM3 mode has more wake up sources from the timer. The QE128 device stop modes do not allow exit using the timer module and therefore cannot be used in this case. Between the two QE128 devices, the MCF51QE128 has the advantage as the core handles the wait instruction differently and more efficiently. The S08 has an indigenous wait instruction; however, wait mode is not native to the ColdFire® V1 CPU. The ColdFire V1 core does not differentiate between stop and wait modes, the core reads both as stop. The difference between the two is at the device level. The system option register has two bits that control the effect of a stop instruction on the MCF51QE128, stopE and waitE, which is a more flexible environment for the microcontroller designers to maximize the power saving.

The QE128 devices gain an advantage by being able to use stop3 with only the RTC using ERCLK, enabled in the sleep section, whereas the MSP430FG4619 is limited to LPM3, since LPM4 cannot clock the RTC. The MC9S08QE128 has the advantage with the lowest sleep current of all three devices.

Overall, assuming a button press is executed every cycle, the 32-bit microcontroller is the most power efficient. The 8-bit microcontroller consumes 7 percent more power and the 16-bit microcontroller consumes double that of the 32-bit microcontroller.

Implication: How did this happen? The electrical specifications don't show this result. In this case, there is no minimum nor maximum requirement to the bus speed, since the analog signal being monitored is slow changing temperature. The other main inputs from the outside world are the button presses, which are also not affected by bus frequency. By taking advantage

Energy Comparison Plot of the Thermostat Example

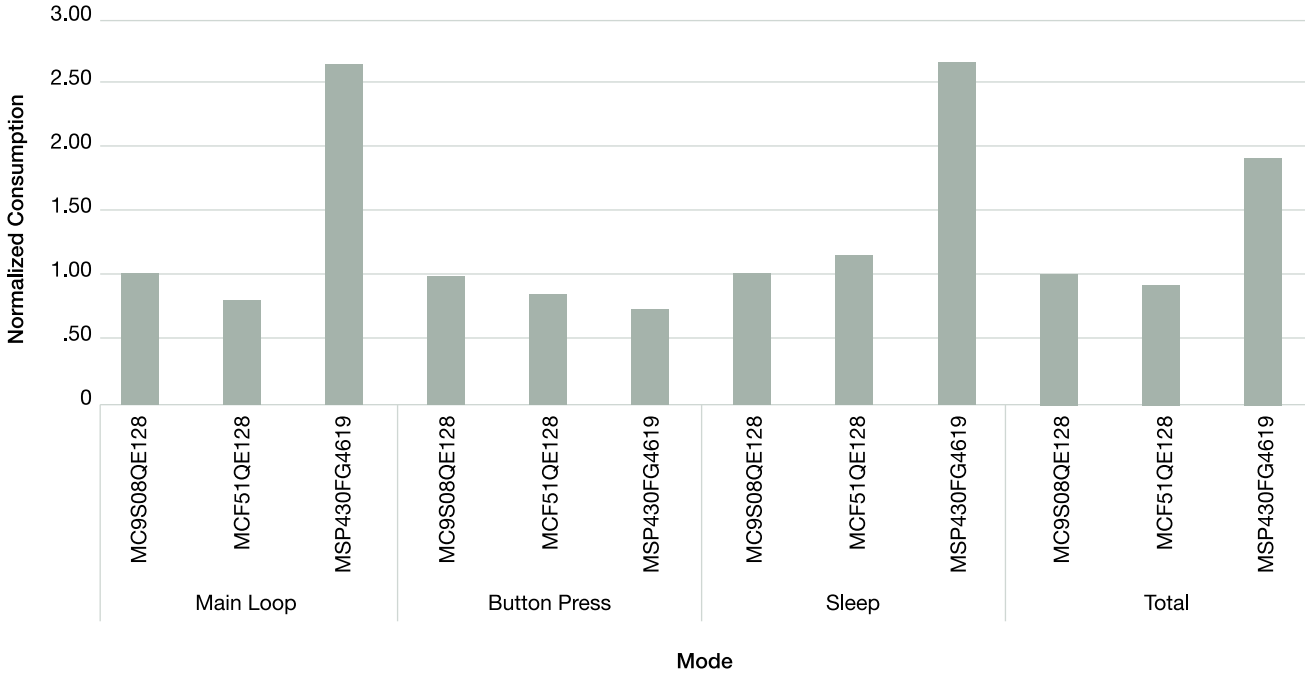

Figure 3

of the extra performance on an ultra-power-efficient 32-bit microcontroller, 8-bit and 16-bit microcontroller benchmarks are beaten. This highlights that a simple comparison of Idd values in the documentation, without thought into what modes can be used, how often and how long the system will be in those modes, cannot necessarily provide the designer with accurate performance and power consumption parameters for the end application.

Choose the Right 32-bit Device

The ColdFire® V1 family of products (part numbers MCF51x) is the industry benchmark in low-power 32-bit microcontrollers. Another extremely popular low-end, power-efficient 32-bit microcontroller core is the CortexM3 from Luminary Micro. Benchmarking data between the ColdFire V1 core and this subset of ARM7 core (LM3S811) shows that the ColdFire V1 has the advantage on execution of basic functions, data management and math; whereas the CortexM3 has the advantage in code density. This means that the ColdFire V1 code may take up more memory, but the functions will execute faster, allowing more time to be spent in the low power modes.

Executing at the same frequency, the MCF51QE128 consumes ~50 percent less current than the LM3S811 in run mode, which when added to the sub 1 uA stop modes means that the MCF51QE128 achieves average power consumption of less than 1 uA, depending on the period. The LM3S811 average current consumption would be the order of mA, highlighting how important choosing the right core and device is.

Conclusion

Time is well spent when engineers think about module use, speed and timing when planning their system design. Not only are the prices of 8-bit and 32-bit microcontrollers converging, but also their power characteristics, meaning that designers can take advantage of 32-bit's superior processing power. Microcontroller manufacturers are also investing time and money in developing innovative techniques and solutions that enable their customers to meet the conflicting goals of increased performance and reduced power consumption. The complete picture of process technology, circuit and module design, component design, platform design and system and application software is in place, ready to be used. Tools, such as Freescale's Battery Life Calculator, available for download at **www.freescale.com/lowpower**, are available to help estimate the microcontroller's probable power performance in the target application before the designer invests in tools and software.

It is also critical to know to what voltage level the application must work, as different microcontroller devices operate down to different voltages, with some restrictions on frequency, temperature and flash programming ability. By remaining in low-power states as long as possible, waking up only when needed then returning to sleep mode as soon as possible, a power efficient system can be built with increased performance.

Inga graduated from the University of Strathclyde with an honors degree in electronics and electrical engineering. Her Technical Marketer and Applications Engineering career has predominantly been with the 8-bit and ColdFire MCUs with a focus on Consumer and Industrial markets. She has numerous articles and application notes published.

Oliver Jones

Proximity Sensing
Adapting to the Environment

Abstract

An interesting alternative to mechanical switches is the proximity sensor. The term “proximity” refers to the fact that there is no contact between the medium you’re trying to detect (i.e., finger, liquid, metal, etc.) and the actual sensing element. Most likely a plate of glass or plastic separates the two. Although you are likely to touch the separating element, there is no physical contact with the sensor.

Proximity sensing technology enables adaptive controls, alleviates isolation issues, improves overall application robustness, generates almost unlimited design flexibility and fosters new functionalities.

Proximity Sensing Technologies

Table 1 describes some of the common sensing technologies.

Additional options include Hall effect, magnetoresistive, radar, sonar and others.

The two technologies most used today for mechanical switch replacement are optical and capacitive sensing, however, as market interest indicates, capacitive is the most versatile and flexible. This article describes some of the theory behind capacitive detection and shows how this theory can be applied to the human-machine environment.

Common Sensing Technologies

Technology	Detection	Mode	Advantages	Disadvantages
Inductive	Metal	Induced electromagnetic currents	• Operates in harsh conditions • Rapid response time	• Short range • Detects only movement • Difficult array setups
Ultrasonic	Virtually all objects	Sound wave echo	• Long range • Measure distance	• Cost • Dead zone • No idea of size/shape
Photoelectric	Solid objects	Reflection or absorption of light different to background	• Medium range	• Possibility of interference • Cost • Pb in fog/smoke/non-transparent materials
Capacitive	Objects capable of absorbing or creating electric charge	Permitivity variation to background	• Simple array construction • Detect metal and non-metal	• Short range • Object properties

Table 1

Capacitor Basics

A capacitor is a device made up of two electrically conducting materials (called electrodes), each at a different potential, separated by a non-conductive material (insulator). The physical value of a capacitor depends on the dielectric constant of the insulator, the relative permittivity of free air, the area of each electrode and the distance separating the electrodes. This value corresponds to the amount of energy the capacitor is able to hold.

Applying a voltage to one electrode that is different to that present on the other induces an electric current through the capacitor, which decreases as the charge builds on the electrode. This potential difference creates an electric field between the electrodes.

Capacitive measurement techniques

Time constants: Input a step function to an RC network where R is fixed, measure the time the output takes to achieve a given voltage.

Phase shifts: Input a periodic signal, measure the delay, due to the capacitance, on the output signal.

Frequency modulation: Design a circuit whose frequency depends on the charge and discharge of a capacitor.

Amplitude modulation: The amplitude of an ac waveform changes due to an RC network, where R is fixed.

Below are simplified schematics of how to perform these measurements (Table 2).

In the real world, the challenge is finding the trade-offs between sensitivity, robustness, noise immunity and cost.

Measuring RC time constants off a square wave function is without doubt the simplest and least expensive solution. However, the drawbacks are sensitivity, detection frequency/speed and electromagnetic noise, since you're typically injecting a mono-pulse step function with a given repetition rate for delay averaging purposes.

Freescale has chosen this technique for a new family of MPR08x proximity sensors based on our S08 microcontroller. It provides the optimal compromise between performance and cost—ideal for keypad, tactile screens and simple button replacement.

Phase shifts have similar sensitivity issues, but tend to have faster response times. Again, noise may be an issue. This measurement technique can easily be integrated into an MCU but does need some external components.

Frequency modulation is a good solution for discrete designs, especially when using square/triangular waves. An F-to-V converter then gives information that is easily interpreted by an MCU. The drawback is noise.

Amplitude modulation is quite design intensive, however it gives the best performance in terms of electromagnetic robustness, since you can easily adapt this technique to sine-waves. The sensitivity is similar to that of frequency modulation.

Freescale has built a portfolio of products based on a small signal sinusoidal excitation. Due to the virtually perfect sine-wave, the resultant electromagnetic interference spectrum is best in class. This portfolio has been in full production for quite a number of years, demonstrating excellent robustness and performance.

Techniques for Capacitive Detection

	Simplified Electrical Representation	Receive Signal Format	Advantages	Disadvantages
RC Time Constant	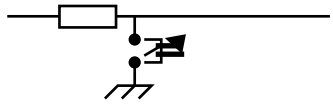	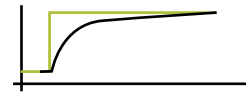	• Simple • Feasible in discrete • Cost	• Sensitivity (resolution) • Response time • EMI
Phase Shift	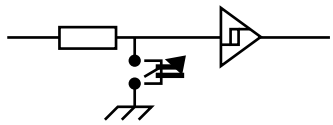		• Simple with MCU • Cost	• EMI • Sensitivity (resolution)
Frequency Modulation	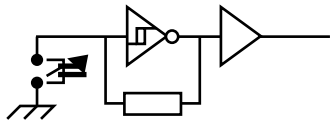		• Discrete solution	• EMI • Environmental sensitivity
Amplitude Modulation	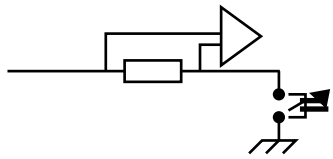	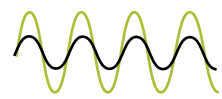	• Shielding • Sensitivity • EMI • Flexibility	• System design complexity

Table 2

Real-World Solution

So if all we need to do is measure capacitance, where's the problem? Since the capacitance changes with the environment, just about anything will influence the measurement—insects or mud, tropical climates or desert dryness, children's toys or even a sack of potatoes. The key to resolving these issues is how you calibrate your sensing system.

Not only can the external environment impact the measurement, but also the design of the measuring system can play an important part in the sensitivity and dynamic range. Unwanted capacitance (or parasitic capacitance) can be created by the chassis (fixings, metal housings, etc.) or by routing the electrode path close to other signals (ribbon cables, PCB routing, etc.). Although there may be certain applications where you want to detect this, such as tamper proofing or security detection, this is more of an inconvenience than a benefit for the vast majority of uses.

Two options exist to overcome disturbance issues: either you ensure that the A/D part of the capacitor equation is so small that the result has little or no impact or you shield the measurement channel. We have seen previously that an electric field is created between two points having a different potential, therefore by creating a shield circuit with nearly the same amplitude and phase as the electrode signal ensures that there is little or no potential difference between the two signals, thereby canceling out any electric field. By ensuring sufficiently low shield impedance, the parasitic capacitors that now exist between the shield and the chassis, GND signals, etc., can be charged and discharged without affecting the signal amplitude.

Applying the Theory: Making Life Easier and Less Power Hungry

Optimizing the man-machine environment

When man and machines work together, there is often a physical limit or exclusion zone that constrains the machine. This limit is often defined as the limit of "inconvenience" for the operator, that is to say a position that the operator would normally have to stretch to reach. The underlying objective of this is to ensure that under no circumstances can the machine get too close to someone without that person making a deliberate choice.

However, what can be considered "distance of security" in one case, such as a robotic tool, can be interpreted by the operator as being just a little too far to be comfortable.

Imagine the improved convenience and machine performance if the robot was able to adapt to the operator's position. By detecting the operator's presence at a given distance, the robot could safely adjust its position with respect to the operator.

Safety measures can also be enhanced in applications where user presence must be validated before operation, such as a lawnmower. If the user slips or loses control of the lawnmower in any way, the mower would stop operating as quickly as possible. Another example is an industrial stamping machine where the user must be detected at a safe distance from the equipment prior to its activation.

The concept of protecting people can equally be applied to protecting sensitive equipment, such as a camera. If it's dropped, using proximity detection would enable the equipment to detect the absence of a human presence and place itself into a more secure state, such as retracting the lens.

Automatic door openers

One of the most common applications for presence detection is the automatic door. Typically, as you approach a door you are detected by an optical sensor, or your weight closes a contact in the floor.

The electric field sensor can be integrated into the floor and can detect the presence of a person through different substances (wood, tile, carpet, etc.). There are no moving parts and the sensor is impervious to rust and virtually indestructible, making it a suitable replacement technology for the mechanical pressure sensor. The physical nature of the electrode ensures a well defined and limited sensing area, unlike that of an optical solution where you need to define a volume and sensitivity threshold.

Alternatively, proximity sensors can be embedded in the wall or other object to be activated only by voluntary movement. This also allows the door to be opened without any physical contact.

Optimizing access control can also lead to benefits in energy consumption. Minimizing the time a doorway remains open ensures the shortest possible exchange between hot or cold outside air with the conditioned air in the building.

Occupant and presence detection

If you want to check how many people are on an aircraft, how many seats are left in a cinema or how many beds are occupied in a hospital ward you can either count the number of tickets sold or the number of people present, or you can let the seat or bed, each with proximity sensing technology, detect by itself whether it is occupied or not.

By using multiple electrodes per seat, not only will a person be detected, but also his/her size and position will be measured. This is particularly useful when employed in conjunction with automotive airbag safety systems.

Energy consumption in battery powered equipment

There is general concern about the amount of energy wasted by electronic equipment when not in use. Displays and lights that remain lit and equipment that continues to draw power, even when turned off, are just a couple of examples.

Rather than setting a certain time limit before extinguishing backlights or putting equipment in standby, why not detect the presence of the user and adapt the energy consumption accordingly?

Battery powered applications can remain in stand-by mode until a proximity sensor detects the approach of a user's hand. The device then automatically powers up. Then, as the hand moves away, the interface can return to a stand-by low-current mode.

Ice Detection

The dielectric properties of water are altered as it changes state from gaseous to liquid to frozen. Therefore, for instance, as water vapor between two electrodes changes to ice, the capacitance value across those two electrodes will vary.

This phenomenon can be used to detect any ice build up in a freezer, helping prevent the igloo effect, where ice actually acts as an insulator. Under extreme conditions ice build up will prevent the compressor from cooling the freezer sufficiently, resulting in wasted energy and spoiled food.

Choosing the Right Technology

When considering which technology to use for which application, here's a very rough guide: the RC technique is best suited to applications expecting a "1" or "0" response. The amplitude modulation allows the user to identify and monitor the "fuzzy" bit between the "1" and the "0," or more accurately, the change in state. Here is a simple table that outlines which technology is best applied to which applications:

Choosing the Best Technology	
Amplitude Modulation	**RC Technique**
Liquid level (continuous)	Touch panel
Distance	On/Off switches
Dielectric properties	Discrete presence detection
Excellent EMI performance	Liquid absence/presence
Presence up to 15cm	
Touch panels in harsh environments	

Table 3

Freescale's Solutions

Freescale has been working with electric field measurement in harsh, security conscious environments for over 10 years, with particular attention to occupant detection in an automotive environment. In addition to the product portfolio, we provide evaluation kits that allow fast and simple experimentation and system construction.

The portfolio comprises three product ranges:

- An analog ASSP providing the highest sensitivity
- An MCU-based solution with IP developed to perform calibration, filtering and other debounce algorithms targeting touch panel solutions
- A software package for S08 and ColdFire V1 products that customers can integrate with their own application software to enable simple button replacement.

Additional Applications

The technology described above can be used to enhance security and automated equipment awareness in the following examples:

- Replacing the traditional mechanical dead man's switch with proximity sensor technology
- Allowing a robotic system to detect the presence of a human or animal to modify machine speed and movement accordingly
- Integrating access control sensing into flooring or walls

There are many other opportunities to apply electric field proximity sensing technology, including:

- Hiding light-switches behind the plaster board
- Placing electrodes behind glass to develop interactive touch-screen applications
- Liquid volume and level detection
- Access control and anti-pinch functions

References

www.freescale.com/proximity

Datasheet : MC33941.pdf, MC34940.pdf, MPR084.pdf

Application notes : AN1985.pdf, AN3456.pdf

Oliver Jones is a Product Marketing Specialist in the Consumer & Industrial Go-To-Market team in EMEA. Based in Toulouse, France, he's been with Motorola/Freescale for 11 years, occupying Product Engineering, Program Management and Marketing positions for Analog and Sensor Products.

Prashant Bhargava

Power-Aware Verification Using CPF

A New Dimension in Low-Power Verification

Abstract

With progressively shrinking transistor sizes, power wasted by leakage current rises exponentially and cannot be ignored. Designs use various low-power design techniques to save dynamic and leakage power. Until now, designs had very little or no verification of these low-power techniques and relied on back-end work or netlist simulations to find issues. But those methods are used very late in the design process and result in many engineering change orders (ECOs). Moreover, functional defects can be introduced with these techniques. Detecting these bugs is not an issue, but finding these bugs early in the design cycle will save time and money (no ECOs). Common Power Format (CPF) provides an interface for performing low-power verification. In this article, we will see how CPF fits into the design cycle and helps verify the low-power features and catch bugs faster. This article will also explain how CPF helps verify power switching and isolation at register transfer level (RTL), and it concludes by illustrating an example from a recently concluded live project.

Introduction

Low-power design techniques

As designs move toward smaller and smaller technologies (90nm, 65nm and so on), leakage current has become a significant contributor to overall power consumption. Designs are using new methodologies and design techniques to reduce the leakage current, including power gating, logic isolation, state retention and clock gating. The aim of these techniques is to cut off power or clocks to the portions of the design that are not required to function in a particular low-power mode, thus saving power. The flip side of these techniques is that they might alter the functionality of the design and introduce some fatal bugs. Hence, not only is the detection of these bugs important, but also detecting them as early as possible in the design cycle to save cost and time.

Need/importance of verification

The power intent of a design or SoC comes into the picture very late in the design cycle. It is only when the back-end flow starts that the power-related checks are made. The system's architect specifies the power intent, but that might become distorted by the time it reaches the back-end flow. RTL and its verification are oblivious to power-related functionality and are never checked. The only power-aware check, so to speak, is done for clock gating. Since RTL does not contain isolation, cells, level shifter, power pins, etc., these are not verified in simulation. Normally, the power intent of the design is captured and verified by Conformal Low Power (CLP), which requires a netlist, or, even better, a power connected netlist. Any bug found at this stage leads to an ECO, which can be expensive to correct in time, effort and silicon. See Figure 1. Since CLP is a static tool, it fails to verify the dynamics of power sequencing and isolation values.

Gate-level simulations (GLS) are also performed late in the design cycle and could lead to heavy ECOs.

This is where CPF comes to the rescue. With CPF we can verify the dynamic nature of power sequencing and isolation at RTL level, and bugs found therein can be easily corrected. But before moving to the CPF-based design flow, we look at the basic elements of a low-power design.

Elements of Low-Power Design

Any low-power design must have the following features verified to ensure they are functioning correctly and not affecting any other functionality.

Isolation

Isolation refers to holding the values of the signal coming from an OFF domain into an ON domain to a known inactive value so that any unknown value does not corrupt the working of the ON domain. See Figure 2. To do so, three types of isolation are available: high isolation, where the signal is isolated to a 1 value, low isolation, where the signal is isolated to a 0 value, and hold isolation, which holds on to the value of the signal just before entering the low-power mode. A verification engineer should ensure that the isolation cells are not powered by the switchable power, or else they will fail their purpose.

Retention

Retention is needed when the system requires the state to be the same after powering up as it was before switching OFF. Such cells aid in faster recovery from low-power modes. If required, these cells can be completely powered OFF by removing VDD (as shown in Figure 3).

While using these cells in the design saves on leakage power during power down, it also maintains the value of the registers, etc. where needed. For the verification engineer, it adds some extra work as he or she must ensure that retention is enabled before power is removed (i.e. Ret signal is driven before Pwr switch is flipped. See Figure 3). There should be no change in the output of such cells during low-power mode, and only the correct portion of the design has been retained. And lastly, during power up, retention should be removed once power is restored. It is removed only after the OFF domains have been reset and I/Os are powered up.

Level shifters

Level shifters are used on signals that are transitioning from one power domain to another. Though CPF specifies which signals should have level shifters, current simulators cannot simulate a level shifter and hence we cannot verify if the directions of level shifters are correct or if level shifters are missing. This is best performed using CLP and outside the scope of this article.

Impact of Bugs Found During a Typical Design Cycle

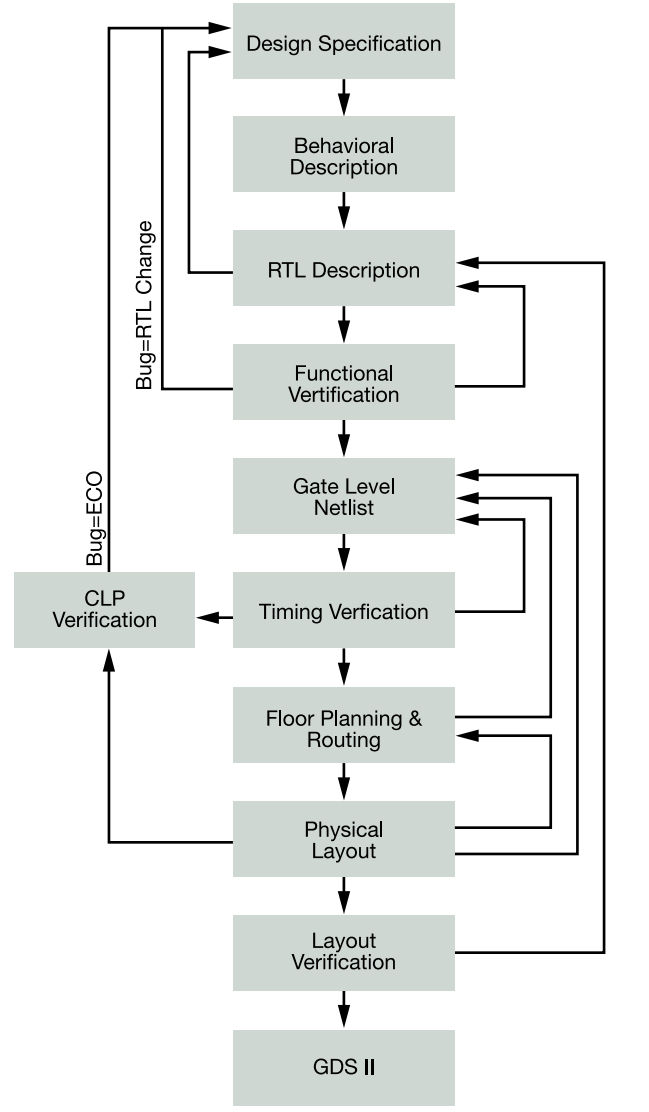

Figure 1

Isolation Cell Block Diagram and Timing

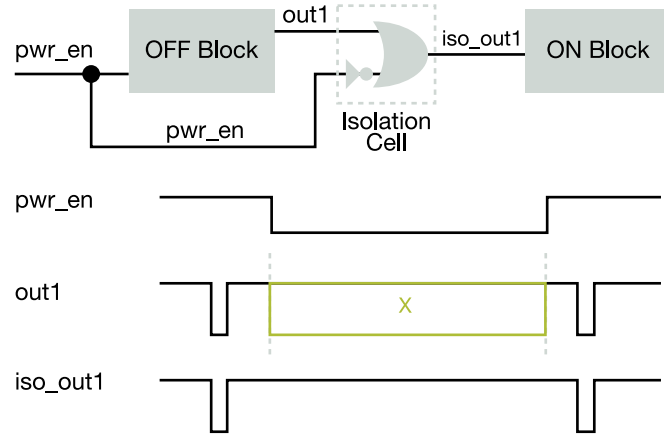

Figure 2

Power sequencing (or power cycle)

Any design with switchable power domains will have a power cycle. Power gating is one of the most effective methods for reducing leakage power. When power is removed from a block inside the SoC, its outputs start floating and need to be isolated to prevent them from corrupting the always ON logic. When power is restored, the isolation is removed.

A power cycle in any low-power design will have the following four stages:

- Power down sequence
- Power OFF state
- Power up sequence
- Power ON state

The above stages are shown in Figure 4.

State Retention Cell Block Diagram and Timing

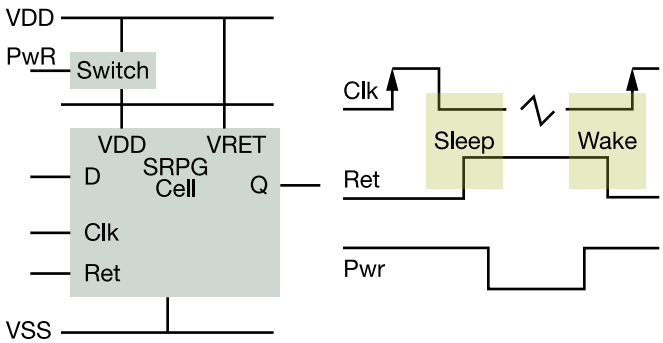

Figure 3

Power Cycle or Sequence Diagram

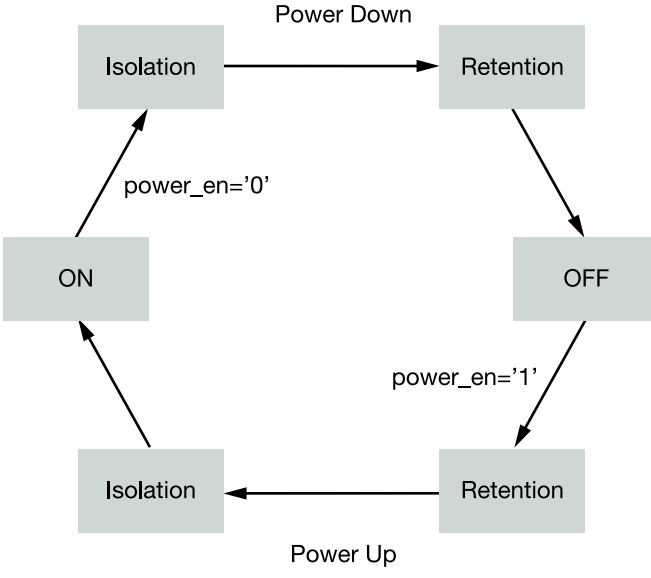

Figure 4

Power down sequence

This is the sequence followed whenever the SoC's normal operating voltage is removed. In such a condition, the design should ensure that signals from an OFF domain are isolated and the states of critical flops/memories are retained before the voltage is completely removed. Here the task of the verification engineer is to check that all required signals are being isolated or retained and the powering down sequence is being followed. If the power is removed before the isolation can be applied, it is possible for unknown values to propagate into the always ON domains. This would increase the leakage current, so it is important to verify that the SoC is not powered OFF erroneously.

Power off state

During this phase of the power sequence, most of the SoC is powered OFF. Only blocks intended to run during the OFF state are consuming power (e.g. from battery). These blocks, which are powered ON, are very sensitive to signals coming from OFF domains. The verification engineer needs to check if all isolation values are being correctly applied and that isolation is not missing on any signal. For example, if a signal is isolated to its active state, it might cause the always ON logic to behave erroneously or even get hung up.

Power up sequence

This is the sequence followed when the SoC power is being ramped up (i.e., the power switch is flipped ON). During this phase, the design needs to ensure that power is stable before the isolation and retention are removed and that the primary I/Os are powered up before the core of the SoC. Here, the verification engineer will have to check the sequence of applying the power, the removal of isolation/retention and that powered ON logic outputs are being correctly driven. The power ON sequence might need to reset the system, so the verification engineer should check that the logic was correctly generated and that no logic is left without a reset.

Power on state

This is the normal operating state of the SoC. The SoC should be able to work the same way it did before it entered the low power mode. The verification engineer should ensure that normal operation can resume after power ON and isolation and retention have been completely removed. The verification engineer should also check that the CPU is able to boot up again and has serviced any interrupts that occurred during the low-power mode before starting normal operations. Any memories that have retained state during low-power mode should be checked for content validity by reading them back.

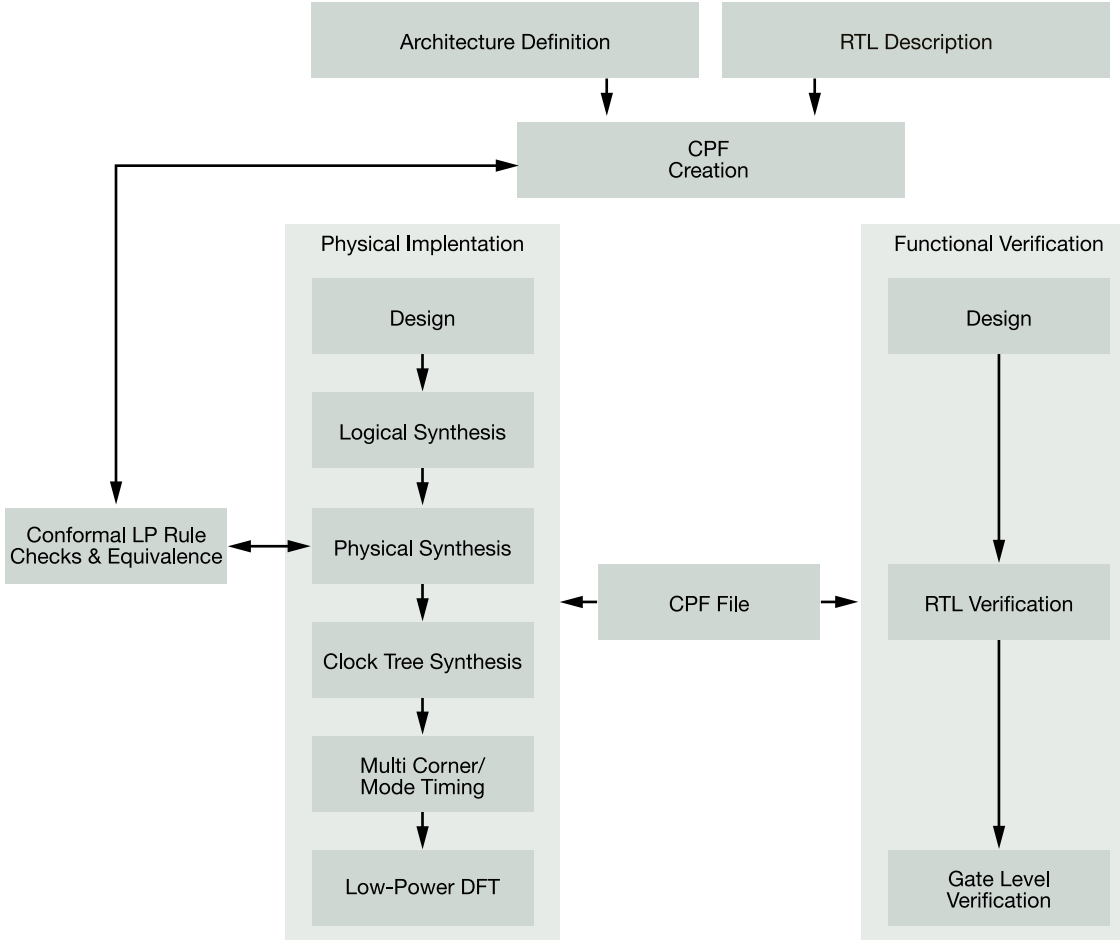

Figure 5

CPF-based Design Flow

What is CPF?

CPF is a file format in which the entire power intent of the design is captured. It is a single file that is used across all stages of design, verification and back-end flows. It is a Tcl-based file that is easy to edit. Most tools can read CPF, thereby ensuring a consistent power intent in the design. The file contains the description of the low-power architecture, which domains can be powered off, which to isolate and which registers to retain. In addition to this information, the power control signals that govern when to retain, isolate and switch power are also contained within the CPF file.

CPF-based design flow

The design flow is slightly modified when CPF comes into the picture. The design's power intent is specified by the systems architect and systems integrator. The CPF file is created during the architecture definition phase along with the contribution from the RTL integrator. This CPF file is the input for all functions in the design cycle. The functional verification uses this file as a reference to check the design for correct power sequencing, isolation and retention. When a simulation is run with CPF file, the same RTL, which does not have any power related implementation, will be simulated as if all power intent is present.

During the physical implementation stage, the CPF is used by the back-end tools to insert power gating logic, actual isolation and retention cells inside the design. This power connected netlist is again verified using CLP, which generates its constraints by reading in the CPF file. This way CLP can be used to verify the physical implementation and the CPF file itself. Gate-level simulations using the CPF file will verify the dynamic behavior of the power connected netlist. See Figure 5 for the modified Design Cycle.

Traditional Verification

Traditional approach and its limitations

To verify the low power logic in the traditional way (without using the CPF) would require the testbench elements to mimic the low-power behavior. In most designs, isolation and retention wrappers are added by the design team. The verification engineer will need to identify all of the registers that are retained and signals that are being isolated to do the simulations. The testbench logic would be needed to retain these registers during low-power mode and then restore the value on power ON. For power OFF, all nodes inside the OFF domains will be required to be driven to 'x'. Now, one can imagine the kind of infrastructure needed to do these simulations. When the design is big, the task becomes overwhelming.

Gate-level simulations, which have the isolation and retention cells inside the netlist, would still require the testbench to model the power OFF behavior and drive 'x' to all OFF nodes. Though circuit level simulations can model power switching, they are late in the design flow, like gate-level simulations, and any bugs found will lead to an ECO just before tape out and delay the process. We cannot afford to take this risk, and the costs involved may be high.

CPF Simulations

CPF-based approach eliminates the drawbacks of the traditional verification and removes the dependency on netlist to perform power-aware verification. With CPF-enabled simulations, no additional testbench or design changes are needed. The simulator reads in the CPF file and simulates the RTL as if actual isolation or retention cells are present. Even the OFF domain internal signals and output ports are driven to 'x' automatically. After power ON, the design should be able to recover from the 'x' state. See Figure 6. This removes that burden from the verification engineer who can then concentrate on finding real issues in the design rather than building the infrastructure for verification, as is the case in traditional approach.

In case a bug is found, a missing isolation for example, the verification engineer does not have to wait for the fix and can simply add the isolation rule in CPF and move ahead.

Another advantage of CPF is that it can be read in by formal tools. The same set of verifications can be done using formal analysis, as in dynamic verification, but much faster and more thoroughly.

Low-power requirements of live project

The design used for CPF simulations had the following low-power features:

- Most of the SoC was powered OFF in standby except for timer modules and on-chip RAM
- A power management block monitored the supply and indicated to the regulator/switch block when the supply was OFF
- The standby regulator compares this supply against the battery supply and controls the switching between the main supply and battery
- The timers block is never disabled and the on-chip RAM is used to preserve critical data needed by processor

The task of the verification engineer was to check the normal operation of timers during multiple transitions in and out of the low-power mode (battery mode) and to ensure that the RAM contents were not lost.

Simulations of live project with CPF and its results

The CPF file for this design had information about the power switching control and blocks that would be switched OFF in low-power mode. Since isolation and retention logic was already added to the SoC by the design team, the CPF file did not include this information. Test cases targeting the timer module and RAM were run with CPF. The verification infrastructure was unchanged except for a few assertions added to check isolation, retention and power sequencing. Some near-fatal bugs have been found using CPF-based low-power verification and are detailed below.

Bugs/issues found

Problem—Signals coming directly from an OFF block into the timers (always ON) block.
Resolution—Logic generating these signals was moved inside the timers block or had isolation put on some signals.

Problem—On-chip RAM became corrupted on power up.
Resolution—The modeling of power-up sequencing was not correct in the power management block and was confirmed in the block's design review.

Power Gating Simulating Example

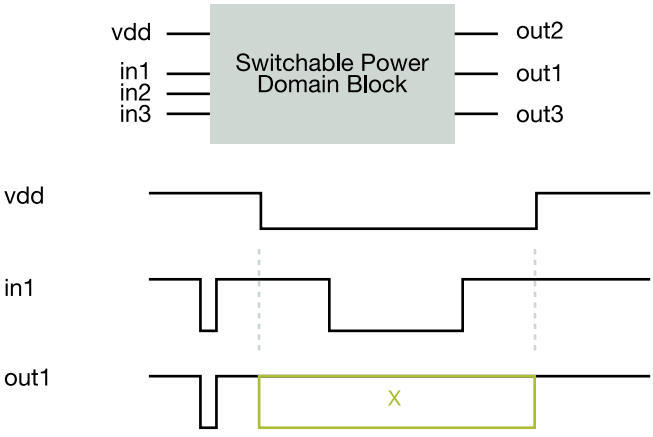

Figure 6

Problem—Modeling issues with analog models of PLL.

Resolution—Since all internal signals are driven to 'x' in CPF simulations, the initial block signals cannot revive from the 'x' situation. These 'x' propagates through the design, corrupting the whole SoC during power-up. The resolution to this problem is to re-code the initial as an assign statement, which can recover from an 'x' situation as it is a continuous assignment. Initial block executes once.

Comparison Between CPF and CLP

Conformal Low Power (CLP) is a powerful tool that can do static checks on the system on chip (SoC) netlist complete with power connections, to verify that the power intent in the netlist matches the specified power properties mentioned in a constraints file and that the netlist has proper power domain crossing with proper level shifting and isolation. However, this tool is limited in certain ways that are overcome by CPF. The differences between the two are elaborated in Table 1.

Based on the Table 1 arguments, we cannot do away with either one of the two. Therefore it is necessary that both CPF and CLP are a part of the design cycle to ensure a complete verification of low-power designs.

Limitations with CPF simulations

CPF as a language supports many low-power constructs, but CPF support in simulator poses some limitations:

- Power connectivity checks cannot be done using CPF. Since only block-level switching information is provided in CPF, power connectivity in design cannot be done
- Level shifter checks cannot be done using CPF because they are related to power connectivity

Conclusion

With new low-power design techniques being developed as designs move toward sub-micron technology nodes, the verification of these techniques has become very important. The traditional approach is cumbersome and might miss many critical power-related scenarios. CPF-enabled simulations (dynamic or static) provide a superior solution to verify these low-power techniques. CPF allows us to find bugs early in the design cycle with very little or no change to testbench infrastructure. This allows the verification engineer to focus on issues rather than trying to model the low-power switching in testbench. CLP does a static check on the design's power intent, while CPF does a dynamic check of the same. But this does not mean CPF can do away with CLP. Both are needed to do a complete power intent verification for a SoC.

References

[1] Noah Bamford, Milind Padhye, "Verification of Low Power Design using CPF"—Freescale Semiconductor (September 2007) [www.cdnusers.org]

[2] Michael Santarini, "Techniques for Low Power ASIC Design"—EDN (24 May 2007) [http://www.edn.com/index.asp?layout=article&articleid=CA6442429]

[3] "Low Power Simulation Guide"—Cadence Docs in 5.83 installation (Jan 2007)

[4] "Power Forward Initiative"—White Paper from Cadence

[5] "CPF Information"—Si2 website (www.si2.org)

Comparison Between CPF and CLP

Common Power Format (CPF)	Conformal Low Power (CLP)
CPF captures the complete power intent. CPF helps develop constraints for CLP.	CLP needs constraints to be manually entered and might have some errors.
CPF touches the entire design cycle.	CLP is confined to netlist verification.
Isolation values can be checked in CPF. Missing isolations can be found.	Isolation values not checked in CLP. Missing isolation can be found.
CPF checks the dynamic nature of low-power design. It can check the power sequencing and operation of SoC after exiting the low-power mode.	CLP is a static check and cannot check the dynamics of low-power design.
CPF cannot be verified for completeness unless a simulation or CLP is run.	CLP reads in the CPF to make its constraints. Hence, it can verify the completeness of CPF file.
CPF simulation speed and coverage is test case specific. CPF with formal analysis can speed up execution and coverage. Netlist simulation using CPF can be slow and memory intensive.	CLP is fast and thorough. Can work on both RTL and netlists.

Table 1

Prashant Bhargava is a design lead for the Microcontroller Solutions Group in Freescale's India Design Center (IDC). He has been with the company since 2006. He has bachelor's degree in electronics.

Ferenc Koplyay

The Orlando 10
Green Ideas Coming to Life

2008 International FTF Design Challenge

Green engineering is here to stay. And, it is up to us to learn how to best apply green engineering to existing technologies and create new applications that epitomize "going green."

To tap into the vast worldwide knowledge bank, Freescale developed the 2008 International Freescale Technology Forum (FTF) Design Challenge, which rewards the most inventive green embedded designs with cash prizes and high-profile recognition. Professional engineers and university students are invited to submit design papers to our panel of judges. The first regional challenge, the FTF Design Challenge Americas, narrowed the submissions down to 10 finalists, with four of the teams made up of college students.

The 10 finalist projects are listed below, in no particular order. Each illustrates how green ideas can come to life.

Clean Water Diverter

The clean water diverter is a simple idea for conserving water that is normally recycled through a water treatment system without serving a useful purpose. Essentially, the system takes the water a user runs from the tap to get hot water, which normally runs down the drain, and diverts it to a holding tank for the toilet.

The clean water diverter was submitted by Mark Donkers from OES, Inc. in London, Ontario, Canada. The clean water diverter bill of materials (BOM) includes:

- Operator interface
- Dual-channel water diverter
- Holding tank monitor
- Single-channel water diverter

Each module communicates through an MC13211 2.4GHz RF transceiver system-in-package (SiP) integrated with an 8-bit MC9S08GT microcontroller (MCU) that performs all logic functions. The system automatically diverts captured water from the holding tank to the toilet tank for flushing.

Digital Sun Distiller

Using energy from the sun to distill potable water from waste water is not what most people would consider a "high tech" enterprise. The distiller is essentially a glass-topped box that gathers condensed water vapor created by the sun's heat. However, instead of gathering the reusable water in jars or bottles, the digital sun distiller can intelligently distribute the water where and when it is needed.

Submitted by Jose Luis Rojas from iluxon S.A. de C.V. in San Luis Potosi, Mexico, the digital sun distiller, with additional sensor technology, can estimate potable water production for a day. If weather variables are programmed into the MCU (such as forecasted hours of sunlight), the digital sun distiller can estimate weekly and monthly water totals. The BOM includes:

- MCF51QE128—Flexis™ 32-bit MCU
- MC13203—2.4 GHz RF transceiver for ZigBee® applications
- MPXV5004—integrated pressure sensor
- MC13213—2.4 GHz RF transceiver SiP with MC9S08GT 8-bit MCU

The digital sun distiller is particularly useful for developing countries, where access to potable water and water management systems may be limited. "The product has been designed with very low-cost components, and it is an attractive alternative in almost any economy," says Jose. "With a small investment, it is possible to create something very helpful for the planet and for society."

ZigBee-based Wireless Sensor Network for Greenhouse Automation

Greenhouse management involves continuous monitoring and activation of different units to maintain a suitable environment for plant growth. Performed manually, these tasks can be time consuming and detrimental to ideal growing conditions. This submission proposes a wirelessly enabled automated system to manage watering systems, fertilizing, CO^2 injection, heating and cooling systems, supplementary lighting and other functions.

The ZigBee-based automation system was submitted by Fernando Gasca, Jaime Sanchez, Jose Maria Ruiz, Marco Tulio Gonzales and Alfredo Cabral from of CODE Ingenieria SA de CV in Toluca, Mexico. It is designed to form one network per greenhouse, each with four network units:

- Dock device, using an MCF5208 ColdFire® MCU to perform most of the processing tasks and network configuration procedures. If a range extender is needed, the dock will work as a ZigBee router
- Remote control, with graphics LCD and capacitive touch keypad controls
- Sensor node, with default in-board temperature, humidity and light sensors and additional analog-to-digital inputs for customer sensors
- Actuator node, with in-board latching power relays for powering diverse units, such as irrigation, CO^2 injection, heating, lighting, etc.

The network architecture follows a star topology, and each unit employs a 2.4 GHz MC13213 RF transceiver with BeeStack ZigBee protocol stack for wireless communication.

Wind Turbine Active Control System

As an alternative source of energy, wind power represents a more environmentally friendly option with the potential to curb pollution while allowing for sustained economical development. To gain the most from wind-generated electricity, an electronic control system for horizontal-axis wind turbines can ensure proper alignment with the wind and correct pitch values are maintained at all times, even under constantly shifting conditions.

The wind turbine active control system submitted by Oscar Roberto Salinas, Aracely Madrid, Diego Rodriguez, Adriin Chaparr and Oscar Ricardo Noi, who are students at Tecnologico de Monterrey in Chihuahua, Mexico, includes the following BOM:

- Flexis 32-bit MCF51QE128 MCU
- Three MMA7260Q accelerometers
- Eight MPXV5004G pressure sensors

The system is designed to monitor wind speed and direction and continually adjust the turbine facing and blade pitch accordingly. For instance, the system processes not only the current wind direction but also the previous direction to predict the wind direction in the next instant. This detection enables smooth and timely changes in orientation. A similar process can be used to adjust the pitch of the rotor blades.

Sunlight Efficiency Detector

The sunlight efficiency detector provides a way to monitor the efficiency of each photovoltaic (PV) module in a PV array to ensure the entire system is operating at peak performance levels. This will eliminate costly and time-consuming manual inspections to hunt down faulty PV modules.

Submitted by Thomas Moulton of FACTS Engineering, LLC, in New Port Richey, Florida, the system involves attaching a low-cost 8-bit MC9S08QG8 MCU to each PV module. Powered by the module itself, the QG8 will use its integrated analog-to-digital converter (ADC) to monitor the module's voltage output, pinpointing panel inefficiencies that can be caused by dirt buildup, poor connections or module degradation. The QG8 can then verify the PV's status by comparing the gathered data with a database of historical information.

The on-chip universal asynchronous receiver/transmitter (UART) will be used to communicate the data to an MPC8313 PowerQUICC™ II microprocessor in a multi-drop network. The processor will periodically poll each QG8 in the entire solar array to read its state, thus allowing for continuous monitoring of the array's performance at the level of its smallest field-replaceable unit.

When asked about his experience with the FTF Design Challenge, Tom said, "It is a unique way to get a concept we had been developing moving forward at a fast pace. The green aspect also attracted our attention because so much can be done in this area, and I feel this is a vastly un-tapped source of new technology. The challenge of competing against our peers for recognition of quality ideas is exciting."

Intelligent Stirling Generator

A Stirling engine is a closed-cycle heat engine that can be used to provide mechanical or electrical power for a number of applications. Theoretically, it can offer perfect efficiency; however, that has proven difficult to achieve. The intelligent Stirling generator submitted by Alejandro Macario of Villa Cuauhtemoc, Mexico, uses thermodynamic principles, special magnetic arrays and electronic components to improve efficiency for such applications as a portable generator.

The project's BOM includes:

- Flexis 8-bit MC9S08QE128 microcontroller
- MPC18730 power management IC
- MC33976 dual-gauge driver
- MMA7260Q accelerometer
- MPXV5004G pressure sensor
- Five temperature sensors

He has basically created a free-piston Stirling engine that generates electricity more efficiently:

- Using magnetic motors dramatically reduces energy loss through internal friction
- Employing Halbach arrays to multiply the magnetic field on the coil generator
- Electronically controlling the engine with a flexible 8-bit microcontroller to maximize performance according to the working conditions

ɔwer

system is most efficient when the sun's rays strike the instrument's surface, whether a PV module or a solar collector, at a perpendicular angle. An electronic solar follower using vertical and horizontal actuators can enable the PV array or the collector to follow the sun across the sky to maintain maximum efficiency.

Pablo Di Guilio from San Francisco, Argentina, and a student at Universidad Tecnologica Nacional FRSFCO, submitted a plan for a simple, low-power solar follower. Key components include:

- Flexis 8-bit MC9S08QE128 microcontroller to calculate optimal positioning for the entire solar power system
- Timer and date clock
- Pulse width modulation (PWM) actuator control

Such a system can improve power yield by as much as 35 percent and can optimize the natural illumination provided by a fiber optic solar collector. "After the FTF Design Challenge I'm planning to continue researching and developing my prototype," says Pablo, "to create a version with better and more efficient performance."

Flexible Fuel Engine Control Unit

The flexible fuel engine control unit (FFECU) adds new capabilities to the standard engine control unit (ECU). Submitted by Erwin Saavedra of ECATES21 in Bucaramanga, Colombia, this project enhances the ECU with:

- The flexibility to scale to almost any size internal combustion engine
- The ability to work with any fuel mix, targeting the incorporation of bioethanol
- The capacity to adapt to changes in environmental factors and to adjust differing engine characteristics, thanks to a neuron-fuzzy (NF) algorithm. A database with engine types is stored to provide initial parameters to the NF network

Controlled by a Flexis 32-bit MCF51QE128 microcontroller, the FFECU incorporates sensors, drivers and other discrete components. The complete solution includes:

- A common closed-loop control algorithm to make real-time adjustments based on inputs from the oxygen sensor
- Knock detector based on an accelerometer
- External memory for adaptive parameters
- Graphical user interface for parameter input

The goal is to create a low-cost adaptive module that can be retrofitted with existing engines without any type of characterization, though the engine may need modifications, such as nickel plating, for ethanol operation.

Gas-saving Automotive Injection System

Engine knocking is simply inconsistent fuel combustion, but it can actually damage an engine. This can be alleviated by injecting water into the air intake to cool the combustion chamber. Water injection systems have been around for some time, but this project goes to another level.

Submitted by a team of University of North Florida students, Kelvin LeBeaux, Harold Rivera and Dangkhoa Ly, the system not only allows the drivers to substitute water for 20 percent or more of the fuel, but also allows precisely controlled combinations of several different fluids to determine the most efficient results. The system is controlled by a 16-bit S12XE microcontroller, using its ADC to convert the analog signals from the system's pressure sensors and the XGATE co-processor for intermediate storage and automatic polling of CAN, serial communication interface (SCI) and serial peripheral interface (SPI) ports.

This injection system will help improve fuel mileage, reduce greenhouse gas emissions and clean carbon build-up in the engine.

Flexible Energy Management and Automation System

The flexible energy management and automation system (FEMAS) is a combination home automation and home energy management solution, combining familiar light switches with outlet replacements and a central controller. It uses a proprietary IEEE® 802.15.4 protocol to provide both lighting and appliance automation and energy usage monitoring.

Submitted by Charles Lord of Triangle Advanced Design and Automation, PLLC in Cary, North Carolina, the heart of the system is the Flexis 8-bit MC9S08JM60 USB microcontroller, which will carry out all day-to-day functions without host PC interaction. The system itself can be easily upgraded to a more powerful Flexis device, the pin-compatible MCF51JM128 ColdFire embedded controller, to provide more computing power.

The FEMAS also includes PC software that makes cost-saving recommendations, including lower wattage bulbs, dimming circuits, turning off unused or unnecessary circuits, etc., making it a complete home energy management system.

Conclusion

These finalists certainly demonstrate that there is an energetic community of designers working on new green ideas for eco-friendly applications. Engineering for energy efficiency is a vibrant, growing field that has an exciting future.

Ferenc Koplyay is responsible for technical content development for global Freescale Technology Forums. He has been a Freescale employee since 1995 and has managed various PowerQUICC product offerings. Prior to Freescale, Ferenc was the Director of Technical Support for Andrew Corporation's Networking Division. He holds a bachelor's degree in computer science from The University of Texas at Austin.

Wireless Connectivity Family
Product Summary

Wireless Connectivity Products

Transceivers

Device	Supply Voltage V	Supply Current @ 1% Duty Cycle (Typ) mA	Standyby Curent (Typ) uA	Frequency Band GHz	Sensitivity @ 1% PER (Typ) dBm	Control Interface	Data Rate (Spec) Kbps	TX/RX Switch	MAC Options	Packages
MC13201FC	2.0–3.4V	30, TX; 37, RX	500	2.4–2.5	-91	SPI	250	Yes	SMAC	32 QFN
MC13202FC	2.0–3.4V	30, TX; 37, RX	500	2.4–2.5	-92	SPI	250	Yes	SMAC, IEEE® 802.15.4, BeeStack™	32 QFN

System in a Package

Device	Supply Voltage V	Supply Current @ 1% Duty Cycle, CPU @ 2 MHz (Typ) mA	Standyby Curent (Typ) mA	Frequency Band GHz	Sensitivity @ 1% PER (Typ) dBm	Data Rate (Spec) Kbps	TX/RX Switch	MAC Options	Packages	Flash	RAM	Core	Interfaces and Peripherals
MC13211	2.0–3.4V	31.1, TX; 38.1, RX	0.2–0.675	2.4–2.5	-92	250	Yes	SMAC	71 LGA	16 KB	1 KB	HCS08	I²C, SCI (2), Timer/PWM(2), KBI, 8-ch. 10-bit ADC, Up to 32 GPIO
MC13212	2.0–3.4V	31.1, TX; 38.1, RX	0.2–0.675	2.4–2.5	-92	250	Yes	SMAC, IEEE 802.15.4	71 LGA	32 KB	2 KB	HCS08	I²C, SCI (2), Timer/PWM(2), KBI, 8-ch. 10-bit ADC, Up to 32 GPIO
MC13213	2.0–3.4V	31.1, TX; 38.1, RX	0.2–0.675	2.4–2.5	-92	250	Yes	SMAC, IEEE 802.15.4, BeeStack	71 LGA	60 KB	4 KB	HCS08	I²C, SCI (2), Timer/PWM(2), KBI, 8-ch. 10-bit ADC, Up to 32 GPIO

Platform in a Package

Device	Supply Voltage V	Supply Current @ 1% Duty Cycle, CPU @ 1 MHz (Typ) mA	Standby Current (Typ) uA	Frequency Band GHz	Sensitivity @ 1% PER (Typ) dBm	Data Rate (Spec) Kbps	Data Rate TurboLink™ (Spec) Kbps	Tx/RX Switch	MAC Options	Packages	Flash	RAM	ROM	Core	Interfaces and Peripherals
MC13224V	2.0–3.4V	“20 TX; 20 RX”	0.3–0.675	2.4–2.5	-95	250	No	Yes	SMAC, IEEE 802.15.4, BeeStack	99 LGA	128 KB	96 KB	80 KB	ARM7 TDMI-S	12-bit ADC, Buck Converter, I²C, SPI, SCI, SSI/I2S, 4–16-bit Timers/ PWM(2), KBI, 2–8ch. Up to 64 GPIO

Digital Signal Controller Family
Product Summary

DSC Products

Device	MHz	Flash (KB)	RAM (KB)	I²C	SCI/QSCI	SPI/QSPI	CAN	PWM	PWM Fault Inputs	12-bit ADC	Analogue Comparator	12-bit DAC	Quad Decoder	16-bit Timers	Package	Temp Range		Additional Perihperals
																-40°C to +105°C	-40°C to +125°C	
MC56F8011	32	12	2	1	1 x SCI	1 x SPI	—	1 x 6-ch.	4	2 x 3-ch.	—	—	—	4	32LQFP	Y	N	On-chip relaxation oscillator
MC56F8013	32	16	4	1	1 x SCI	1 x SPI	—	1 x 6-ch.	4	2 x 3-ch.	—	—	—	4	32LQFP	Y	N	On-chip relaxation oscillator
MC56F8014	32	16	4	1	1 x SCI	1 x SPI	—	1 x 5-ch.	3	2 x 4-ch.	—	—	—	4	32LQFP	Y	N	On-chip relaxation oscillator
MC56F8023	32	32	4	1	1 x QSCI	1 x QSPI	—	1 x 6-ch.	4	2 x 3-ch.	2	2 (Internal)	—	4	32LQFP	Y	N	On-chip relaxation oscillator
MC56F8025	32	32	4	1	1 x QSCI	1 x QSPI	—	1 x 6-ch.	4	2 x 4-ch.	2	2 (Internal)	—	4	44LQFP(.8p)	Y	N	On-chip relaxation oscillator
MC56F8036	32	64	8	1	1 x QSCI	1 x QSPI	1	1 x 6-ch.	4	2 x 5-ch.	2	2 (Internal)	—	4	48LQFP	Y	N	On-chip relaxation oscillator
MC56F8037	32	64	8	1	2 x QSCI	2 x QSPI	1	1 x 6-ch.	4	2 x 8-ch.	2	2	—	4	64LQFP	Y	N	On-chip relaxation oscillator
MC56F8322	60	48	12	—	2 SCI	2 SPI	1	1 x 6-ch.	1	2 x 3-ch.	—	—	1 x 4-ch.	8	48LQFP	N	Y	On-chip relaxation oscillator
MC56F8323	60	48	12	—	2 SCI	2 SPI	1	1 x 6-ch.	3	2 x 4-ch.	—	—	1 x 4-ch.	8	64LQFP	N	Y	On-chip relaxation oscillator
MC56F8335	60	80	12	—	2 SCI	2 SPI	1	2 x 6-ch.	4 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	128LQFP	N	Y	
MC56F8345	60	144	12	—	2 SCI	2 SPI	1	2 x 6-ch.	4 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	128LQFP	N	Y	
MC56F8346	60	144	12	—	2 SCI	2 SPI	1	2 x 6-ch.	3 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	144LQFP	N	Y	
MC56F8347	60	144	12	—	2 SCI	2 SPI	1	2 x 6-ch.	4 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	160LQFP	N	Y	
MC56F8355	60	280	20	—	2 SCI	2 SPI	1	2 x 6-ch.	4 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	128LQFP	N	Y	
MC56F8356	60	280	20	—	2 SCI	2 SPI	1	2 x 6-ch.	3 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	144LQFP	N	Y	
MC56F8357	60	280	20	—	2 SCI	2 SPI	1	2 x 6-ch.	4 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	160LQFP	N	Y	
MC56F8365	60	560	36	—	2 SCI	2 SPI	2	2 x 6-ch.	4 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	128LQFP	N	Y	
MC56F8366	60	560	36	—	2 SCI	2 SPI	2	2 x 6-ch.	3 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	144LQFP	N	Y	
MC56F8367	60	560	36	—	2 SCI	2 SPI	2	2 x 6-ch.	4 + 4	4 x 4-ch.	—	—	2 x 4-ch.	16	160LQFP	N	Y	

8-bit Product Family
Product Summary

General-Purpose Products

HCS08 and RS08 Families

Device	Flash	RAM	USB	ADC Channels 10-bit	ADC Channels 12-bit	SCI (UART)	ESCI	SPI	I²C	ACMP	Timer	Clock Type	Package DFN/QFN	Package QFP/LQFP	Package TSSOP	Package SOIC	Package DIP	Dev Tools DEMO	Dev Tools EVB	Applications/Additional Features *All RS08, S08 and HC08 Products Include COP, LVI, POR and KBI
MC9S08AW60	60 KB	2 KB		16		2		√	√		6 + 2-ch.	ICG w/FLL	48	64, 44				√		High integration, flash programmable to 5V
MC9S08AW32	32 KB	2 KB		16		2		√	√		6 + 2-ch.	ICG w/FLL	48	64, 44				√		High integration, flash programmable to 5V
MC9S08AW16	16 KB	1 KB		16		2		√	√		4 + 2-ch.	ICG w/FLL	48	64, 44				√		High integration, flash programmable to 5V
MC9S08AC128	128 KB	8 KB		16		2		2	√		6 + 6 + 2-ch.	ICG w/FLL	48	80, 64				√		High integration, flash programmable to 5V
MC9S08AC96	96 KB	6 KB		16		2		2	√		6 + 6 + 2-ch.	ICG w/FLL	48	80, 64				√		High integration, flash programmable to 5V
MC9S08AC60	60 KB	2 KB		16		2		1	√		6 + 2 + 2-ch.	ICG w/FLL	48	64, 44, 32				√		High integration, flash programmable to 5V
MC9S08AC48	48 KB	2 KB		16		2		1	√		6 + 2 + 2-ch.	ICG w/FLL	48	64, 44, 32				√		High integration, flash programmable to 5V
MC9S08AC32	32 KB	2 KB		16		2		1	√		6 + 2 + 2-ch.	ICG w/FLL	48	64, 44, 32				√		High integration, flash programmable to 5V
MC9S08AC16	16 KB	1 KB		8		2		√	√		6 + 2-ch.	ICG w/FLL	48	44, 32				√		High integration, flash programmable to 5V
MC9S08AC8	8 KB	1 KB		8		2		√	√		6 + 2-ch.	ICG w/FLL	48	44, 32				√		High integration, flash programmable to 5V
MC9S08GB60A	60 KB	4 KB		8		√		√	√		3 + 5-ch.	ICG		64				√	√	High performance, flash programmable down to 1.8V
MC9S08GT60A	60 KB	4 KB		8		√		√	√		2 + 2-ch.	ICG	48	44			42	√	√	High performance, flash programmable down to 1.8V
MC9S08GB32A	32 KB	2 KB		8		√		√	√		3 + 5-ch.	ICG		64				√	√	High performance, flash programmable down to 1.8V
MC9S08GT32A	32 KB	2 KB		8		√		√	√		2 + 2-ch.	ICG	48	44			42	√	√	High performance, flash programmable down to 1.8V
MC9S08JM60	60 KB	4K	2.0		12	2		2	√	√	1 x 2-ch., 1 x 6-ch.	MCG w/ PLL and FLL	48	64, 44				√		USB S08 device with high performance and integration
MC9S08JM32	32 KB	2K	2.0		12	2		2	√	√	1 x 2-ch., 1 x 6-ch.	MCG w/PLL and FLL	48	64, 44				√		USB S08 device with high performance and integration
MC9S08JM16	16 KB	1K	2.0		12	2		2	√	√	1 x 4-ch., 1 x 2-ch.	MCG w/PLL and FLL	48	44, 32				√		USB S08 device with high performance and integration
MC9S08JM8	8 KB	1K	2.0		12	2		2	√	√	1 x 4-ch,. 1 x 2-ch.	MCG w/PLL and FLL	48	44, 32				√		USB S08 device with high performance and integration
MC9S08QG8	8 KB	512B		8		√		√	√	√	2-ch., MTIM	ICS	8, 16, 24		16	8	16	√		High performance, low voltage, small package
MC9S08QG4	4 KB	256B		8		√		√	√	√	2-ch., MTIM	ICS	8, 16, 24		16	8	16, 8	√		High performance, low voltage, small package
MC9S08QA4	4 KB	256 KB		4						√	1-ch., MTIM	ICS	8			8	8	√		Low-end S08 device with 1.8V to 3.6V op range
MC9S08QA2	2 KB	160 KB		4						√	1-ch., MTIM	ICS	8			8	8	√		Low-end S08 device with 1.8V to 3.6V op range
MC9RS08KA8	8 KB	254B		12					√	√	1 x 2-ch., MTIM	ICS				8	8	√		More integration, new RS08 core with small MCUs
MC9RS08KA4	4 KB	126B		12					√	√	1 x 2-ch., MTIM	ICS				8	8	√		More integration, new RS08 core with small MCUs
MC9RS08KA2	2 KB	62B								√	MTIM	ICS	6			8	8	√		Ultra-low end, new RS08 core for small MCUs
MC9RS08KA1	1 KB	62B								√	MTIM	ICS	6			8	8	√		Ultra-low end, new RS08 core for small MCUs
MC9S08GT16A	16 KB	2 KB		8		√		√	√		3 + 2-ch.	ICG	48, 32	44			42	√	√	High performance, flash programming down to 1.8V
MC9S08GT8A	8 KB	1 KB		8		√		√	√		3 + 2-ch.	ICG	48, 32	44			42	√		Flash programming down to 1.8V, small package
MC9S08QD4	4 KB	256B		4							2 + 3-ch.	ICS				8	8	√		Low-end, flash programmable to 5V
MC9S08QD2	2 KB	128B		4							2 + 3-ch.	ICS				8	8	√		Low-end, flash programmable to 5V
MC9S08LC60	60 KB	4 KB			8	√		2	√	√	2 x 2-ch.	ICG w/FLL		80, 64				√		Integrated Liquid Crystal Display (LCD) driver with high segment count
MC9S08LC36	36 KB	2.5 KB			8	√		2	√	√	2 x 2-ch.	ICG w/FLL		80, 64				√		Integrated Liquid Crystal Display (LCD) driver with high segment count
MC9S08SH32	32 KB	1KB		16		√		1	1	√	2 x 2-ch, MTIM	ICS			28, 20,16	28		√		Low-end S08 device with 2.7V to 5.5V op range
MC9S08SH16	16 KB	1KB		16		√		1	1	√	2 x 2-ch, MTIM	ICS			28, 20,16	28		√		Low-end S08 device with 2.7V to 5.5V op range
MC9S08SH8	8 KB	512B		12		√		√	√	√	2 + 2-ch.	ICS						√		Low-end S08 device with 2.7V to 5.5V op range
MC9S08SH4	4 KB	256B		12		√		√	√	√	2 + 2-ch.	ICS	24		20, 16	8	20	√		Low-end S08 device with 2.7V to 5.5V op range
MC9S08QE128	128 KB	8 KB			24	2		2	2	2	1 + 6-ch., 2 + 3-ch.	ICS	48	80, 64, 44				√	√	Ultra-low-power S08 device with 1.8V to 3.6V op range
MC9S08QE64	64 KB	4 KB			24	2		2	2	2	1 + 6-ch., 2 + 3-ch.	ICS	48	80, 64, 44, 32				√	√	Ultra-low-power S08 device with 1.8V to 3.6V op range
MC9S08QE8	8 KB	512B			10	1		1	1	2	2 + 3-ch.	ICS w/FLL		32	16	20, 28	16	√		Ultra-low-power S08 device with 1.8V to 3.6V op range
MC9S08QE4	4 KB	256B			10	1		1	1	2	2 + 3-ch.	ICS w/FLL		32	16	20, 28	16	√		Ultra-low-power S08 device with 1.8V to 3.6V op range

32-bit ColdFire® Product Family
Product Summary

Device	Freq. (MHz)	MIPS @ Max Freq.	Cache (KB)	SRAM (KB)	Flash (KB)	DMA	GPT*	PIT	10/100 FEC	Crypto	USB	CAN	I²C	UART/USART PSC	SPI	ADC	Other	GPIO Max	Package
MCF51QExx	50	46	—	8	32, 64, 128	—	1 x 6-ch., 16-bit, 2 x 3-ch.,16-bit	—	—	—	—	—	2	2 SCI	2	12-bit	Ultra-Low-Power 2x KBI (8-ch.)/ 2x ACMP	54	LQFP 80, LQFP 64
MCF51JMxx	50	46	—	16	64, 128	—	1 x 6-ch., 16-bit, 1 x 2-ch., 16-bit	—	—	1	Full-Speed Device/Host/ OTG	1	2	2 SCI	2	12-bit	8x KBI, 1x ACMP	66	LQFP 80, QFP 64, LQFP 64, LQFP 44
MCF51ACxxx	50	46	—	16, 32	128, 256	—	1x 2-ch., 16-bit	—	—	—	—	1, 2	2	2 SCI	2	24-ch., 12-bit	2x 6-ch., 16-bit FTM, 2x ACMP, CRC, COP	70	LQFP 80, QFP 64 LQFP 64
MCF540x	40, 54, 166	50, 159	4K I, 8K I/D	8, 16	—	2-ch., 16-ch.	2-ch., 16-bit, 4-ch., 32-bit	2	1	—	—	—	1	2, 3	QSPI	—	—	30, 50	QFP 160, LQFP 144, MAPBGA 196
MCF521x	66, 80	76	2K I/D	16, 32, 64	128, 256, 512	4-ch.	4-ch., 16-bit, 4-ch., 32-bit, 8-ch., 16-bit	2, 4	—	—	—	1	1	3 UART	QSPI	10-bit, 12-bit	—	44, 56, 142	LQFP 64, MAPBGA 81, MAPBGA 256
MCF522xx	50, 66, 80	46, 76	—	4, 8, 16, 32	64, 128, 256	4-ch.	4-ch., 32-bit, 4-ch., 16-bit	2	—	—	Full-Speed Device/Host/ OTG	—	1, 2	2 UART, 3 UART	QSPI	12-bit	RTC w/32 kHz Osc	35, 46, 52	MAPBGA 81, LQFP 100, QFN 64, LQFP 64
MCF5223x	50, 60	46, 57	—	32	128, 256	4-ch.	4-ch., 16-bit, 4-ch., 32-bit	2	1	—	—	1	1	3 UART	QSPI	12-bit	EPHY, RTC	43, 76	LQFP 80, LQFP 112, MAPBGA 121
MCF5227x	120, 160	114, 159	8K I/D	128	—	16-ch.	4-ch, 32-bit	2	—	—	Full Speed Device/Host/ OTG	1	1	3 UART	DSPI	Touchscreen, 12-bit	SPI Boot Flash, Crossbar Switch	47, 55	LQFP 176, MAPBGA 196
MCF523x	80, 100, 150	144	8 KB I/D	64	—	4-ch.	4-ch., 32-bit	4	1	—	—	1, 2	1	3 UART	QSPI	—	16-ch. eTPU 32-ch. eTPU	79	QFP 160, MAPBGA 256
MCF52xx	120, 140	107, 125	8 KB I	96, 128	—	4-ch.	2-ch., 16-bit	—	—	—	High-Speed On-the-Go	2	2	2 UART 3 UART	QSPI	12-bit	IDE, Audio, I²S, USB OTG (FS)	49, 60, 64	LQFP 144, MAPBGA 160, MAPBGA 225
MCF5253	140	125	8 KB I	128	—	4-ch.	2-ch., 16-bit	—	—	—	High-Speed Device/OTG	2	2	3 UART	QSPI	12-bit	IDE, I²S	60	MAPBGA 225
MCF527x	66, 100	63, 96	8 KB I/D, 1K I	4, 64	—	2-ch., 4-ch.	4-ch., 16-bit, 4-ch., 32-bit	4	1	—	Full-Speed Device	—	1	2 UART 3 UART	QSPI	—	PLIC, TDM Soft HDLC	48, 78	QFP 160, MAPBGA 196
MCF527x	133, 166	159	16 KB I/D	64	—	4-ch.	4-ch., 32-bit	4	1, 2	—	Full-Speed Device	—	1	3 UART	QSPI	—	—	53, 73. 74	MAPBGA 196, MAPBGA 256
MCF528x	66, 80	76	2 KB I/D	64	256, 512	4-ch.	4-ch., 32-bit, 8-ch., 16-bit	4	1	—	—	1	1	3 UART	QSPI	10-bit	—	150	MAPBGA 256
MCF5307	66, 90	75	8K U	4	—	4-ch.	2-ch., 16-bit	—	—	—	—	—	1	2 UART	—	—	—	16	FQFP 208
MCF532x	240	211	16K I/D	32	—	16-ch.	4-ch., 32-bit	4	1	—	Full Host, Full OTG	1	1	3 UART	QSPI	—	SVGA LCD	64, 94	MAPBGA 196, MAPBGA 256
MCF537x	180, 240	158, 211	16K I/D	32	—	16-ch.	4-ch., 32-bit	4	1	—	Full Host, Full OTG	—	1	3 UART	QSPI	—	—	46, 62	QFP 160, MAPBGA 196
MCF5407	162, 220	316	16K I, 8K D	4	—	4-ch.	2-ch., 16-bit	—	—	—	—	—	1	1 UART, 1 USART	—	—	—	16	FQFP 208
MCF5445x	166, 200	256, 308	32K I/D, 32K D	32	—	16-ch.	4-ch., 16-bit	—	1, 2	—	Full Host, Full OTG	—	1	3 UART	DSPI	—	SSI, SBF, PCI, ATA	93, 132	MAPBGA 256, PTEPBGA 360
MCF547x	200, 266	308, 410	32K I, 32K D	32	—	16-ch.	4-ch., 16-bit	2	1, 2	—	High-Speed Device	—	1	4 PSC	DSPI	—	PCI	77, 93	PBGA 388
MCF548x	166, 200	256, 308	32K I, 32K D	32	—	16-ch.	4-ch., 16-bit	2	1	—	High-Speed Device	2	1	4 PSC	DSPI	—	PCI	77, 93	PBGA 388

It's the best connected 15-year-old in the MCU industry.

ColdFire® Architecture. The perfect 32-bit compilation.

The ColdFire Architecture has a 15-year heritage steeped in innovation. With more than 100 individual devices, ColdFire has one of the broadest 32-bit portfolios in the industry. Whether it's USB, Ethernet, CAN or ZigBee®, ColdFire connects and controls millions of networked products in homes, offices and factories around the world. ColdFire also enables your design to keep pace with today's eco-trends thanks to its industry-leading, ultra-low-power 32-bit MCU. And to make its own ecosystem even stronger and more vibrant, we are now openly licensing ColdFire V1. Find out what ColdFire can do for you.

Meet ColdFire at freescale.com/coldfire101 and save 50% on the JM Demonstration Kit.*

freescale™ semiconductor

* While supplies last.