

White Paper



Green Embedded Computing and the MPC8536E PowerQUICC® III Processor

Introduction

Growing energy demands from embedded electronics and increasing evidence of dramatic global climate change are generating greater environmental and cultural pressure for green solutions in embedded computing applications. This, coupled with continued expectations for higher performance embedded computing with each new product generation, despite environmental concerns, is impacting the future usage models of embedded processing systems.

Although traditional offline powered equipment, such as appliances, HVAC and lighting systems, dominate electric equipment energy consumption, embedded electronics and online equipment, such as printers, storage, networking infrastructure and data centers, are increasingly consuming a larger share of our energy resources. To balance the performance required for powerful new electronic applications with rising concerns over energy consumption, environmentally aware “green” movements and government regulations and programs are driving manufacturers to develop intelligent strategies for optimizing performance within specific energy budgets.

Traditional embedded computing platforms have been designed for maximum work load with little regard to the cyclical work profile across hourly, daily, weekly or extended time intervals. However, new generation high-performance systems are shifting from power provisioning to energy efficiency across varying workloads.

The U.S. Environmental Protection Agency's (EPA) ENERGY STAR initiative is a good example of a governmental conservation strategy that captures the priority shift from power and performance to energy efficiency. In the early phase, the ENERGY STAR requirements were focused on two-staged power management, that being running and not running. More recently, ENERGY STAR initiatives have concentrated on energy budgets based on typical usage profiles.

For example, office automation equipment, such as printers, are allotted fixed weekly kilowatt limits based on advertised print speed. To acquire an ENERGY STAR qualification, equipment manufactures must match the energy limits across the advertised page rates as the equipment transitions through its advertised power usage profile—standby, management and printing modes. The printer is a good example of cyclical workload, as it tends to spend much more time in a ready-to-print state or performing low-workload management services than it does for higher energy consumption printing states. However, other embedded applications can employ similar profiling to reduce the energy waste and costs, too, including home network gateways, industrial processing plants and telecommunications systems.

Growing demand for high-performance embedded computing poses significant challenges for equipment manufacturers' energy efficiency priorities. Incorporating media content processing and Web services into embedded systems is driving up CPU processing requirements into gigahertz-level clock rates. In addition, rich media content requires high-speed networking and fast memory interfaces.

Lowering power consumption, in the face of these increasing performance demands, requires advanced energy management schemes from new product development engineers. A simple strategy to design lower power consuming electronics begins to address the green embedded computing challenge. However, larger gains will come from creating flexible systems that can pace workload with energy consumption in an intelligent and efficient manner.

System Implications of Green Computing: Design Considerations


Quiescing the system

CPUs, ASICs and I/Os all contribute to the overall energy efficiency of a green computing system. It is important for the system designer to understand that different areas of the system may require a disproportionate share of power throughout the workflow. Understanding that will allow him or her to model the system from an energy consumption perspective, employing more elegant energy savings techniques, such as turning off different ASICs and I/Os or portions of the processor to pace energy consumption with the workload.

In most cases, all the work performed in embedded computing applications is done in cycles—a combination of active states, management states and dormant states that are dynamically administered to most effectively optimize energy-efficient performance on demand. This is true for such applications as high-speed printing, home routers and all sorts of WAN managed systems.

- Active state: All I/Os are running and the core is operating at full speed. Performance and connectivity are maximized.
- Management state: The CPU is not needed at full speed and some I/Os are unused. Performance is minimized and connectivity is maximized.
- Dormant state: The system does not require CPU processing, and many I/Os can be turned off. Performance is zero and connectivity is minimized.

Fig 1: Embedded Computing System-Level Work-Flow States



The processor

Today, there is connectivity everywhere you turn. What's important to remember is that the world is not just connected, it's always connected. This "net effect" is why carefully managing the different performance cycles, pacing energy to fulfill application needs without losing the connection, is so important. The processor is the heart of the system and must be network-aware in all three system states, even in the dormant state, because dormant does not mean dead.

During the embedded application design process, an intimate understanding of the application's work flow is a critical prerequisite for energy pacing. There will be times when high frequencies will be required, but these will likely be very short cycles in the work flow. The vast majority of the time, the processor is being asked to perform relatively low-performance tasks and can do so as easily at 600 MHz as it can at 1400 MHz. Cutting processor frequencies during these management periods saves watts, which can in turn be used by the ASIC or I/O that are working harder. Throttling down the processor for the balance of the work flow, tailoring the power requirements for each cycle, will lead to a far more efficient embedded computing application. Remember, however, that during frequency throttling it may be necessary to change PLL ratios without a system reset to maintain the "always" connected status.

Understanding the workflow needs along with the code profiles also allows more efficient use of the processor's low-power modes, such as doze, nap, sleep and deep sleep. Low-power modes are key energy efficiency components that work in conjunction with frequency throttling.

Devices connected to processor


Throughout the workflow, energy demands vary among system components. Some devices need more power than others, and power sharing management must be tightly controlled.

Big gains in energy efficiency can be realized by optimizing the dormant state. With fast recovery from the dormant state, external devices to the processor can be powered off without compromising system objectives. However, signal and leakage contention are special design considerations for managing powered-on and powered-off devices.

Software requirements

To ensure fast recovery from management or dormant states it is important that software saves previous known states in non-volatile memory so the processor can retrieve those states upon entering a more active state. The various functions that the OS performs before the system enters a dormant state are outlined in Figure 2.

Fig 2: System States



Thermal considerations

Generally, the cooler a system runs, the less power it consumes. This is critical during the active mode when the system is working the hardest. So, a thorough understanding of the thermal properties of the embedded design will help the system developer design for more efficient power management. The developer must answer such questions as:

- What is the expected airflow in the system?
- Can the airflow be regulated based on system state? For example, enabling a fan during an active state or slowing down or disabling a fan during a management state.
- What are the workload changes that cause the chip to work harder?

Because of the advanced power management features and the low nominal platform and core voltages, Freescale's MPC8536E processor is a gigahertz-plus SoC that can be thermally controlled without a fan.

A Freescale Solution—MPC8536E PowerQUICC III Processor

The MPC8536E PowerQUICC III processor is a highly integrated system-on-chip (SoC) for green embedded computing applications. It supports low-power states nap, doze, sleep and deep sleep and incorporates advanced power management features to help OEM product developers meet government efficiency initiatives, such as ENERGY STAR, Top Runner and Energy Using Products (EuP).

Wake on important network event

To provide overall processor energy efficiency, it's important not only to minimize power in low-power modes, such as sleep and doze, but also to maximize useful work performed in an active state and reduce the time penalty incurred when transitioning between states.

In an embedded networked application, the system spends much of the time in a low-power mode and wakes up in response to an external event. If the system takes too long to wake up, the window for acting on the event that caused the wake up may have closed.

Lossless packet operation in a networked environment is a method for ensuring that critical packets initiate the wake-up sequence and that no targeted packets are lost. A common example is wake-on-ARP (address resolution protocol). Office networks typically have traffic 24 hours a day, even when the offices are not occupied. If the controlling processor wakes up on every packet, it will be constantly awake and have no time to enter a low-power mode. Nevertheless, there are times when the processor needs to wake up to service and process certain packets.

An ARP packet can find a host's hardware address, given its network layer address is known. When the system receives an ARP packet that is destined for it alone, it triggers a wake-up to respond as per protocol specifications.

The MPC8536E PowerQUICC processor can wake on any targeted network event. It can enter deep sleep mode where the system is dormant, yet the enhanced three-speed Ethernet (eTSEC) still operates and no packets are ignored. At the same time, DDR is in self-refresh mode, but it can still be accessed if needed. The eTSEC's Receive Filter is configured to drop packets that don't need to be processed, but packets that need processing, such as ARP packets destined for the correct address, are written to DDR, and the eTSEC wakes the system from deep sleep for processing. The eTSEC Receive Filter can be programmed to accept and wake on whatever packets are interesting for a particular system usage configuration.

This is the best of both worlds—operating at ultra-low power the vast majority of the time, yet with no penalty of reduced functionality (dropped packets) because the system can wake and respond as needed.

Split power planes

To support low-power deep sleep modes while still supporting lossless packet operation, the processor requires split power planes. Power must be supplied to parts of the die, such as the eTSEC, which needs power to operate its Receive Filter to analyze incoming packets. However, the vast majority of the die is unused in deep sleep, so to minimize total power (both static and dynamic), voltage can be removed from that portion of die.



Separate power planes enable these powered and unpowered portions to coexist on the same chip. The power planes have separate pins, separate (or isolated) power supplies on the system, and their power distribution networks are routed completely independently within the SoC. The MPC8536E processor has two power planes—one for the e500 core and L2 cache arrays and the other for the rest of the SoC.

Dynamic power gating

Dynamically turning portions of the die on or off must be coordinated through both the SoC and the system, telling them when it is safe to apply and remove power and to indicate when this process has been completed. On the MPC8536E processor, this is controlled by the power management controller (PMC).

When entering deep sleep mode, the PMC deasserts POWER_EN output to indicate that internally the SoC is prepared for the voltage to be removed and that the voltage regulator on the system can remove power. When exiting deep sleep, the PMC asserts POWER_EN to indicate that the voltage regulator needs to apply power. After voltage ramp-up, the voltage regulator asserts POWER_OK input to the SoC to indicate that voltage has ramped up and is stabilized. If, however, the voltage regulator does not have this capability, counters within the PMC can be programmed with values sufficient to guarantee that power has stabilized.

Figure 3 illustrates the timings of the PMC from WAKE_EVENT to first DDR access. Note that POWER_EN to POWER_OK is system design dependant and POWER_OK to READY is programmable to anything greater than 100 μ s.



Isolation cells within the SoC

Isolation cells are required whenever there are logic paths between power domains that don't always operate simultaneously (i.e., there may be a powered-down power domain). Isolation cells ensure that unknown values (X values) do not propagate from a powered down domain into an operational domain.

Isolation cells can also be used to ensure that the system does not waste power driving logic from a powered domain into an unpowered domain. The isolation cells themselves must remain powered at all times as well as the PMC logic used to control the enablement of the isolation cells.

The MPC8536E processor has isolation cells as described above between the core power domain, containing e500 core and L2 cache arrays, and the platform power domain, which includes eTSEC and other peripherals.

Dynamic frequency scaling

In many applications, high performance during periods of activity should be balanced with low power consumption when there is less workload. Microprocessor cores typically operate at higher maximum frequencies than the rest of the SoC. Therefore, power consumption can be best minimized by controlling core frequency. Software can dynamically increase or decrease the core's clock frequency while still allowing the rest of the SoC to continue operating at the previous (fixed) frequency.

The MPC8536E processor implements this process, known as dynamic frequency scaling, or JOG, which can be used to alter the frequency of the e500 core. There are two JOG options:

- Prior to entering deep sleep, the core can be configured to wake up in a different core frequency. There is no overhead to wake from deep sleep at a frequency different from that which was used to enter deep sleep. This is useful, for instance, if upon waking up from deep sleep the system will have to wait for mechanical motors or other functions to warm up before it performs any useful tasks. During that wait, there is no need for the core to be running at a high frequency.
- The other option is to perform a frequency change without waiting for an external deep sleep event to wake up. Just configure the new desired core frequency and then issue the frequency change. For example, after the core has performed some intensive data processing, it may not be completely idle, but maximum processing power is not required. Or, if new incoming work will be coming soon, power can be conserved in the interim.

Saving energy with the MPC8536E PowerQUICC III processor

To demonstrate how low-power modes and advanced power management features can reduce overall system energy consumption, consider an office networked print system that has two workload patterns, 9 a.m. to 5 p.m. (33% of each day) and 5 p.m. to 9 a.m. (67% of each day).

From 9 a.m. to 5 p.m., the print system has the following workload characteristics each hour:

Time per hour	System Task	System State	MPC8536E Mode	% of time from 9 a.m. to 5 p.m. in this mode
5 minutes	Print (3 jobs)	Active	Run	8.3%
20 minutes	Management Services	Management	JOG	33.3%
35 minutes	Network Standby	Dormant	Packet-Lossless Deep Sleep	58.3%

For this example we will say the processor's total power consumption in deep sleep mode is 0.7W, while in JOG mode at 600 MHz it is 5.5W, and in active mode at 1400 MHz it is 8.0W. Average power consumption in this time period is:

$$8.3\% \times 8.0W + 33.3\% \times 5.5W + 58.3\% \times 0.7W = 2.91W$$

From 5 p.m. to 9 a.m. the print system is largely idle and in deep sleep mode, responding only to ARP packets and other packets that are destined for the print system. Even though the system is idle, the network will be moving, on average, multiple packets per second, even if the office is empty. Most of these packets are not targeting the print system.

For this example, packets that are targeting the print system, which need to be serviced, arrive, on average, once every 60 seconds. It takes 100 ms to service each of these packets, including voltage ramp-up and ramp-down times. Therefore, every 60 seconds the system spends 0.1s in active mode (which is 0.17% of the time) and 59.9s in dormant mode (99.83% of the time). Average power consumption in this time period is:

$$0.17\% \times 8.0W + 99.83\% \times 0.7W = 0.712W$$

Combining these two workload patterns over a 24-hour period, the overall average power is:

$$33.3\% \times 2.91W + 66.6\% \times 0.712W = 1.44W$$

In comparison, a legacy system without green embedded computing features would need to continually operate at the maximum frequency, consuming a constant 8.0W. Therefore, the green strategy in this example would use just 18 percent of the power of such a legacy system.

Conclusion

Embedded computing applications are all around us, everywhere we go. Designers are severely challenged to continue feeding the industry with increased product performance while adhering to constantly shrinking energy budgets.

In this paper we have explored how energy-saving features integrated on the MPC8536E PowerQUICC III processor can significantly reduce energy consumption in a single printing application. Multiply these savings by the millions of embedded applications across the globe and we can see just how much green embedded computing can reduce the energy load on power grids worldwide, providing a positive economic impact, saving resources and reducing the cost to our environment.

How to Reach Us:

Home Page:

www.freescale.com

Power Architecture Information:

www.freescale.com/powerarchitecture

e-mail:

support@freescale.com

USA/Europe or Locations Not Listed:

Freescall Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
1-800-521-6274
480-768-2130
support@freescale.com

Europe, Middle East, and Africa:

Freescall Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:

Freescall Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064, Japan
0120 191014
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescall Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:

Freescall Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
1-800-441-2447
303-675-2140
Fax: 303-675-2150
LDCForFreescallSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescall Semiconductor products. There are no express or implied copyright license granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescall Semiconductor reserves the right to make changes without further notice to any products herein. Freescall Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescall Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescall Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescall Semiconductor does not convey any license under its patent rights nor the rights of others. Freescall Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescall Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescall Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescall Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescall Semiconductor was negligent regarding the design or manufacture of the part.



Freescall and the Freescall logo are trademarks or registered trademarks of Freescall Semiconductor, Inc. in the U.S. and other countries. All other product or service names are the property of their respective owners. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
© Freescall Semiconductor, Inc. 2009.
Document Number: MPC8536EWP
REV 1

