



# Intel<sup>®</sup> Technology Journal

The Original 45nm Intel Core™ Microarchitecture

## Power Improvements on 2008 Desktop Platforms

# Power Improvements on 2008 Desktop Platforms

Paul Zagacki, Client Platform Architecture and Planning, Intel Corporation  
Vidoot Ponnala, Client Platform Architecture and Planning, Intel Corporation

Index words: low power, platform power management, energy efficiency, desktop power, desktop platform

Citations for this paper, Paul Zagacki, Vidoot Ponnala “Original 45nm Intel® Core™2 Processor Performance” Intel Technology Journal. <http://www.intel.com/technology/itj/2008/v12i3/7-paper/1-abstract.htm> (October 2008).

## ABSTRACT


Idle and low-utilization platform power management have been key deliverables for mobile platforms for many years. The resulting platforms based on Intel® Centrino® technology have delivered increasing performance and capabilities while continuing to increase the overall battery life of the mobile platform. These same principles have gained importance also in desktop platforms as corporations strive to reduce the cost of deploying platforms in support of their efforts to address global climate change and deliver more energy-efficient computing. Further, the Energy Star specification for computers was adopted on July 20, 2007 [1], adding idle power targets to desktop platforms. The Energy Policy Act, adopted by Congress on July 27, 2007 [2] now requires that federal agencies buy equipment that is Energy Star qualified (effectively making Energy-Star compliance mandatory for some percentage of desktop platforms).

In this paper we present platform- and silicon-component power data that demonstrate advances in desktop platform power management enabled on 2008 platforms, built with the Intel Q45 Express Chipset, originally referred to by the codename “Eaglelake,” and Intel processors based on the original 45nm Intel® Core™2 Quad microarchitecture, originally referred to by the codename “Yorkfield/Wolfdale.” For instance, a 2008 desktop platform with more advanced power management can demonstrate (when correctly configured) a 16-percent reduction in AC idle power when compared with the same platform enabled with 2007 platform power-management techniques. In addition to providing an overview of how technologies such as deeper C-states or Serial ATA link power management (familiar in mobile platforms) dramati-

cally reduce silicon and platform power, we also demonstrate the effectiveness of these technologies in desktop platforms, from operating-system power-management settings to devices such as USB keyboard mouse or multimedia card readers. Finally, we make recommendations on how to configure a desktop platform (hardware and software) to make the most of the new features found on our 2008 platforms.

## INTRODUCTION


Desktop platforms are experiencing an evolution in what is seen as their primary value to consumers. Until recently, the quest has been exclusively directed at achieving higher and higher workload performance at lower and lower cost to the end user. However, driven by concerns over the environment and greenhouse gas emissions and the increasing importance of Energy Star compliance [1,2], the focus is shifting from the exclusive pursuit of performance and cost improvements to energy-efficient performance (EEP) [3,4]. EEP is the intersection of performance or capabilities with the delivery of those capabilities, using the least amount of energy. In simple terms, for many workloads, this can mean getting the job done as fast as possible and then getting the system into an idle or low-power state. Most of a system’s time over the course of a day is spent at or near idle utilization and power levels, as demonstrated in Figure 1.



**Figure 1: System power levels during a 9-hour workday following the EEP 2.0 Model demonstrate that 75 percent of time is at or very near idle utilization.**


The EEP 2.0 workday, the source of the data in Figure 1, assumes several hours of work represented by runs of the productivity benchmark Sysmark 2007 (which has idle time built into it), followed by idle in the form of an employee break, and by some sleep time on longer breaks [3,4].

Considering all the time the system is less than 10 percent used in any given workday, and that the system is likely to be idle or asleep overnight (unless it is in a batch environment), it is important to focus our attention on idle and low-utilization power improvements. The fact that most desktop platforms are lightly loaded most of the time coupled with the Energy Star Version 4.0 [5] focus on idle power is key to understanding the motivation behind many of the platform power-management features implemented on our 2008 desktop platforms. Idle is defined as the average state of the platform after the operating system (OS) has loaded and the system has been given adequate time to quiesce any activities (15 min in the Energy Star Version 4 specification). Without getting too deep into the topics of ACPI and OS power management, readers are encouraged to reference the ACPI 3.0a specification and the Windows\* power-management whitepaper for more background information on idle and processor power states [6]. These are shown in Figure 2.



**Figure 2: Processor/platform power states from the ACPI 3.0a specification.**

In order to better understand what techniques are required to deliver the most energy-efficient desktop platforms, it is important to first understand where all the power goes in a typically configured desktop platform. Throughout, we focus our attention on a typically configured Intel® vPro™ desktop platform with integrated graphics. (Figure 3 shows a typical platform hierarchy with the key components.)




**Figure 3: Key components in a 2007 corporate platform built with Q35 Express Chipset GMCH and ICH9.**

We focus on integrated graphics primarily because discrete graphics cards cover a very wide range of power, and the performance extracted from this additional power is not typically necessary for office- or productivity-type applications such as Outlook\*, PowerPoint\* or Excel\* (while energy efficiency is quickly becoming a key attribute for these platforms).

All of the components and their interfaces shown in Figure 3 are powered either directly from the silver box power supply (as is the case for peripherals such as serial advanced technology attachment, or SATA, drives) or through voltage regulators (VRs) built into the desktop motherboard. For example, in the current platform generation, the processor has at least three individual power rails (core, front-side bus (FSB), and phase locked loop (PLL) supplies); the graphics and memory controller hub (GMCH) may have up to five supply voltages (many of which are shared across various regulators); the I/O controller hub (ICH) has seven individual rails; and dual data rate (DDR) memory requires three rails. All of these component supplies are derived from either a dedicated 12-V processor rail or from the 12-V/3.3-V/5-V rails out of the silver box power supply utilizing about fourteen separate regulators (more if the platform supports the manageability engine, a key component of Active Management Technology). With all of this voltage regulation on a desktop motherboard, it is clear that power delivery is one of the biggest sources of efficiency losses in a platform. Some components may have three voltage regulation steps from AC to final DC supply, prior to the voltage being seen by the silicon. Each such stage loses a little power in the translation.

The pie chart in Figure 4 demonstrates the areas that we need to focus on to maximize energy efficiency in desktop platforms. Power delivery conversion losses, major platform silicon components (processor, GMCH, and ICH silicon), and SATA hard drives are among the primary power consumers in a desktop platform.



**Figure 4: Approximate component power consumption (including losses) for a 45-W AC idle platform.**

Throughout the remainder of this paper we describe several power innovations on the corporate desktop platform based on the Intel Q45 Express Chipset and the latest Intel processors based on original Intel Core™2 Quad, 45nm microarchitecture. We explore innovations in efficiency through the use of phase shedding on the processor VR. We also look at the impact on idle and active power from deploying deeper C-states than those deployed in previous-generation Intel Q35 chipset-based platforms. We demonstrate some common issues that USB devices present for platform power management and explore solutions and best known practices already in use on mobile platforms. Further, we explore how to get the most out of these improvements with the right OS configuration settings.

## ARCHITECTURE

For the purposes of our research for this paper, we used a single Intel vPro desktop platform with a 45nm Intel Core 2 Quad (Yorkfield) processor with a 1333 FSB at 2.83 MHz, a Q45 GMCH with DDR3, and the ICH-10 I O controller hub. This platform is very similar to that shown for the previous-generation platform in Figure 3. We utilized knobs built into a debug BIOS to allow us to enable the platform power features for the previous-generation platform (Intel Q35) and then enabled all the features that are new to the Intel Q45 Express Chipset platform. Using a single platform to create a baseline for the previous generation and then demonstrating the improvements in power management on the Intel Q45 Express Chipset platform helps to remove variation in all the other components between the systems.

## DESKTOP POWER INSTRUMENTED REFERENCE PLATFORM

Many of the power measurements referenced throughout the remainder of this paper were taken on a specially designed and instrumented Intel vPro desktop reference validation platform (RVP). As mentioned earlier, the processor has three power rails, the GMCH has five rails, the ICH has seven rails, and the system memory has three individual power rails. All of these individual voltage rails are instrumented with low-loss, high-accuracy sense resistors and voltage sensing points. All this instrumentation is fed into a high-speed data acquisition card for real-time display and offline analysis. This gives us visibility into the DC power required by each of the major components on the platform along with information on the current state of the platform. The C-state information (refer to Figure 2) is also measured by logging the coordination

signals used between the processor, GMCH, and ICH to set up entry and exit points from these C-states.

The AC power consumed by the platform is measured by using an AC power meter that connects to the wall power outlet at one end and to the silver box power supply at the other end. The data collection was done with a Windows Vista\* SP1 OS using the balanced mode of native OS power-management techniques under various workload conditions (although we focused mostly on idle as our workload).

## **INTEL CORE 2 QUAD PROCESSOR FAMILY**

### **Power features**

The Yorkfield processor is the first four-core desktop processor family, based on Intel's 45nm silicon process technology, that doubles transistor density while providing major improvements in switching leakage power. Beyond the improvements in energy efficiency gained from 45nm process technology, the Yorkfield processor family supports a key additional feature beyond what was implemented in earlier steppings. Core power states below C2 (stop grant) are now supported on the 2008 desktop platform providing a significant hook for components in the platform to opportunistically manage their power.

A core power state such as deeper sleep (C4) is a platform-level decision, so we needed all the key silicon components in the 2008 platform to have access to the feature. The Yorkfield processor family initiates the C4 request for the OS's idle handler through either an I/O read to a specific location, or through an MWAIT instruction. The Q45 GMCH and ICH10 then coordinate the details for the rest of the platform, and the ICH10 asserts the appropriate signal to indicate that the platform is entering C4. The key thing to understand here is that C4 is a state that comprises a coordinated effort among all the major silicon components in the platform.

The C4 power state allows the processor to drop to a very low voltage while still maintaining all the processor's state information. This reduction in core voltage produces a dramatic reduction in transistor leakage, the primary component of idle power. From a behavioral standpoint, the primary differences between C4 and C2 are that the processor no longer responds to snoop requests in C4, because the core voltage is too low to service them and the latency to exit C4 is somewhat longer (about 60 uSeconds). If there is any activity the processor must respond to, the chipset will initiate an exit from C4, in some cases exiting to C2 to


respond to a snoop request and then quickly re-entering C4.

In addition to the core voltage reduction for the processor on entry into C4, the processor also has the ability to tune the efficiency of its own power delivery through the use of a power status indicator (PSI#) signal. This signal is asserted on entry into C4 by the processor and consumed by the core voltage controller, thereby allowing the core VR to tune the power-delivery efficiency to match the processor's expected current consumption.

### **VOLTAGE REGULATOR DOWN 11.1—POWER STATUS INDICATOR (PSI#)**

In the pie chart shown in Figure 4, a significant amount of power is shown as processor VR losses. This loss comes from converting 12 V coming out of the power supply to the voltage being requested by the processor (typically between 0.8 V and 1.2 V). A typical desktop processor VR is split into 2–4 individual phases, depending on the maximum current requirements, and it is most efficient in the middle of its supported current range (0–90 A plus in a 3-phase design). The efficiency drops off very quickly at light loads, effectively making the processor appear to draw more power than it really requires to operate in that mode.

In the 2008 desktop platform we tackle this light-load efficiency problem with an optional feature for Voltage Regulator Down (VRD) 11.1 controllers that allow the processor to give the VR an indication of the current demand it expects over a period of time. When the processor asserts PSI# on entry into C4, the VR can turn off phases of the voltage regulation to improve the power-delivery efficiency. The chart in Figure 5 shows the improvement in efficiency of a typical reference design achieved by the use of the PSI# signal. The improvement in efficiency available with a PSI#-enabled VR design translates to a 20–30 percent improvement in the processor's power consumption under idle conditions.



**Figure 5: Improving power efficiency at light loads on VRD 11.1 controllers.**

## INTEL® Q45 EXPRESS CHIPSET

### Power optimizations

Similar to the Yorkfield processor family described above, the Intel Q45 Express Chipset family has many optimizations to enhance its power-efficient performance. The Intel Q45 Express Chipset is also on a new process technology (a 65nm derivative vs. a 90nm technology used for Intel Q33 in the 2007 platforms). This process technology allows for scaling of leakage and dynamic power through support of a lower core voltage (from a typical 1.25 V to 1.1 V), reduced gate leakage characteristics, and reduced switching capacitance on the newer process technology. The Intel Q45 Express Chipset memory controller also provides improved support for DDR3, which uses less voltage (1.5 V vs. 1.8 V) and power for similar frequency DDR2 technology.

Beyond the process technology shift for the Intel Q45 Express Chipset family and support for a more power-efficient DDR3, there are architectural innovations tied to the implementation of C4 that we consider briefly here.

First, the memory controller in Intel Q45 Express Chipset, in response to a C4 entry, takes advantage of the longer expected idle periods to put memory into a self-refresh state. On a DDR3 device—a dual in-line memory module (DIMM) consists of many devices—this can mean a third of the current consumption requirement versus the best possible device state supported on the previous-generation chipset. Considering that a DIMM can be made up of eight to sixteen individual devices, each consuming an average

of approximately 34 mA per device, one third less current can mean a savings of hundreds of mWs per DIMM.

On entry into C4, the Intel Q45 Express Chipset also reduces its own idle power consumption over the previous generation on the platform by dynamically powering off memory DLL circuits, tristating memory I/O logic buffers, and powering off host, memory, and PCI Express internal clocks.

## RESULTS

All the power-saving techniques described in the previous sections implemented on the 2008 desktop platform contribute to a more energy-efficient platform. Figure 1 shows that a system spends most of its time at or near idle, and power saved under these conditions is critical to achieve energy efficiency. In this section we present DC component and AC platform power data to demonstrate the power-management improvements of the 2008 platform. We also explore the impact of USB devices, such as the keyboard and mouse, on platform power management. Finally, we present some best known methods for getting the most out of the new platform power-management features.

The 2008 platform achieves significant idle and low utilization power savings from its support of deeper C-states (up to C4 state versus the 2007 desktop platform which supports only C2). The support for deeper C-states allows enabling of additional power-management techniques for major components in the desktop platform, ultimately leading to a 16-percent reduction in platform power (varies from platform to platform depending on the components in the platform) in idle and a 5-percent reduction when running Sysmark 2007.

Figure 6 illustrates the AC power savings achieved as well as the DC savings at the component level under idle conditions. Note that there is nearly a 60-percent power savings for the Yorkfield processor, mostly attributed to a dramatic reduction in core voltage when in C4. Beyond the processor itself, since C4 is a synchronized effort between major silicon components in the platform, the Intel Q45 Express Chipset achieves a 28-percent power improvement. This DC savings will vary from processor to processor and chipset to chipset due to natural distributions of static current and the way C4 voltage settings are optimized by manufacturing on a part-by-part basis. The system memory achieves a 60-percent power improvement, because the chipset puts the memory in a self-refresh state as previously explained. These component-level power savings, along with the efficiency improvements on

VRD11.1 power delivery (described earlier), add up to an approximate 16-percent power improvement at the platform level.




Figure 6: Platform and component power savings.

The shift in C-states residency between 2007 and 2008 platforms is illustrated in Figure 7. Generally, C-state residency numbers give an indication of how much time as a percentage the platform spends in each C state. The 2007 platform has the ability to utilize only C2, and Figure 7 shows that the platform spends 99 percent of its time in C2 under idle conditions. Much of that time spent in C2 converts directly to C4 time on the 2008 platform, but not all of it. On the 2008 platform, Figure 7 shows a remainder of 11 percent in C2 and C4 state residency of 88 percent. This change in the distribution of deep C-state residency numbers results in the considerable power savings illustrated in Figure 6. A 100-percent C4 state residency is not achievable because of certain break events such as interrupts, wake up events caused by drivers, or because of the polling architecture of USB devices.

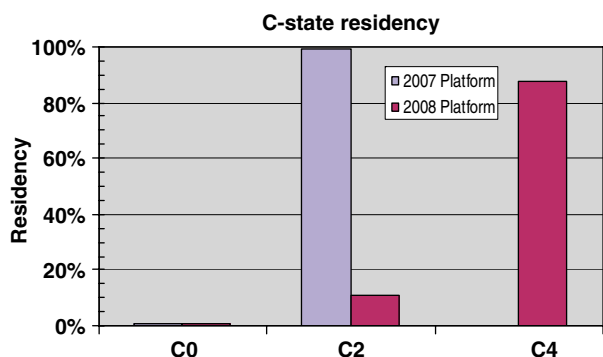


Figure 7: C-state residency.

Power savings on Sysmark 2007 on a suite-by-suite basis and on average are illustrated in Figure 8, which shows a 5-percent improvement in AC platform power

on average for the 2008 platform. Average here is the arithmetic mean of the power consumed by each component of the benchmark.

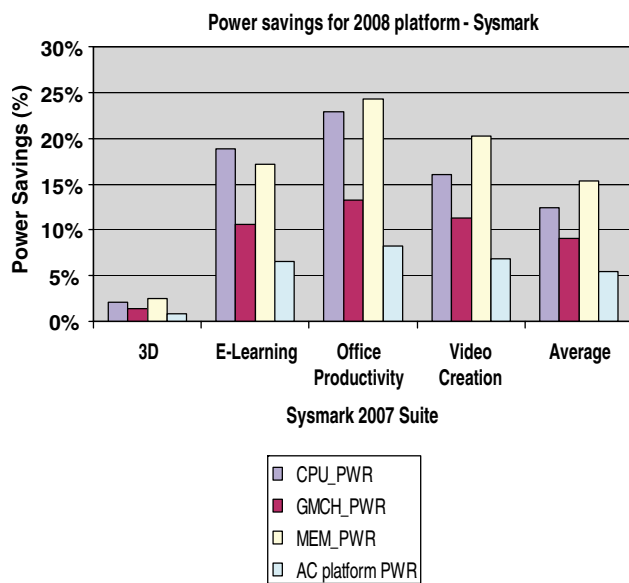


Figure 8: Sysmark power savings.

### USB IMPACT ON PLATFORM POWER AND C-STATE RESIDENCY

An additional power-savings opportunity can be had by enabling the USB selective suspend feature in the OS [7]. Selective suspend is a low-power mode defined in the USB 2.0 specification [8], that allows the USB hub driver to turn off USB ports when they are in idle. Figure 9 illustrates these power savings. Observe that the Yorkfield processor shows a 76-percent power savings in these conditions (an additional 16 percent over what was shown in Figure 7). This, along with the Intel® Q45 Express Chipset’s 34-percent improvement and the system memory’s 73-percent power savings, help to achieve a 22-percent power improvement at the AC platform level.

This improvement in the component power is also reflected in the C-state residency data with USB selective suspend enabled, as shown in Figure 10. C4 state residency now is just over 98 percent compared to 88 percent achieved before without USB selective suspend enabled, and this translates into the power savings seen in Figure 9.




Figure 9: Platform and component power savings with USB selective suspend enabled.

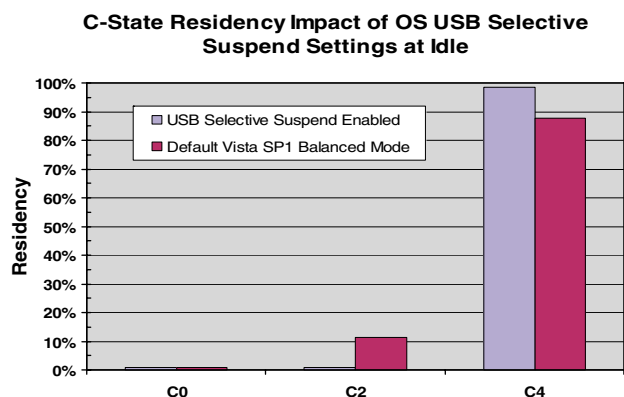


Figure 10: Effect of USB selective suspend on C-state residency.

Enabling this feature in the OS also helps to increase the active power savings as illustrated in Figure 11. For Sysmark 2007 we see a 7-percent improvement in the AC platform power on average as compared to the 5-percent improvement seen earlier (refer to Figure 8).

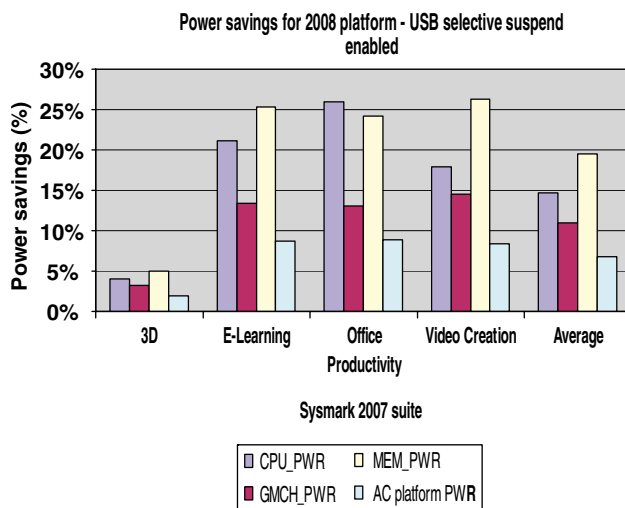



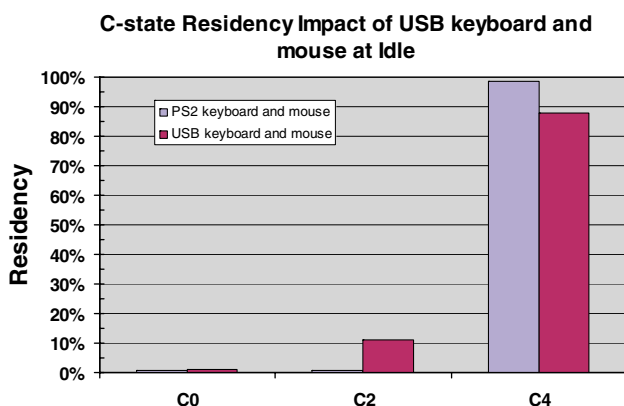
Figure 11: Sysmark power savings with USB selective suspend enabled.

In addition to the USB hub device behavior, USB devices such as the keyboard and mouse have an adverse impact on the platform power management because of the polling architecture of USB devices [9]. USB devices cannot initiate a transfer or transfer data without being polled by the host first. Because of this, the processor has to constantly come out of the C4 state just to poll the USB devices and check whether they have any data to transfer, thereby incurring a power penalty, especially in idle conditions. So while a 76-percent power savings was achieved by using a PS2 keyboard mouse and by enabling USB selective suspend in the OS, only a 69-percent power savings for the processor is observed when a USB keyboard mouse are connected to the platform with USB selective suspend enabled, and similar behavior is observed for the other components as well. As a result, only a 16-percent power savings is seen for the AC platform power. Figure 12 illustrates the power savings seen when the PS2 and USB keyboard and mouse are connected to the 2008 platform.



**Figure 12: Power savings with PS2 and USB Keyboard/Mouse—USB selective suspend enabled.**

This change in power savings can also be seen from the C-state residency numbers as illustrated in Figure 13. There is clearly a change in the C-state residency when a USB keyboard and mouse are added to the system—88-percent C4 residency and 11-percent C2 residency compared to the 98.4-percent C4 residency and 0.8-percent C2 residency seen with a PS2 keyboard mouse. This large shift in the C2 state residency explains the reduced power savings seen at the platform and component level.



**Figure 13: C-state residency with USB keyboard/mouse—USB selective suspend enabled.**

Figures 6 and 9 demonstrate how the 2008 platform is improved with respect to platform power management and along with Figure 12 establish that we get the most power savings when we use the 2008 platform with the USB selective suspend feature enabled in the OS and with a PS2 keyboard and mouse.

## CONCLUSION

The new-generation desktop systems built with the Eaglelake chipset along with the Yorkfield/Wolfdale processors provide significant power-management techniques that help to drastically reduce power at the platform and component level. This paper has illustrated the power savings that can be achieved by using power-management features in the new platform and illustrates how users can take full advantage of the potential of the system to build more energy-efficient platforms.

## ACKNOWLEDGEMENTS

We thank the following individuals for their valuable contributions: Sanjeev Jahagirdar, Intel Mobile CPU Architecture; Barnes Cooper, Intel Mobile Platform Group; Jeff Krieger and Loc Mai, Intel Platform Systems Technology Group; Mike Derr, Karthi Vadivelu, and Tuong Trieu, Intel Client Components Group.

## REFERENCES

- [1] <http://www.energystar.gov/>.
- [2] [http://energy.senate.gov/public/\\_files/ConferenceReport0.pdf](http://energy.senate.gov/public/_files/ConferenceReport0.pdf).
- [3] [http://www.intel.com/technology/eep/index.htm?iid=tech\\_ee+body\\_eep](http://www.intel.com/technology/eep/index.htm?iid=tech_ee+body_eep).
- [4] <http://download.intel.com/technology/eep/overview-paper.pdf>.
- [5] [http://www.energystar.gov/ia/partners/prod\\_development/revisions/downloads/computer/Computer\\_Spec\\_Final.pdf](http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/Computer_Spec_Final.pdf).
- [6] <http://www.microsoft.com/whdc/system/pnppwr/powermgmt/ProcPowerMgmt.mspix>.
- [7] <http://msdn.microsoft.com/en-us/library/ms793200.aspx>.
- [8] <http://www.usb.org/developers/docs/>.
- [9] Barnes C. *et al.* Making USB a more Energy-Efficient Interconnect. Intel Technology Journal, Vol. 12, No. 1, 2008.

## AUTHORS' BIOGRAPHIES

**Paul Zagacki** is a Principal Engineer in the Client Platform Architecture and Planning Group within the Digital Enterprise Group at Intel. Paul holds a B.S. degree in Computer Science from the University of Michigan. He joined Intel in 1994 and has worked on performance modeling for microprocessor architectures, software and benchmark analysis and optimiza-

tion, design convergence, and desktop component and platform power optimization. Paul has four issued patents. His e-mail is paul.zagacki at intel.com.

**Vidoot R. Ponnala** received his B.Tech degree in Electrical Engineering from CVR College of Engineering, Hyderabad, India and his Masters degree in Computer Engineering from the University of Wisconsin, Madison. He joined Intel as a system architect intern in Intel's Mobility Group in 2006 and currently is a Platform Power Architect in Intel's Digital Enterprise Group. Since joining this group in 2008 he has been responsible for platform power and performance analysis for client systems. His current interests include processor and systems architecture. His email is vidoot.r.ponnala at intel.com.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Any codenames featured in this document are used internally within Intel to identify products that are in development and not yet publicly announced for release. For ease of reference, some codenames have been used in this document for products that have already been released. Customers, licensees, and other third parties are not authorized by Intel to use codenames in advertising, promotion or marketing of any product or services, and any such use of Intel's internal codenames is at the sole risk of the user.

Intel's trademarks may be used publicly with permission only from Intel. Fair use of Intel's trademarks in advertising and promotion of Intel products requires proper acknowledgement.

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

Any code names featured in this document are used internally within Intel to identify products that are in development and not yet publicly announced for release. For ease of reference, some code names have been used in this document for products that have already been released. Customers, licensees, and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

Intel Corporation uses the Palm OS<sup>®</sup> Ready mark under license from Palm, Inc.

LEED - Leadership in Energy & Environmental Design (LEED<sup>®</sup>)

Copyright © 2008 Intel Corporation. All rights reserved.

Additional legal notices at:  
<http://www.intel.com/sites/corporate/tradmarx.htm>.