



Intel® Technology Journal

Intel® Centrino™ Mobile Technology

**Mobile PC Platforms Enabled with
Intel® Centrino™ Mobile Technology**

Mobile PC Platforms Enabled with Intel[®] Centrino[™] Mobile Technology

Gordon Chinn, Mobile Platforms Group, Intel Corporation
Sanjiv Desai, Mobile Platforms Group, Intel Corporation
Eric DiStefano, Mobile Platforms Group, Intel Corporation
Krishnan Ravichandran, Mobile Platforms Group, Intel Corporation
Shreekant (Ticky) Thakkar, Mobile Platforms Group, Intel Corporation

Index words: Bluetooth[™], Intel[®] Centrino[™] mobile technology, Intel[®] WCS, WLAN, battery, coexistence, electronic cooling, notebook, power, validation, wireless

ABSTRACT

Mobile PC platforms enabled with Intel[®] Centrino[™] mobile technology incorporate innovative technologies to significantly increase the mobility of notebook PCs. In this paper, we describe the following innovations:

- How to reduce interference between wireless communication technologies using the Intel[®] Wireless Coexistence System (Intel[®] WCS).
- Two techniques for significantly extending average platform battery life.
- Simple thermal solutions for thinner and lighter form factors.

Finally, we outline a comprehensive platform-level validation process that helps deliver a stable platform with Intel quality.

INTRODUCTION

Platforms enabled with Intel Centrino mobile technology incorporate innovative technologies to significantly increase the mobility of notebook PCs. This paper describes innovations in reducing interference between wireless communication technologies, extending average platform battery life, and in driving a thermal solution for thinner and lighter form factors.

When multiple wireless technologies such as Bluetooth^{*}, Wireless LAN, and Wireless WAN are embedded in the same mobile system, the risk of interference between these radio frequency (RF) technologies greatly increases because of their close proximity to each other. In this paper we discuss general technical RF challenges for mobile platforms and describe innovative techniques used on platforms enabled with Intel Centrino mobile technology. These techniques mitigate interference by using the Intel Wireless Coexistence System (Intel WCS).

We also discuss new and advanced platform power-management features for improved battery life. Two new techniques are described: asynchronous voltage regulator (VR) control with a power status indicator (PSI), and multiphase Intel[®] Mobile Voltage Positioning (IMVP).

Platforms enabled with Intel Centrino mobile technology also incorporate state-of-the-art techniques to enable compact, lightweight thermal solutions in smaller form factor systems. The design includes an efficient remote heat exchanger, an air flow set aside for system cooling, and system venting, that together cool the processor and the system as a whole, using only a single fan. The latest interface materials, heat pipe technology, and fans provide the required performance in a small package.

Finally, Intel validation best-known methods (BKMs) are applied to ensure that these platforms deliver optimum

[®] Intel and Intel Centrino are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

^{*} Other names and brands are the property of their respective owners

[®] Intel is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

system performance with multiple wireless technologies operating concurrently within the enterprise as well as when connected to hotspots outside the enterprise.

WIRELESS TECHNOLOGIES

Platforms enabled with Intel Centrino mobile technology incorporate innovative radio frequency (RF) technologies to significantly increase the mobility of notebook PCs. These technologies include Wireless Personal Area Networking (WPAN), e.g., Bluetooth; Wireless Local Area Networking (WLAN), e.g., 802.11a and 802.11b; and Wireless Wide Area Networking (WWAN), e.g., General Radio Packet Service (GPRS). Integrating multiple wireless technologies greatly enhances connectivity to the enterprise Intranet and the Internet, as well as to peripherals such as PDAs, printers, and headsets.

Radio Frequency Challenges in Mobile Notebook PCs

When incorporating RF technologies into a digital computing environment such as a notebook PC, traditional system designers are faced with new challenges previously relegated to RF engineers. Antennas must be added for each frequency band. Since notebook PCs generally consist of many metallic components such as framing structures, hard drives, displays, antenna radiation patterns can be greatly distorted. This distortion can potentially cause significant RF performance variations, as the notebook PC's physical orientation is varied relative to the location of the intended communicating device such as an access point (AP), a peripheral such as a printer, and even another notebook PC.

These RF technologies consist of very sensitive receivers and potentially high-power transmitters. When multiple RF technologies are embedded (co-existent) in the same notebook PC, the risk of interference between these technologies is greatly increased as the transmitters of one radio can couple with the sensitive receivers of another radio, even if they operate in different frequency bands. Technologies such as Bluetooth and 802.11b have even greater interference risks since they both operate in the same 2.4GHz frequency band. Antenna isolation techniques mitigate some of these interference issues, but with multiple antennas and decreasing form factor sizes, optimal placement of antennas is not always possible. New interference-mitigating technologies are then required. For example, Intel Centrino mobile technology includes the Intel Wireless Coexistence System (Intel WCS), which significantly mitigates the interference between 802.11b and Bluetooth technologies.

Intel Wireless Coexistence System

While antenna isolation provides some interference mitigation between 802.11b and Bluetooth radios, performance is still impacted to some degree. For example, 802.11b data throughput is degraded by Bluetooth interference, even with 40dB of antenna isolation, as shown by the lower curve in Figure 1. The upper curve shows the ideal throughput when no Bluetooth interference is present.




Figure 1: 802.11b throughput with Bluetooth interference

To further mitigate interference between Bluetooth and 802.11b, Intel WCS was developed as one of the Intel Centrino-enabling technologies. Intel WCS consists of a combination of antenna isolation techniques, a channel exchange (Figure 2), and priority signaling between an Intel PRO/Wireless Network Connection 802.11 solution and a third-party Bluetooth module. Phase 1 of Intel WCS has been implemented: it mitigates Bluetooth interference and restores 802.11b data throughput nearly completely as shown by the chart in Figure 1.

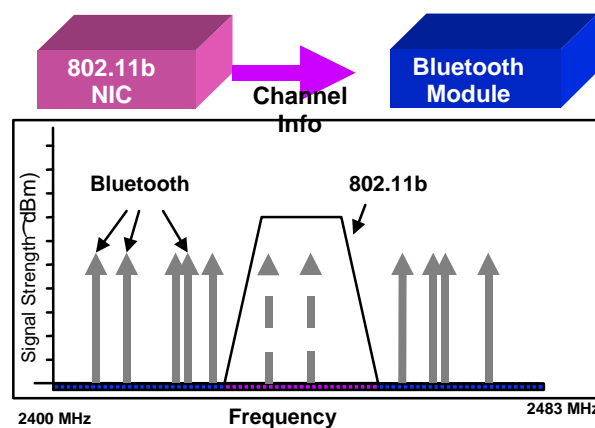


Figure 2: Intel WCS Channel Exchange

Intel WCS is designed to complement the Adaptive Frequency Hopping (AFH) interference mitigation algorithm being developed by the Bluetooth Special Interest Group (SIG). AFH will mitigate the impact of

802.11b on Bluetooth data throughput, but only between AFH-compliant Bluetooth devices as shown in Figure 3.

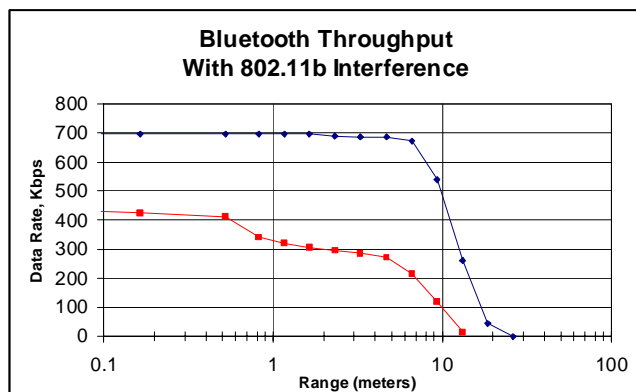


Figure 3: Bluetooth throughput with 802.11b interference

Phase 2 of Intel WCS will add Bluetooth priority signaling from the Bluetooth module to the Intel PRO/Wireless network connection, resulting in a restoration of connection reliability for both AFH and non-AFH (legacy) Bluetooth devices as shown in Figure 4.

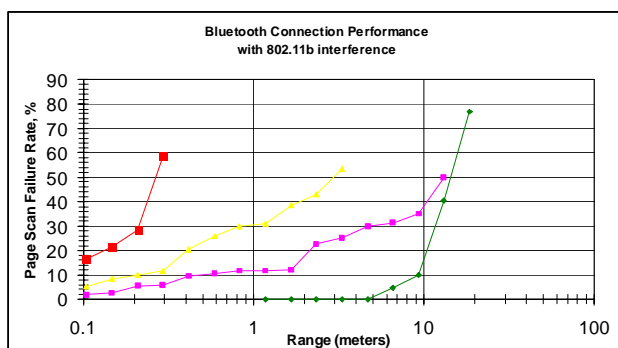


Figure 4: Bluetooth connection performance

Bluetooth Enabling

Multiple third-party Bluetooth silicon and module vendors have been enabled to be compatible with Intel WCS, including [SiliconWave](#) and [Cambridge Silicon Radio](#). Extensive verification and validation testing has been completed with these silicon vendors, providing pre-validated system solutions to the customer.

In summary, integrating multiple RF technologies into mobile notebook PCs provides new challenges to systems designers, including antenna gain uniformity and interference mitigation. Intel WCS, an Intel Centrino mobile technology, provides powerful interference mitigation between 802.11b and Bluetooth and enhances AFH.

EXTENDED BATTERY LIFE TECHNIQUES

The battery life of a mobile PC is a function of the power consumed by each of its components and the capacity of the battery. Figure 5 shows the breakdown of the power consumption of a typical notebook while running the battery life benchmark Ziff-Davis BL4.01.




Figure 5: Platform average power distribution

As can be seen, some of the largest power consumers are the display (LCD), power supply, hard disk drive (HDD) and graphics. It is also noteworthy that the CPU is one of the lowest consumers of battery life power (only consumes 6%) due to the advanced power-management built into the processor (Intel SpeedStep® technology, QuickStart, Deep Sleep, etc.) To learn more about processor power management, please see “Intel® Pentium® M Processor Power Estimation, Budgeting, Optimization, and Validation” in this issue of the *Intel Technology Journal*.

On Intel Centrino mobile technology, Intel developed and enabled the following *techniques* to reduce the power consumption of some of these subsystems and to help increase the battery life:

- Asynchronous Voltage Regulator (VR) Control with Power Status Indicator #
- Multiphase Intel Mobile Voltage Positioning (IMVP) technology

Asynchronous Voltage Regulator Control with PSI#

The fully power-optimized CPU represents a very complicated load to the voltage regulator (VR). This is because its current draw can range from a few mAs to several tens of Amperes, depending on the workload put on the processor. The transitions from low to high currents can also occur rapidly (in ms). This makes it difficult for the VR designer to maintain high-power

conversion efficiency in the full operating range. Typically, the efficiency is maximized at the highest end of the current range. However, the power-conversion efficiency of the VR drops off quickly toward the low-power condition. This is because the standby or quiescent (zero output loading) power of the VR is approximately 0.2-0.5W, which is comparable to the low-power CPU load, causing the efficiency to be <50%.

In a battery life benchmark, such as the ZDBL4.01, the processor spends >80% of the time in the low-current state (C3). Therefore, it is very important to maximize the power-conversion efficiency of the VR during the C3 state to extend the battery life of the notebook. The Asynchronous VR control with PSI# is a method used to increase the efficiency of the VR during the C3 state, without affecting the efficiency during the high-power state.

A VR is used to convert an unregulated high input voltage, either the AC/DC adapter or battery, to a suitable regulated DC voltage rail to power the core of the CPU. In this particular application, the regulated voltage is $1.35V \pm 7.5\%$, including VR DC error, noise, and transient response error, etc. A switching voltage regulator is used to produce high-power conversion efficiency. Due to the high current demand of the CPU, this VR operates in a continuous synchronous topology.




Figure 6: Synchronous VR

Referring to Figure 6, a VR is used to regulate the battery voltage (8.1-12.6V, 10.8V nominally) down to 1.35V for the CPU core rail. A VR controller monitors the output voltage, V_{CORE} , and compares it to an internal 1.35V reference.




Figure 7: Non-synchronous VR with CPU power status indicator

Referring to Figure 7, the logic driving the gate of Q_{SYN} is gated, in addition to CO, by a Power Status Indicator (PSI). PSI, when LO, keeps the gate of Q_{SYN} low regardless of CO level. When PSI is HI, the Mosfet driver drives the gate of Q_{SYN} normally. In this application, the STP_CPU# signal can be used in place of a PSI.

The Geyserville-III transition (Performance Optimized Mode to Battery Optimized Mode) also requires additional modifications due to the voltage differences between the two modes. Figure 8 shows the final circuitry that is used in Intel Centrino mobile technology.




Figure 8: STP_CPU# override with VID code

Optimized Intel Mobile Voltage Positioning

A CPU enabled with Intel SpeedStep technology operates at a different CPU core frequency. For example, when operating with an AC adapter plugged into the notebook, battery life is not an issue, as power is always available (although power consumption can limit performance as the internal temperature may rise.) The CPU is placed into a Performance-Optimized Mode (POM), where the core frequency is high (i.e., 1GHz). When operating from the battery as the input power source, the CPU is placed into the Battery-Optimized Mode (BOM) where the core frequency is dropped to a lower value (i.e., 600MHz) to reduce power consumption. The CPU power demand for POM is higher than that of BOM. More importantly, the current demand for POM is much higher than for BOM. This means the output-decoupling requirement for POM is much worse than that for BOM. The CPU voltage tolerance is $\pm 7.5\%$ for both POM and BOM because the Intel Mobile Voltage Positioning (IMVP) load line remains constant for both POM and BOM. Therefore, when switched to BOM, using the GMUXSEL signal

(with low indicating that the CPU is in BOM and high indicating POM), the output voltage can be “shifted” down to reduce power consumption.

Intel Mobile Voltage Positioning Design Implementation

Figure 9 shows positive offset voltage is controlled by $R_{\text{OFFSET-P}}$. A small-signal mosfet switch is added in series with this resistor and ground. During POM, GMUXSEL is high connecting $R_{\text{OFFSET-P}}$ to ground, providing a positive offset voltage. During BOM, GMUXSEL is low turning off the mosfet, removing the positive offset voltage.

During DeepSleep state, a negative offset voltage is applied through $R_{\text{OFFSET-N}}$ and pulled down with the STP_CPU# signal.

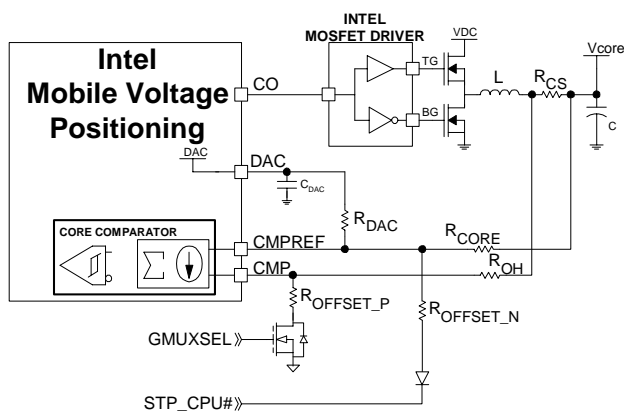


Figure 9: IMVP implementation

RESULTS

The techniques were implemented on two IBM T20 notebooks with 600 MHz Pentium® III processors. All data, unless mentioned otherwise, are averaged power numbers measured during the entire battery discharge (~3-4 hours duration) during a Ziff-Davis Battery Life benchmark, ZDBL4.01.

To establish accurate improvement (if any), a baseline measurement is first established. Two complete ZDBL runs were performed on a new T20 notebook platform with the same battery pack charged overnight each time. An average runtime was calculated along with an average platform/battery power measurement.

CPU Intel Mobile Voltage Positioning Feature Results

Figure 10 shows a battery from a charged status, 12.0V, down to a discharged status, ~8.5V. The average power over the entire ZDBL run of 3 hours, 40 minutes is 11.77W.




Figure 10: Discharge battery voltage/power profile—baseline

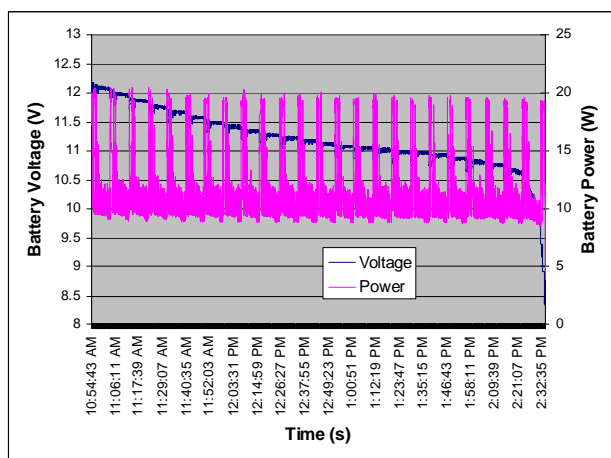


Figure 11: Discharge battery power profile with CPU IMVP VR

Measurements after CPU Intel Mobile Voltage Positioning Modification show the average power over the entire ZDBL run is 11.49W. This is a 280mW power reduction in the CPU during the ZDBL run (see Figure 11).

® Pentium III is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Non-Synchronous Voltage Regulator with PSI# Results

Figure 12 shows the battery discharge power over the entire 3 hour, 15 minute ZDBL run of the IBM T20 notebook used for this study.

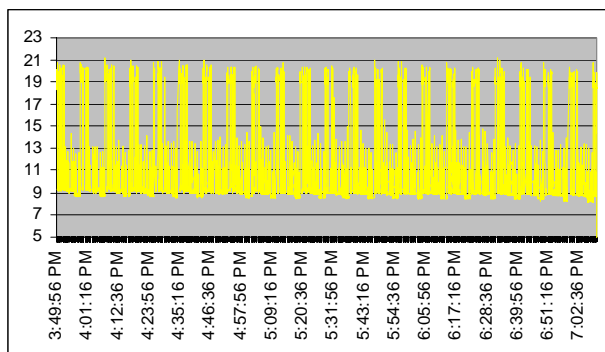


Figure 12: Discharge battery power profile-baseline

The average power over the entire ZDBL run of 3 hours, 19 minutes, is 10.94W.

With the non-synchronous VR with PSI# (STP_CPU#) implemented, the discharge profile is as shown in Figure 13:




Figure 13: Discharge battery power profile with STP_CPU# VR

The average power over the entire ZDBL run of 3 hours, 36 minutes, is 10.36W. This is a 580mW power reduction in the CPU VR.

In summary, as demonstrated in the above experiments, the two features implemented on the IBM T20 can result in over 1W of power savings during the battery life benchmark (ZDBL4.01) run. While the benefits may vary with different OEM systems, the expectation is that these will save power and improve the life of the battery.

THERMAL DESIGN

The Intel Centrino mobile technology platform can be cooled using a simple, compact, light-weight solution. The approach, though simple, is fundamentally important to successful design. In this section, we discuss a strawman set of requirements, the overall cooling approach, cooling performance expectations, and the resulting cooling solution design.

For the purposes of this discussion, we define a set of strawman requirements comprising the system form-factor, a platform power scenario, and a set of boundary conditions. We assume a system thickness of 25 mm (total base and display) with a footprint of 320x260 mm, which accommodates a 15" display. This thin form-factor allows for approximately a 13 mm inside height below the keyboard, for a solution design. We further assume a power scenario in which the processor may be at 24.5W, and all the other components may be at 18.5W, for a total of 43W in the base of the system (see Figure 14).

Platform Power (High Power Application)




Figure 14: Platform TDP power

Finally, important boundary conditions are the maximum temperature difference allowed between the component and the user ambient ΔT_{j-a} . For the CPU, we assume 65°C. Additionally, any given system design must consider chassis surface temperatures as well as acoustic limits on fan noise, for ergonomic constraints.

The overall thermal solution approach is illustrated in Figure 15:




Figure 15: Notebook layout illustration

It includes the now conventional remote heat exchanger approach for component cooling, but with fresh air flow capability set aside to provide cool air to the system. The fresh air flow is important to cool components and especially to keep skin temperatures cool. The passive cooling limit (i.e., dissipation by radiation and natural convection from the base of the system) is approximately 18W, allowing for a chassis skin temperature of roughly 15°C above ambient as is observed in typical system designs; close to the non-processor power assumed above.

The overall cooling performance expectations are identified using the equation below:

$$\theta_{j-a} = \frac{\Delta T_{j-a} - T_{system}}{P_{CPU}}$$

where P_{CPU} is the power we need to cool and ΔT_{j-a} is as discussed above. T_{system} represents the net impact of the power of the rest of the platform on the component, in this case the processor and the thermal solution. The value of T_{system} is dependent upon system design and is best approximated prior to design by using system simulations. Assuming 10°C for the strawman scenario discussed above, the required performance θ_{j-a} is 2.25°C/W.

Figure 16 illustrates schematically the remote heat exchanger cooling solution and important locations of temperatures. The overall performance θ_{j-a} is further approximated by the equation below:

$$\theta_{j-a} \approx \theta_{j-hp} + \theta_{hp-a}$$

where θ_{j-hp} is the thermal resistance from the component junction to the heat pipe and θ_{hp-a} is the remaining resistance from the heat pipe to the ambient temperature through the heat exchanger. θ_{j-hp} includes resistance of conduction through silicon, the interface material and contact resistances, and the resistance of the evaporator (region over the heat source) of the heat pipe. With new thermal interface materials, suitable thermal attach and heat pipe design, a performance of approximately 0.45 C-cm²/W can be achieved on a reference die. For performance verification, direct communication with respective suppliers is recommended. All components of θ_{j-hp} are affected by the source they are cooling and are considered in θ_{j-hp} . Although the Intel® Pentium® M processor uses 0.13 micron process technology, the design is optimized to mitigate the effects of a smaller source,

and the effective θ_{j-hp} translates to approximately 0.9°C/W.

Prototypes from fan and heat exchange suppliers indicate that 1.1°C/W can be achieved for θ_{hp-a} . However, the θ_{hp-a} value must be de-rated to allow for fresh air to be set aside for system cooling. Although the passive limits are close to those accommodated by the maximum passive limits, spreading is not perfect, and an allowance must be made for system constraints on effective passive dissipation. We assume 15% θ_{hp-a} performance set aside to complement system cooling. It is important to note that the air flow is provided directly from the exhaust of the fan and flows underneath the motherboard to best protect skin temperature. Two very important features of the system design provide short and direct air flow paths into the fan and out of the heat exchanger and adequate system air outlet vents, respectively.

Figure 16 shows a design rendition of the solution. The weight and volume of prototypes are approximately 55g and 4.3cm³, respectively, approximately one-third the weight and size of solutions for higher-power processors.

In summary, the thermal solution for Intel Centrino mobile technology platforms is compact and lightweight, utilizing only a single fan to cool the processor and the system.




Figure 16: Thermal solution

PLATFORM-LEVEL VALIDATION

This section outlines the comprehensive platform-level validation process that helped deliver a stable platform with Intel quality. Intel validation BKMs were applied to ensure that a mobile platform with Intel Centrino mobile technology delivered optimum system performance with multiple wireless technologies operating concurrently within the enterprise as well as when connected to hotspots outside the enterprise.

The goal for the platform with Intel Centrino mobile technology was extended validation of all platform components to maximize reliability and interoperability.

® Pentium M is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

This validation was primarily divided into the following three categories (see Figure 17):

1. Validation of Intel and Third-Party Vendor (TPV) Intel Centrino mobile technology-enabled platform components.
2. Validation of Intel Centrino mobile technology-enabled platforms on customer reference boards.
3. End-User system validation.




Figure 17: Systematic approach to validation

Figure 18 shows Intel and TPV delivered components on a platform with Intel Centrino mobile technology. Each Intel component underwent thorough comprehensive system and compatibility validation using Intel validation BKM's traditionally used on processors and chipsets. To ensure high-quality validation of TPV components, Intel worked with TPVs to define comprehensive test plans for validation of their components using Intel validation BKM's.




Figure 18: Intel Centrino technology platform integration and validation

The focus of platform-level validation was the integration of Intel Centrino mobile technology components on an Intel Customer Reference Board platform and the performance of comprehensive compatibility and interoperability validation under mobile-specific stress test conditions as well as under typical usage-model scenarios. Figure 19 shows some examples of platform validation tests.

Intel Platform Validation Tests Overview	
Test Category	Test Examples
Pentium-M Processor and System Power Management Testing	<ul style="list-style-type: none"> Verify C3 and C4 State transitions occur with the wireless device and its software Suspend and resume system from S1, S3, and S4 states many times and verify the Wireless Device and its software is in consistent state
Multiple Wireless Integration Testing	Perform the Banias Processor and System Power States testing with all the wireless elements present in the system
Stress Testing	Verify wireless device and its software is in consistent state while running CPU and network intensive applications like DVD playback, and Packet Blaster

Figure 19: Platform validation test examples

The third type of validation was platform-level validation using Intel Centrino mobile technology-based OEM systems. The focus was testing in a real end-user environment. For example, multiple Intel Centrino mobile technology-based systems with multiple platform configurations were tested for compatibility with industry- standard Wireless LAN (WLAN) access points at various hotspots.

Intel Wireless Verification Program

As part of establishing Intel Centrino mobile technology as the premier Wireless LAN notebook PC client, Intel worked with many leading companies to accelerate the deployment of 802.11 wireless communication capabilities in private and public spaces. The overall goal is to help reduce the two major Wi-Fi adoption barriers: availability and awareness. The program works with traditional and emerging Wi-Fi providers and key location owners to achieve the following:

1. Deploy Wi-Fi in locations relevant to the business traveler, such as airports, hotels, and franchise chains.
2. Verify Intel Centrino mobile technology in the installed infrastructure.

3. Raise the awareness of the availability of the service in these locations by directed co-marketing programs and a signage program.

incorporated into this paper; and Francis Truntzer for co-developing and driving wireless platform validation strategy.

CONCLUSION

Platforms enabled with Intel Centrino mobile technology incorporate innovative technologies to significantly increase the mobility of notebook PCs. We described the following innovations:

- Details are given on how to reduce interference between wireless communication technologies. The Intel Wireless Coexistence System is demonstrated to basically recover all WLAN performance under Bluetooth interference.
- Two new, advanced platform power-management features for improved battery life are described in this paper: Asynchronous VR control with PSI, and multiphase IMVP. These features implemented on an IBM T20 can result in over 1W of power savings during the battery life benchmark (ZDBL4.01) run. While the benefits may vary with different OEM systems, the expectation is that these will save power and improve battery life
- By coupling state-of-the-art thermal techniques with the advanced platform power-management features for lower power consumption, the thermal solution for a typical platform with Intel Centrino mobile technology was demonstrated to be compact and lightweight, utilizing only a single fan to cool the processor and the system.

Finally, the Intel Centrino mobile technology platform is extensively tested and tuned for components of Intel Centrino mobile technology to optimally work together in order to maximize reliability and interoperability and to deliver on all four mobility vectors. Intel Centrino mobile technology continues to undergo extensive security validation with industry-standard security and leading third-party security solutions. Intel is continuing to conduct comprehensive infrastructure verification with the wireless LAN infrastructure ecosystem and public wireless LAN service providers.

ACKNOWLEDGMENTS

The authors thank the following individuals for their contributions: Al Bettner for antenna development and enabling; Marlon Cardenas and Je-Young Chang for driving thermal solution development; Camille Chen for driving WLAN coexistence integration and validation; Randy Durrant for Intel WCS systems design and Bluetooth enabling; Eddie Lin for driving Bluetooth coexistence integration and validation; Don Nguyen for developing the detailed white papers on each of these initiatives, which provided much of the data that is

REFERENCES

For more information on Extended Battery Life Program, please access the following Web site:
<http://developer.intel.com/design/mobile/battery.htm>

AUTHORS' BIOGRAPHIES

Gordon Chinn is the Radio Frequency (RF) Systems Architecture manager in Intel's Mobile Platforms Group (MPG) and is responsible for enabling mobile platform PCs with multiple, concurrent, embedded radio technologies. He has over 30 years of experience in various RF communications development positions in the commercial and aerospace industry. Gordon holds a B.S. and M.S. degree in Electrical Engineering and Computer Science from the University of California at Berkeley. His e-mail is gordon.chinn@intel.com.

Sanjiv Desai is the Mobile Systems and Wireless Architecture manager in Intel's Mobile Platforms Group (MPG). He is responsible for the definition and architecture of mobile platforms at Intel. He has over 20 years of industry experience in various CPU, Chipset, and System design and development positions at Intel, National Semiconductor, and American Microsystems Inc. (Gould). Sanjiv holds an M.Sc. degree in Electrical Engineering from North Carolina A&T State University (USA) and a B.Sc. degree in Electrical Engineering from M.S. University (India). His e-mail address is sanjiv.c.desai@intel.com.

Eric DiStefano is the manager of the Thermal Tools and Technology team in the Mobile Platforms Group. He has been with Intel defining thermal solutions for mobile processors and platforms for five years. Prior to Intel, he worked in the aerospace industry for 15 years developing advanced space propulsion, propellant feed systems, and cooling of various systems including avionics. His e-mail is eric.distefano@intel.com.

Krishnan Ravichandran has a B.S.E.E. degree from IIT Madras, India and an M.S.E.E. degree from Carnegie Mellon University. He joined Intel in 1988 after completing his M.S. degree and has spent his entire career in the Mobile Group, performing a variety of roles (circuit design, chipset lead, CPU Design Manager, Program Manager, Platform Architecture and more recently Extended Battery Life Manager). His e-mail is krishnan.ravichandran@intel.com.

Shreekant (Ticky) Thakkar is the co-director of Mobile Platform Architecture in Intel's Mobile Platforms Group (MPG). He has over 24 years of experience in various development and planning positions at Intel and Sequent (now part of IBM). Ticky established direction for mobile notebook PCs to transition from portable to wireless computing—making the notebooks deliver the best wireless

Internet experience. Ticky holds Ph.D. and M.S. degrees in Computer Science/Engineering from the University of Manchester (England) and a B.S. degree in Computer Science/Statistics from the University of Manchester. His e-mail is ticky.thakkar@intel.com.

Copyright © Intel Corporation 2003. This publication was downloaded from <http://developer.intel.com/>.

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

For further information visit:

developer.intel.com/technology/itj/index.htm



Intel® Technology Journal

Intel® Centrino™ Mobile Technology

The Intel® Pentium® M Processor: Microarchitecture and Performance

The Intel[®] Pentium[®] M Processor: Microarchitecture and Performance

Simcha Gochman, Mobile Platforms Group, Intel Corporation
Ronny Ronen, Microprocessor Research Labs, Intel Corporation
Ittai Anati, Mobile Platforms Group, Intel Corporation
Ariel Berkovits, Mobile Platforms Group, Intel Corporation
Tsvika Kurts, Mobile Platforms Group, Intel Corporation
Alon Naveh, Mobile Platforms Group, Intel Corporation
Ali Saeed, Mobile Platforms Group, Intel Corporation
Zeev Sperber, Mobile Platforms Group, Intel Corporation
Robert C. Valentine, Mobile Platforms Group, Intel Corporation

Index words: Pentium[®] M processor, microarchitecture, power-aware design, branch prediction, instruction fusion, processor bus, SpeedStep[®] technology

ABSTRACT

The Intel[®] Pentium[®] M processor is a key component of Intel[®] Centrino[™] mobile technology platform. It is Intel's first microprocessor designed specifically for mobility. It provides outstanding mobile performance¹ and its dynamic power management enables energy saving for longer battery life.

Designing a mobile processor calls for different power/performance tradeoffs than designing a traditional high-performance processor. In this paper we explain the design philosophy that was adopted by the Intel Pentium M processor's architects to achieve best performance at given power and thermal constraints.

We present an overview of the Intel Pentium M processor's major advanced power-aware performance features including the innovative branch predictor, the dedicated stack manager, the micro-operation fusion, and the Intel Pentium M processor bus.

[®]Intel Pentium M is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

[®]Intel Centrino is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

¹ System performance, battery life and functionality will vary depending on your specific hardware and software configurations.

We next describe the Pentium M processor's Enhanced Intel SpeedStep[®] technology that allows significant reduction in energy consumption without compromising performance.

We conclude with demonstrating the superior performance and power-awareness of the Pentium M processor by comparing it with other mobile processors on a variety of known industry benchmarks.

INTRODUCTION

The distinction between *Mobile* and *Desktop* computing segments is not new. There are several vectors in which these segments differ, two of which are relevant to our discussion: power dissipation and battery life [1].

- *Power, Power Density, and Thermal.* The overall dissipated power, as well as the power dissipated by the chip per unit area, are important factors. Power generates heat. In order to keep transistors within their allowed operating temperature range, the generated heat has to be dissipated from its source in a cost-effective manner. These constraints limit the processor's *peak power* consumption. Peak power consumption limits apply both to desktops and mobile computers. However, the mobile computer's

[®]SpeedStep is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

smaller form-factor and lighter weight decrease the mobile processor's power budget².

- *Battery Life.* Batteries are designed to support a certain Watts x Hours³. The higher the average power is, the shorter the time that a battery can operate. This constraint limits the processor's *average power* consumption, which is a crucial factor for mobile computers but less relevant for desktop computers.

Until not so long ago, the size of the mobile market was significantly smaller than that of the desktop market, causing mobile PC designers to retrofit processors designed to address the desktop market. Desktop processors were designed to achieve the highest performance possible for all user profiles with little consideration of power consumption. Meeting the more restrictive peak and average power constraints of mobile PCs involved a compromise. New processors were adapted to the mobile market by either operating them at a lower voltage and frequency, hence compromising performance, or by delaying the implementation of their mobile version to the next-generation process technology, hence losing time-to-market.

The increased demand for mobile PCs, combined with the growing gap between desktop and mobile peak and average power budgets, made it impractical to continue the trend of using desktop processors for the mobile market. Mobile users expect to have close-to-desktop performance even with the more restricted power and thermal environment. Addressing this need called for a processor designed with the mobile environment in mind. This is where the Intel Pentium M processor comes in.

The Pentium M processor is Intel's new flagship power-aware microprocessor. Upon introduction (in March 2003) its highest performing version ran at 1.6 GHz @1.47V⁴, its Low Voltage version ran at 1.1 GHz @1.18V, and its Ultra Low Voltage version ran at 900 MHz @1.0V. It follows a new design approach with the goal of delivering breakthrough performance at

a lower power budget as well as minimizing the processor's average power for extending battery life. The Intel Pentium M processor includes several new innovative features that enable it to meet its design goals.

This new processor has 77 million transistors implemented on Intel's 0.13 μ CMOS process, with six levels of copper interconnect. Its die size is 84 mm² and its peak power consumption is 24.5 watts at 1.6 GHz. Its 3.2 GB/second processor bus helps to provide the high data bandwidths needed for today's and tomorrow's demanding applications. It fully implements the IA32 instruction set architecture [2], including Streaming SIMD Extension (SSE) and Streaming SIMD Extension 2 (SSE2) targeted for multimedia, content creation, scientific, and engineering applications.

We begin with an overview of the Intel Pentium M processor design philosophy. Then we examine in depth the major innovative power-aware and energy-aware features of this processor. We conclude by comparing the performance and power-awareness of the Pentium M processor with those of other mobile processors.

POWER-AWARENESS PHILOSOPHY AND STRATEGY

Design tradeoffs for the mobile market are rather complicated and involve several challenges:

- Optimizing the design for maximum performance and extended battery life. The challenge lies in how to balance between these conflicting goals.
- Trading performance for power. Performance features—whether increasing instruction-level parallelism (ILP) or speeding up frequency—usually increase power consumption. Power-saving features usually decrease performance. The challenge lies in figuring out how much power one can afford to lose in order to implement a performance feature.

Setting a clear direction for tradeoffs between higher performance and longer battery life and between performance and power was essential for converging the definition of the Intel Pentium M processor.

Higher Performance vs. Longer Battery Life

Early in the design it was realized that the average power consumption for typical usage of the Intel Pentium M processor is only a small portion of the whole platform power consumption—less than 10%. This low average power is mainly due to the ability of the processor to enter lower power states in idle periods and to the Enhanced Intel SpeedStep technology, which significantly reduces power in periods of low processor activity. The majority of the power in the platform is

² Currently (2003), typical desktop processor peak power consumption is about 100W. Typical mobile processor peak power is about 30W.

³ Currently (2003), typical battery capacity is 24-72Wh. A typical Pentium M platform uses 48Wh batteries. A typical platform's power is 13W, of which the processor consumes about 1W on average. Smaller platforms use 24Wh to save weight.

⁴ 1.47V is the highest operating voltage of the 0.13 μ CMOS process on which the processor is implemented.

consumed by other components: the LCD, the hard disk, the memory system, networking components, etc. Under these circumstances, it was clear that the potential gain in system battery life by further reducing the processor's average power would be small⁵. With that in mind, we decided to optimize the design for the highest performance within power and thermal constraints, when the processor is active, and to focus on battery life when it is idle.

It should be noted that the above may change in the future. We expect high-performing processors' average power to grow due to more complex processor logic and higher static power and new demanding workloads. We expect platforms to become smaller, simpler, and more efficient hence consuming less power. Therefore, the portion of the processor's average power as part of the overall platform power consumption will increase.

Trading Performance For Power

The tradeoffs are different if we optimize for higher performance or for longer battery life. For higher performing mobile processors (and in fact, now, in all high-performing processors) the criterion is "*Maximizing performance at given thermal constraints.*" For longer battery life the criterion is "*Minimizing energy per task.*" Below, we explain what each criterion actually means.

Maximizing Performance at Given Thermal Constraints

The processor's thermal map depends on the power consumption, the local power density at various points on the die, the cooling mechanism, and more. At the early stages of the Intel Pentium M processor's microarchitecture definition, when thermal information was not available, we replaced the criterion "*Maximizing performance at given thermal constraints*" with: "*Maximizing performance at a given power envelope*"⁶.

According to this criterion, a microarchitectural feature that gains performance or saves power should be better than simply using voltage/frequency scaling. For a given working point of core voltage V_0 and Frequency F_0 the power consumption of a processor is given by

$$Power_0 = \alpha * C_0 * V_0^2 * F_0$$

where α is the activity factor, $Power_0$ is the power consumption and C_0 is the effective capacitance for a given design. The frequency is usually approximated as being linearly proportional with the operating voltage, namely

$$F_0 \cong K_f * V_0$$

where K_f is the proportion constant. This leads to the cubic dependency of power on the operating voltage

$$Power_{max} = \alpha * C_0 * V_0^3 * K_f$$

The performance at this operating condition is given by

$$Perf_0 \cong IPC_0 * F_0$$

where IPC_0 indicates the Instruction Per Cycle in Frequency F_0 ⁷.

It can be derived from the above formulae (see also [3]) that by increasing the voltage by 1%, for example, one can increase performance by 1% through increased frequency. This would result in a power increase of approximately 3%. Thus, an alternate microarchitectural feature that gains less than 1% in performance for a power increase of 3% or more should be rejected upfront. In general, a microarchitectural feature can be regarded as power-aware, if the % ratio between the power increase and the performance gain is less than 3.

Minimizing Energy Per Task

Energy consumption in general is a sum of two components: active energy and idle energy. Minimizing the idle energy consumption is relatively straightforward and does not involve conflicting design tradeoffs: the processor enters a deep-sleep power state, stops the clocks, and lowers the operating voltage to the minimum allowed to sustain the internal state. Optimizing active energy is more complex. A very slow execution consumes less power for a longer period of time, while heavy parallelism reduces the active time but increases the active power.

$$Energy_{active} = Power_{active} * Time_{active}$$

or

$$Energy_{active} \cong Power_{active} / Perf_{active}$$

This implies that in order to improve overall battery life, the % performance benefit must be greater than the additional power consumed.

⁵ In the ideal case where the 10% power is totally eliminated, battery life would be extended by only 11%.

⁶ The Intel Pentium M processor's design assumed power envelopes from 7W for passively cooled boxes up to 24W for Thin and Light platforms.

⁷ For the sake of this discussion, we assume performance scales linearly with frequency. In reality, mainly due to off-chips accesses (memory and I/O), performance does not scale with frequency.

During the definition of the Intel Pentium M processor we tended to use the stricter criterion in each case:

- Performance improvement features were usually included if they “*minimized energy per task*,” that is, they save energy. Features that pass this criterion improve performance and extend battery life—the ultimate win/win situation. Faster execution also implies longer idle time, allowing active units to be shut off, thus saving even more energy.
- Power-saving features may reduce Instructions Per Cycle (IPC) resulting in a performance loss. Such power-saving features were usually included only when they “*maximized performance at given thermal constraints*,” that is, the performance loss was smaller than would have been achieved by just using voltage and frequency scaling. In practice, by applying this criterion, a tradeoff can be made between saved power and increased frequency, thus squeezing more performance at the peak allowed power.
- Performance-improvement features that met the “*maximize performance at given thermal constraints*” criterion, but failed the “*minimizing energy per task*” one, were carefully judged and in many cases included. Such features do increase performance but consume more energy. This loss is negligible, since, as mentioned above, the processor’s average power as a portion of the overall system average power is relatively small. In fact, the faster execution results in a longer idle time, potentially allowing additional energy savings.
- [Figure 1](#) illustrates the design tradeoffs from a performance feature viewpoint. The magenta area (in the lower right side) indicates a clear win—improving both performance at the power envelope and battery life. The orange area (resides in the upper right side) indicates a tradeoff where constrained performance is preferred over lower battery life. The green area (resides in the lower left side) indicates a tradeoff where improved battery life is preferred over constrained performance. The white area (in the upper left side) is a clear “drop.”




Figure 1: Performance/power tradeoff zones

Now that the tradeoffs are known, we examine the strategies we used to identify power-aware features. Power-awareness means attacking power and energy consumption at all levels:

- Reducing number of instructions per task.
- Reducing number of micro-ops per instruction.
- Reducing number of transistor switches per micro-op.
- Reducing the amount of energy per transistor switch.

Reducing the Number of Instructions Per Task

From an architectural point of view the number of instructions per task is fixed. However, from a microarchitectural point of view, this is true only for the number of retired instructions. With branch prediction, there are many speculated instructions running within the processor that are not retired. A better branch predictor decreases the number of the speculated instructions, thus practically reducing the number of overall processed instructions. Indeed, the extra logic involved in a better branch predictor does consume power, but the gain in the reduced number of instructions exceeds that extra cost.

Reducing the Number of Micro-ops Per Instruction

Out-of-order implementations of the IA32 Instruction Set Architecture (ISA) break macro-instructions into a sequence of one or more simple operations, called micro-operations, or *micro-ops* [4]. Handling and executing each micro-op consumes power. Eliminating micro-ops from the micro-op stream or combining several micro-ops together reduces the overall power. The Intel Pentium M processor micro-ops fusion and dedicated stack engine do exactly that. Here, as well, the gain in reduced micro-ops exceeds the cost of the extra logic involved.

Reducing the Number of Transistor Switches Per Micro-op

This is more straightforward and intuitive. High performance processors run at a high frequency and

provide a high degree of instruction-level parallelism—switching a lot of transistors on the way. There is a clear gain in doing the same operation with a smaller number of switches. In some cases, this is simple and involves only local optimizations such as accessing only the portion of a register or a cache line that is actually needed. In other cases, it calls for a global optimization to decide whether a unit will not be used for the next cycle, and thus can be shut off. Occasionally, such power saving may involve a performance loss.

Reducing the Amount of Energy Per Transistor Switch

Energy per switch depends on the transistor size and type, and on the operating voltage. Smaller transistors and lower operating voltages reduce energy per switch [5]. The effect of microarchitecture here is rather limited. Transistor size and type are tuned so that they meet the timing constraints without wasting unnecessary power. The Enhanced Intel SpeedStep technology reduces the operating voltage at low activity periods, thus reducing the energy per transistor switch. Microarchitecture can help reduce energy per switch by optimizing the amount of interconnect in the processor.

Each strategy affects performance, power, and energy in a different way. In most cases, features that fall under one of these strategies save energy per task. However, performance-improving features are likely to result in increased power consumption, e.g., better branch prediction reduces energy, but also reduces stalls, thus increasing power. This may look bad, but, in fact, it is not. Higher performance at lower energy can be traded for lower power by slowing down the processor either by using voltage/frequency scaling or microarchitectural throttling.

Static Power

The power consumed by a processor consists of active power (used to switch transistors) and static power (leakage of transistors under voltage). In this paper we focus mainly on active power reduction, but it is worth mentioning how the Intel Pentium M processor also reduces static power consumption.

The static power is roughly a function of the number of transistors, their type, the operating voltage, and the die temperature. The Pentium M processor reduces static power by several means:

- *Low-leakage devices.* The processor's 1MB power managed L2 cache, which contains roughly two-thirds of the transistors in the processor, is built with low-leaking transistors. Low-leaking transistors are somewhat slower, thus slightly increasing the cache

access latency, but the significant power saved justifies the small performance loss.

- *Enhanced Intel SpeedStep technology.* This advanced technology significantly reduces the processor voltage (and temperature), hence leakage power, when processor activity is low.

POWER-AWARE FEATURES

The following sections describe several of the Intel Pentium M processor's power-aware features. These features cover all the above-mentioned strategies:

- Reducing the number of instructions per task: advanced branch prediction.
- Reducing the number of micro-ops per instruction: micro-ops fusion and dedicated stack engine.
- Reducing the number of transistor switches per micro-op: the Intel Pentium M processor bus and various lower-level optimizations.
- Reducing the amount of energy per transistor switch: Intel SpeedStep technology.

ADVANCED BRANCH PREDICTION

For high-frequency pipelined microprocessors, branch prediction continues to be one of the biggest ticket items for gaining performance. In the Intel Pentium M processor, the benefits are actually twofold: the decrease in speculative code gains performance and reduces the energy spent per instruction retired.

The advanced branch prediction in the Pentium M processor is based on the Intel Pentium® 4 processor's [6] branch predictor. On top of that, two additional predictors to capture special program flows, were added: a Loop Detector and an Indirect Branch Predictor.




Figure 2: The Loop Detector logic

®Pentium 4 is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

The Loop Detector (Figure 2) analyzes branches to see if they have loop behavior. Loop behavior is defined as moving in one direction (taken or not-taken) a fixed number of times interspersed with a single movement in the opposite direction. When such a branch is detected, a set of counters are allocated in the predictor such that the behavior of the program can be predicted completely accurately for larger iteration counts than typically captured by global or local history predictors.

The Indirect Branch Predictor solves the problematic data-dependent indirect branches. Indirect branches are heavily used in object-oriented code (C++, Java), hence they became a growing source of branch mispredictions. While most indirect branches have a single target at run time, some, such as a case statement in a byte-code interpreter, may have many targets. These targets are chosen in a data-dependent manner.

The Indirect Branch Predictor (Figure 3) chooses targets based on a global control flow history, much the same way a global branch predictor chooses the direction of conditional branches using global control flow history. As can be seen in the figure, it is an adjunct to the normal target prediction device. Targets are always allocated in the Instruction Pointer tagged table along with the type of branch. When a misprediction occurs due to a mispredicted target on an indirect branch, the Indirect Branch Predictor allocates a new entry corresponding to the global history leading to this instance of the indirect branch. This construction allows monotonic indirect branches to predict correctly from the IP-based Target array, and data-dependent indirect branches to allocate as many targets as they may need for different global history patterns, which correlate with the different targets. Entries in the Indirect Branch Predictor are tagged with the hit and type information in the IP-based target array to prevent “false positives” from the Indirect Branch Predictor to lead to new sources of target mispredictions.




Figure 3: The Indirect Branch Predictor logic

The Intel Pentium M processor branch predictor misprediction rate is 20% lower than that of previous generation designs, resulting in as much as 7% in real performance. Approximately 30% of this benefit comes from the combination of the Loop Detector and Indirect Branch Predictor. The Loop Predictor captures a common program behavior and benefits many applications regardless of compilation techniques. The Indirect Target Predictor shows its gains more in specific applications where indirect branches are used to select data-dependent targets. In such applications, when the compilers use calculated branches rather than if-trees made from conditional branches, performance gains can be a few percentage points.

MICRO-OPS FUSION

Out-of-order implementations of the IA32 Instruction Set Architecture (ISA) break macro-instructions into a sequence of one or more simple operations, called micro-operations, or *micro-ops*. A conventional micro-op consists of a single operation operating on two sources. The Instruction Decoder breaks a macro-instruction into multiple micro-ops whenever the macro-instruction operates on more than two sources or when the nature of the operation requires a sequence of unrelated operations. There are quite a few cases of macro-instructions that break into several micro-ops, two of which are store operations and load-and-op (read-modify) operations.

Macro-instructions that store data in memory are decoded as two independent micro-ops. The first operation—*store-address*—calculates the address of the store, while the second operation—*store-data*—stores the data into the Store Data buffer⁸. The separation between the store-data and the store-address operations is important for memory disambiguation. Breaking the store operation into two micro-ops allows the store-address operation to dispatch earlier, even before the stored data are known, enabling resolution of address conflicts and opening the memory pipeline for other loads.

A typical load-and-op macro-instruction consists of two micro-ops: the first operation reads the operand from an address in memory, and the second operation calculates the result based on the data read from memory and the register operand. A load-and-op macro-instruction may have up to three register operands, so it must be implemented by two micro-ops. The atomic operations

⁸ The actual write to memory is done when the store retires. Only then are the data in the respective Store Data buffers written into the specified address.

are inherently serial, and the second operation cannot start until the first operation completes.

Splitting the macro-instruction into multiple micro-ops also has its toll:

- The increased number of micro-ops creates pressure on resources with limited bandwidth (rename, retire) or limited capacity (Reorder-Buffer, Reservation-Station). This pressure eventually results in performance loss.
- Splitting a macro-instruction into more than one micro-op is a complex operation that requires a significantly more capable decoder. Due to its complexity, most implementations opt to have only one complex decoder; all other decoders are left to handle macro-instructions that break only into a single micro-op.
- Delivering more micro-ops through the system increases the energy required to complete a given instruction sequence.

The Pentium M processor features the micro-ops fusion mechanism to reduce this performance and energy cost while maintaining the benefit of the out-of-order execution. With micro-ops fusion, the Instruction Decoder *fuses* two micro-ops into one micro-op and keeps them united throughout most parts of the out-of-order core of the processor—at allocation, dispatch, and retirement. To maintain their non-fused behavior benefits, the micro-ops are executed as non-fused operations at the execution level. This provides an effectively wider instruction decoder, allocation, and retirement. [Figure 4](#) describes the different domains in which the micro-op is fused and unfused.

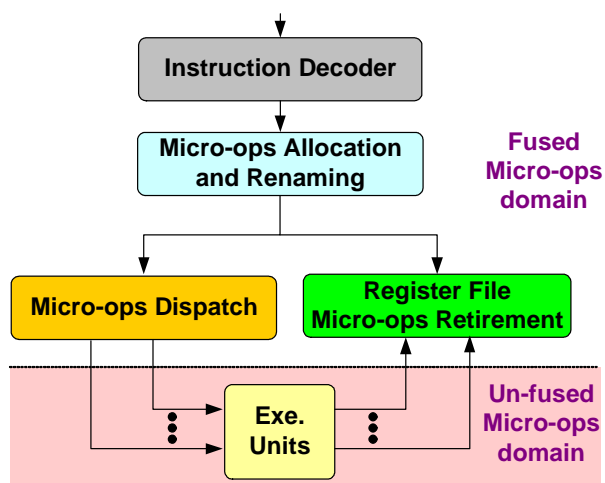


Figure 4: Micro-ops fusion domains

The macro-instruction is decoded into a single fused micro-op by the Instruction Decoder. The fused micro-op is allocated, renamed, and then issued into a single entry in the Reorder-Buffer and the Reservation-Station. To support fused micro-ops, each reservation-station entry can accommodate up to three source operands. When dispatching to the execution units, the Dispatcher controls the separate execution of each portion of the fused micro-op according to the readiness of its sources. In a sense, the Dispatcher treats each portion as if it occupied the whole entry for itself. The Execution of each operation is performed in the same way as a non-fused micro-op with only minor changes made to the execution units.

The fused store operation is depicted in [Figure 5](#).

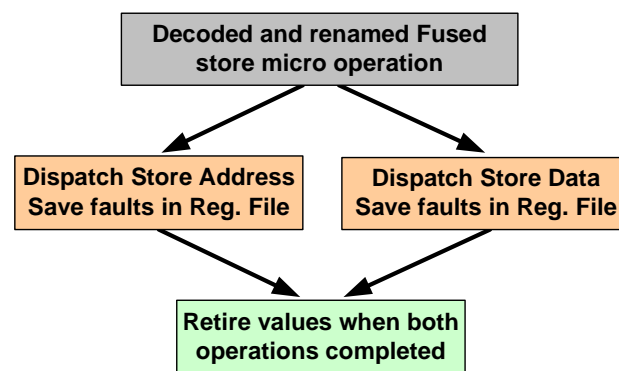


Figure 5: Fused store flow

The two micro-ops making up the fused store micro-op can be issued to their relevant execution units in parallel. The dispatch of the store-address operation to the address-generation unit is performed when its sources (the base and index registers) are ready. The dispatch of the store-data operation to the store data buffer unit can occur independently when its source operand is available. The retirement of the fused store can occur only after both operations complete.

The fused load-and-op operation is depicted in [Figure 6](#).

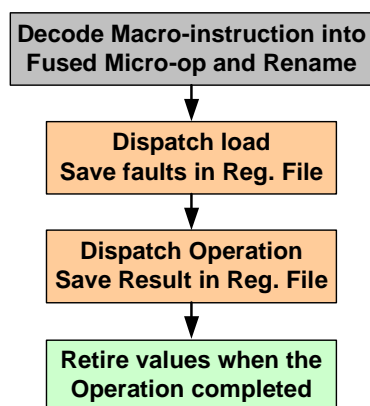


Figure 6: Fused load-and-op flow

The two micro-ops making up the fused load-and-op micro-op are issued serially to the relevant execution units. The dispatch of the load operation is performed when its sources (the base and index registers) are ready. The dispatch of the “op” portion of the load-and-op operation to the execution unit can occur only after the load completes and the other operand is ready. The retirement of the fused load-and-op micro-op can occur only after both operations complete.

We have found that the fused micro-ops mechanism reduces the number of micro-ops handled by the out-of-order logic by more than 10%. The reduced number of micro-ops increases performance by effectively widening the issue, rename, and retire pipeline. The biggest boost is obtained during a burst of memory operations, where micro-op fusion allows all decoders, rather than the one complex decoder, to process incoming instructions. This practically widens the processor decode, allocation, and retirement bandwidth by a factor of three.

The typical performance increase of the micro-op fusion is 5% for integer code and 9% for Floating Point (FP) code. The store fusion contributes most of the performance increase for integer code. The two types of fused micro-ops contribute about equally to the performance increase of FP code.

Delivering less micro-ops through the processor decreases the energy required to complete a given instruction sequence since the same task is accomplished by processing fewer micro-ops. The gain in power consumption is higher than the loss, due to the additional logic required to implement the micro-ops fusion mechanism.

DEDICATED STACK ENGINE

The IA32 Instruction Set Architecture (ISA) features instructions for hardware-assisted stack management that

are typically used to implement a combined parameter and control-flow stack used in high-level programming languages. The ISA provides PUSH, POP, CALL, and RET, which have the obvious parameter stack and control-flow stack behaviors. The ISA dedicates the hardware Stack Pointer register (ESP) as the machine stack pointer, and this register is modified as a side effect of each of these instructions. Sequences of such instructions are quite common, for instance, PUSHing a set of operands and then using a CALL instruction is the standard mechanism for making a Procedure or Function Call.

In traditional out-of-order implementations of the IA32 ISA, these side-effect operations were performed by sending with each stack-related macro-instruction an additional micro-op to update the ESP register. This micro-op adds or subtracts an immediate value to the ESP register.

The Intel Pentium M processor chose to implement the ESP “side-effect” behavior in a more efficient way, using dedicated logic near the superscalar decoders. The idea is to represent the programmer’s view of ESP (ESP_p) at any given point in time by some historic ESP living in the out-of-order execution core (ESP_o) added to a delta (ESP_d) that is maintained in the front end (see also [7]):

$$ESP_p := ESP_o + ESP_d$$

When, for example, a sequence of PUSHes and POPs is encountered in the instruction stream, the dedicated Stack Hardware executes the ESP side-effects in the decoders and updates the ESP_d register. Referring to [Figure 7](#), we can see a superscalar implementation for N+1 decoders passing the accumulated delta value across the decoders and updating the delta register with the result, after the instructions are decoded. The hardware also patches the in-flight ESP_d value into the address syllable of each of the stack referencing micro-ops (patch HW not shown for clarity) so that the address generation unit (AGU) can calculate the proper memory location referenced by ESP_p. This provides the following benefits:

- Dependencies on ESP are removed since the ESP_o value used for scheduling in the out-of-order machine is not changed during the sequence of stack operations. This allows more parallelism opportunities to be realized in the out-of-order execution.
- ESP_d updates are done using a small specialized dedicated adder, thus freeing the general execution units to work on other micro-ops. This allows a higher degree of superscalarity for these instructions

without the cost associated with going to a higher degree of superscalarity for all integer operations. Additionally, since the ESP updating micro-ops have been eliminated, the ALUs are free to be utilized by more complex operations that would have been blocked by the ESP updates, increasing execution bandwidth.

- Updating the delta register in the front end eliminates the ESP updates micro-ops from the out-of-order machine. Thus, power saving is realized since the large adders are not used for small operations, and the eliminated micro-ops do not toggle bits throughout the machine.




Figure 7: The dedicated stack engine logic

There are two complications with the Dedicated Stack Engine. Since it lives in the front of the pipeline, all its calculations are speculative. In order to recover a precise state at any point in the machine's life, the value of ESP_O and ESP_D must be able to be recovered for all instructions in the machine. ESP_O is maintained by the out-of-order core as any other general-purpose register. The Intel Pentium M processor adds an additional table (also shown in Figure 7) that saves the ESP_D value and a code relating to the effect on the ESP_D register for every instruction in the machine. This allows the value of ESP_P to be recovered for all instructions either pre- or post-execution. This allows for handling of either Faults or Traps as defined in IA32.

The second complication occurs when the architectural value of ESP is needed inside the out-of-order machine, for instance, "*XOR ESP, 3*" or, more commonly, when ESP is used in an address syllable. In this case, the decode logic automatically inserts a micro-op that carries out the ESP_P calculation. The ESP_D register can be cleared since the architectural value is now coherent. A sync is not generated when the ESP_D register is zero, so continued usage of ESP as a general-purpose register will have no ill effects.

The Dedicated Stack Engine ESP typically eliminates 5% of the micro-ops from an IA32 program compared to a processor not including this feature, even when

including the ESP synchronization micro-ops. Clearly this makes the out-of-order machine look virtually larger and frees execution bandwidth. However, the major performance gain is to increase the front-end bandwidth on these common instruction sequences to the full width of the superscalar decoders. These 5% of eliminated micro-ops result in a similar decrease in energy per instruction—in line with the Pentium M processor's power-awareness direction.

THE INTEL PENTIUM M PROCESSOR BUS

The Intel Pentium M processor bus was designed to provide a desktop-like performance while consuming significantly less power. Power saving is achieved by the combination of protocol and circuit methods that are unique to the Pentium M processor bus.

The processor bus supports 100 MHz bus clocks with a data rate of 400M transfers per second. It is a latched bus with an in-order queue of 8-pipelined transactions. Designed for mobile systems, the bus is optimized for a uni-processor environment. For example, it allows us to reduce the number of pins to save power:

- There are only 32 address bits that cover 4GB of physical address space.
- The bus does not support dual-processors, since the mobile systems' power budget cannot support dual processors anyway.

The Pentium M processor bus saves power aggressively when idle; it carefully controls its input buffer's sense-amplifiers that sample the activity on the bus. When the bus is idle, all sense amplifiers are disabled and do not consume any power. When the bus is active and address and data are driven on the bus, the input buffers are enabled in advance to ensure all information is captured with no delay.

The bus features many innovative mechanisms to reduce power while maintaining performance. Several of them are described below.

DPWR#: Data Bus Power Control. This is a special signal driven by the 855PM chipset whenever data are transferred to the processor. DPWR# is used to dynamically enable the processor's 64-bit data bus input sense amplifiers and their related controls (~80 signals) only when data are transferred to the bus.

BPRI Control: This is a method to achieve the DPWR# functionality for the address bus. BPRI# is asserted whenever the 855PM chipset attempts to drive the bus. It is used to dynamically enable the 32-bit address bus

input sense amplifiers and their related controls (~40 signals) only when a transaction is issued to the bus.

Low Vtt: The Intel Pentium M processor's I/O buffers work at a low voltage of 1.05V (Vtt). The low Vtt is an essential element to reduce the bus power. However, operating at low Vtt introduces a new set of problems because the I/O buffer is working at the low linear point, which affects the buffer's characteristics. The bus includes a special Resistor Compensation (RCOMP) method to adjust the buffer strength dynamically during run time. It accommodates the impacts of temperature, voltage drift, and bus topology. Thus, at any thermal and power state the Pentium M processor bus has full impedance termination. It has split power planes that allow setting the I/O operating voltage to a fixed value of 1.05V even though the core may be operating at a higher Enhanced Intel SpeedStep technology operating point.

PSI: Power Status Indicator. The Pentium M processor bus provides a signal to reduce the overall platform power (not just the processor power!). This signal is driven by the processor to control the current consumption of the Voltage Regulator (VR) when the processor operates at a low power state.

LOWER-LEVEL POWER OPTIMIZATIONS

This section describes several lower-level mechanisms that demonstrate the power awareness of the Intel Pentium M processor.

A simple, yet effective method that was pursued in order to reduce power was to identify idle logic and shut it off. This was done locally and globally.

Locally, the design was thoroughly reviewed for any inefficiency during idle states. The goal was to gate the clocks as much as possible. Ideally, the clocks should be shut off for each pipe stage separately. However, if such a naive approach is used, the added complexity may sometimes outweigh the gain. In these cases the logic is shut off for the entire unit only at the end of the operation, resulting in a smaller power saving.

Globally, idle time identification is done at a higher microarchitectural level, when the unit alone cannot identify the idle period. For example, the first stage of a unit is always kept awake in order to respond to incoming messages. So, the goal was to create a few central controllers that can microarchitecturally identify or predict idle periods and instruct the units to reduce power (either by shutting off their clocks or by disabling parts of their logic). Also, the prediction logic should

allow operations to resume seamlessly with no performance penalty.

One example for such a power predictor is the "Allocate stall" predictor. Whenever the Reorder-Buffer is full, the Allocator stalls the pipeline. However, the Allocator cannot tell if the Reorder-Buffer will remain full on the next cycle. It therefore needs to re-evaluate the stall condition every cycle. It turns out that in many cases when the Reorder-Buffer is full, it stays so for very long periods. Therefore the power penalty in this case is high. A specialized logic was defined to collect information from the Reorder-Buffer and other units in order to predict the nature of the next cycle. This logic instructs the Allocator to hold on to the stall condition and shut off its clocks.

Another type of a microarchitectural power feature identifies the logic that is necessary for a specific operation and activates only that part of the processor. Here are two examples for an implementation of such a power feature.

In a conventional processor, all the units of a specific execution port electrically share the same source bus wires. However, power can be saved if instead of driving the sources to all the execution units (EUs), only the wires that belong to the target EU are driven. Therefore, the Pentium M processor execution units were divided into a few segments (stacks), and a special logic was created in order to control the flow of data to every stack according to its actual destination (see [Figure 8](#)).




Figure 8: Execution units stacking

A generic processor operates on several different data types with different widths. In an IA32 processor there are integer operations, operating on 32 bits, multimedia operations, operating on 64 bits or 128 bits, and floating-point operations, operating on 80 bits. The most common instructions are integer instructions that access only a limited set of registers and use only 32-bit data values. Toggling a wider bus and reading from a bigger register file consumes more power than is actually

required. The Intel Pentium M processor saves power by identifying integer operations in advance and activating only the appropriate hardware (see [Figure 9](#)). The savings include the narrower buses to and from the EU during dispatch and writeback, and also other elements in the renaming logic that are not accessed while an integer operation is executing. This effectively transforms the processor into a 32-bit machine that utilizes only resources needed for integer operations while operating on integer data types.




Figure 9: Early identification of EU width

ENHANCED INTEL SPEEDSTEP TECHNOLOGY

High-performance processors tend to have high power consumption during execution. This is a simple derivative of the active power equation

$$Power = \alpha * C * V^2 * F$$

where V is the core voltage, F is the operating frequency, and α is the activity factor. However, proven mobile usage models, indicate that typical usage is bursty in nature, requiring high performance only for short bursts of time. Average power reduction can be achieved by switching voltage and frequency to a lower operating point, when demand is low. The efficiency of the solution depends on the ability to execute this operation-point switch frequently and efficiently, to track demand.

Previous generations of Intel mobile processors implemented the Intel basic SpeedStep technology [8]. It switches both voltage and frequency between two distinct states: Lowest Frequency Mode (LFM) and Highest Frequency Mode (HFM) by using the platform C3-idle state. While achieving low-power operation during LFM, the basic SpeedStep architecture does not fully address demand-based switching needs. The long system unavailability time during transitions limits the switching frequency due to interaction with streaming devices such as Audio Codec '97 (AC '97) and the Universal Serial Bus (USB). Additionally, having only two fixed operating points limits operating point

optimization according to the load. The Intel Pentium M processor introduces a multi-point Enhanced Intel SpeedStep technology optimized for demand-based switching.

The Enhanced Intel SpeedStep technology attempts to address the following challenges:

- *Minimizing system and processor unavailability.* Operating point switching requires voltage to be transitioned over a wide range (e.g., from 0.9V to 1.5V). Physical limitations of the power delivery system translate this demand to over 100 μ s delay. A full clock generator Phase-Locked-Loop relock requires approximately 30 μ s. The architecture needs to ensure system memory access unavailability will not exceed 10-15 μ s, to match isochronous device needs.
- *Self-managed voltage and frequency stepping.* The Enhanced Intel SpeedStep technology requires the migration of the mechanism from the chipset into the processor. This introduces two challenges: (a) how to sequence the operation when the processor clock is halted and (b) how to prevent loss of system events, such as interrupts and snoops, previously blocked by the chipset during the transition.

[Figure 10](#) depicts the high-level block diagram of the Enhanced Intel SpeedStep technology instantiation in the Intel Pentium M processor.

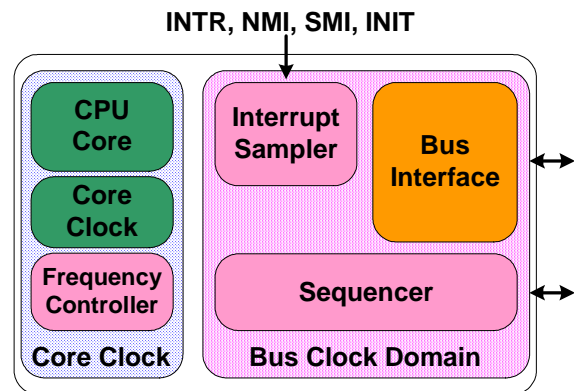


Figure 10: Enhanced Intel SpeedStep technology block diagram and clocking

The Enhanced Intel SpeedStep technology uses three novel principals to address the challenges stated above:

1. *Voltage-Frequency switching separation.* Unlike previous architectures, the Enhanced Intel SpeedStep technology separates the voltage and frequency transition stages ([Figure 11](#)).

Voltage is stepped in short increments, preventing clock noise and allowing processor execution during the voltage transition stage. Thus, system memory and the processor are made available during the longest segment of the operating point transition, thereby minimizing unavailability time to only the frequency transition stage.




Figure 11: Enhanced Intel SpeedStep® technology transition sequencing

2. *Clock partitioning and recovery.* During the Enhanced Intel SpeedStep technology transition, only the core clock and Phase-Locked-Loop are stopped, while the bus-clock is kept running. The Enhanced Intel SpeedStep technology logic was partitioned such that only the command interface and core controls operate on the core clock, while the sequencer and interrupt interface operate on the bus clock. Thus the logic can be kept active constantly, even though the core clock has been halted.

Additionally, the clock circuitry of the Pentium M processor was designed to utilize the active bus-clock to shorten core-clock relock time considerably. Thus core-clock restart time is set to only 10µs, minimizing the processor inactive time.

3. *Event blocking.* Interrupts, pin events, and snoop requests sent during the frequency transition stage must not be lost, even though the core clock is not available to serve them.

The Enhanced Intel SpeedStep technology logic samples all pin events when the core clock is stopped. These are re-sent to the processor once the core clock is available, preventing loss of events.

Bus events (such as snoops and interrupt messages) are blocked off using the native BNR# protocol, which captures the bus for the frequency transition period. Thus bus and pin events are not missed; they are serviced once the core is capable and running.

Consequently, the Enhanced Intel SpeedStep technology provides the Pentium M processor a flexible, multi-point operating mode, completely self managed, and with a very low CPU and memory unavailability time, which optimizes its power and performance according to demand.

PERFORMANCE

The Intel Pentium M processor architecture delivers breakthrough mobile performance and enables extended battery life in notebook PCs.

Comprehensive information about this processor performance can be found in [9]. In this paper, we choose to demonstrate the power-awareness of the Pentium M processor. We measured the processor performance and average processor power consumption in various operation modes on several benchmarks and compared its performance and efficiency with those of other mobile processors in similar configurations⁹. Efficiency reflects energy per task. Benchmark efficiency is measured by dividing the benchmark performance (1/execution-time) by the average processor power of that benchmark. As will be shown, the Pentium M processor exhibits higher performance and superior efficiency.

The set of benchmarks includes (see more in [Table 1](#)):

- Mobile Representative Office Productivity Workload
- Internet Experience workload
- SPEC CPU 2000 V1.2 [10]

We compared the Pentium M processor with the following: (see [Table 2](#) for detailed system configurations):

- Intel® Pentium® M processor (1.6 GHz/600 MHz)
- Mobile Intel® Pentium® 4 Processor - M (2.4/1.2 GHz)
- Mobile Intel® Pentium® III Processor - M (1.2 GHz/800 MHz)

All system run the Windows® XP® operating system. The operation modes used are as follows:

⁹ The information included in this section was prepared specifically for this paper to provide insight into the success of the design criteria using known benchmarks as a workload. Some measurements were collected on reference boards that are not publicly available. Therefore, these results should be considered only as an estimate for relative performance and efficiency.

- Always On (Max Frequency)
- Portable/Laptop (Adaptive Frequency)
- Maximum Battery (Min Frequency)

All scores are normalized to the Intel Mobile Pentium 4 Processor - M (orange bar) score. The efficiency-performance over power-is obtained from the benchmark performance score divided by the average processor power for the duration of that benchmark.

Always On Mode

In the Always On mode the processor always runs at its highest frequency. This mode is mostly used when the system is connected to an AC power source. This mode demonstrates the inherent performance and power-awareness of the Intel Pentium M processor without utilizing the Enhanced Intel SpeedStep technology.




Figure 12: Always On mode performance and efficiency

Figure 12 presents comparative performance and efficiency results in the Always On mode.

In this mode, the Intel Pentium M processor performs equal or better (2%-25%) than the Mobile Intel Pentium 4 Processor - M on all benchmarks, and is significantly more efficient (2X-3X) than it. This shows that the Pentium M processor power-awareness philosophy works: it does more work and consumes significantly less power in the same thermally constrained environment. The Intel Mobile Pentium 4 Processor - M, designed for higher performance at higher power envelopes, cannot exploit its full performance potential in this thermally restricted environment and has to slow down.

SPECint_base2000 and SPECfp_base2000 are of particular interest here. The Pentium M processor

exhibits nearly double the efficiency advantage over the Intel Mobile Pentium 4 Processor - M on most of the tests in these benchmarks. Several tests, *179.art* and *300.twolf*, exhibit an even greater efficiency gain (over 8X and 5X respectively) mainly due to the large 1MB power-managed L2 cache of the Intel Pentium M processor.

Portable/Laptop Mode

In Portable/Laptop mode, frequency and voltage changes depend on the application demand. This mode is the normal usage mode when the system is not connected to an AC power source. This mode demonstrates the effectiveness of combining the performance and power-awareness of the Pentium M processor with the energy-saving nature of the Enhanced Intel SpeedStep technology to provide end users with breakthrough mobile performance and extended battery life.

In this mode, all three processors operate between their highest and lowest frequency operating points, depending on the amount of work to be done. For processor-intensive workloads, each processor operates at its highest operating voltage and runs at its maximum frequency: the Intel Pentium M processor @ 1.6 GHz, the mobile Intel Pentium 4 Processor - M @ 2.4 GHz, and the Mobile Intel Pentium III Processor - M @ 1.2 GHz. When there is no activity (idle period), each processor runs at its lowest frequency and voltage to conserve energy: the Intel Pentium M processor @ 600 MHz, the Mobile Intel Pentium 4 Processor - M @ 1.2 GHz, and the Mobile Intel Pentium III Processor - M @ 800 MHz. Using the efficient switching algorithms of the Enhanced Intel SpeedStep technology, the Pentium M processor is transparently switched between the highest and lowest frequency and voltage states, giving the user the best of both worlds: maximum performance under demanding applications and lowest power during idle periods.




Figure 13: Portable/Laptop mode performance and efficiency

*Other brands and names are the property of their respective owners.

Figure 13 presents comparative performance and efficiency results in the Portable/Laptop mode. Results show that in this mode, the Pentium M processor performs equal or better (0%-30%) than the Mobile Intel Pentium 4 Processor - M (a similar advantage to the Always On mode). However, the relative efficiency over the Mobile Intel Pentium 4 Processor - M in benchmarks that exhibit periods of lower activity went up, e.g., from 3X to over 4X on the mobile representative Office Productivity workload. This improved efficiency results from the much lower power consumption of the Intel Pentium M processor at its low frequency mode @ 600 MHz compared with the power consumption of the Mobile Intel Pentium 4 Processor - M @1.2 GHz.

Maximum Battery Mode

In the Maximum Battery mode the processor runs at its lowest frequency. This mode is usually used when the user is away from an AC power source for a long time. This mode demonstrates the ability of the Intel Pentium M processor to minimize energy consumption when longer battery life is crucial.




Figure 14: Maximum Battery mode performance and efficiency

Figure 14 presents comparative performance and efficiency results in the Maximum Battery mode. When all three platforms are locked at their lowest processor frequencies, the Intel Pentium M processor running at 600 MHz draws much lower power than the other two processors. This puts the Pentium M processor in a better position to get more work done for the power consumed under the workloads tested. The processor does compromise performance in this mode: it is about 20% slower than the Mobile Intel Pentium 4 Processor - M. However, it is extremely more efficient—about 5X more—allowing it to do significantly more work with the same energy.

Table 1: Benchmark description

Mobile Representative Office Productivity Workload: Targeted to evaluate notebook user experience under popular business-oriented applications in a Microsoft Windows* operating environment. Some usage models represented in this productivity workload include applications from Microsoft Office XP* (i.e., Word 2002, Excel 2002, PowerPoint 2002, Outlook 2002), McAfee* VirusScan*, Adobe* Photoshop*, WinZip* and others.

Internet Experience Workload: Measures PC client performance under a range of popular Internet technologies such as SSL*, XML*, VML*, Java*, etc. using applications such as Adobe Acrobat*, Apple Quicktime*, Cycore*, Cult3D*, Macromedia Flash*, Windows Media Player*, and RealNetworks RealVideo*.

SPEC CPU2000: The industry-standard benchmark that evaluates compute-intensive integer and floating-point application performance [10].

Table 2: System configurations

Platform	Dell Latitude C610	Intel Reference Platform	Intel Reference Platform
CPU	Mobile Intel® Pentium® III Processor-M	Intel® Pentium® M Processor	Mobile Intel® Pentium® 4 Processor-M
CPU Core Freq (MHz)	1200/800	1600/600	2400/1200
CPU Bus Freq	133	400	400
L2 Cache (KB)	512	1024	512
Chipset	Intel 830M	Intel 855PM	Intel 845
Mem Size (MB)	512		
Mem Type/Speed	PC133	DDR 266	
Mem CAS Latency	2-2-2	2-3-3	
Graphics Core	ATI Mobility Radeon M6	ATI Radeon 9000	
Graphics Mem	16MB	64MB	
Gfx Driver	6.13.10.3293	6.13.10.6200	
Screen Resolution	1024 x 768 x 32bpp 60Hz		
HDD Mfr/Model	IBM IC25N040ATCS05-0		
HDD Size, Buffer, RPM	40GB IDE 8MB 5400RPM		
OS, Build, File System	WinXP, SP1 5.1.2600, FAT 32		
LAN	Intel ICH3 Integrated Ethernet Ctrl.	Intel ICH4 Integrated Ethernet Ctrl.	Intel ICH3 Integrated Ethernet Ctrl.

* Other brands and names are the property of their respective owners.

CONCLUSION

The Intel Pentium M processor is Intel's first microprocessor designed specifically for the requirements of tomorrow's mobile PCs. It provides uncompromised performance while observing the thermal and energy requirements and limitations of the mobile platform. Performance-enhancement features were included only if proved to be power-efficient. The processor features many novel power-aware performance mechanisms such as advanced branch prediction, micro-operation fusion, a dedicated stack engine, and the optimized Pentium M bus. It also features the Enhanced Intel SpeedStep technology to reduce energy consumption.

These unique features enable the Pentium M processor to deliver breakthrough performance and enable extended battery life thereby providing users with a superior mobile experience.

ACKNOWLEDGMENTS

The authors thank all the researchers, architects, designers, and validators who took the Intel Pentium M processor from a vision to a real product. Special thanks go to Wilson Huang from Intel's Mobile Platforms Group who conducted the performance and power measurements used in this paper.

REFERENCES

- [1] R. Ronen, A. Mendelson, K. Lai, S.L. Lu, F. Pollack, and J.P. Shen, "Coming Challenges in Microarchitecture and Architecture," in *Proceedings of the IEEE*, Vol. 89, No. 3, March 2001, pp. 325-340.
- [2] *IA-32 Intel Architecture Software Developer's Manual Volume 1: Basic Architecture* at <http://developer.intel.com/design/pentium4/manuals/245470.htm>
- [3] D. M. Brooks et al., "Power-aware microarchitecture: design and modeling challenges for next-generation microprocessors," *IEEE Micro*, Vol. 20, Issue: 6, Dec. 2000, pp. 26-44.
- [4] D.B. Papworth, "Tuning the Pentium® Pro microarchitecture," *IEEE Micro*, Vol. 16-2, April 1996, p. 8.
- [5] C. Mead and L. Conway, *Introduction to VLSI systems*, Addison-Wesley Publishing Company, Boston, Dec. 1980.
- [6] G. Hinton, D. Sager, M. Upton, D. Boggs, D. Carmean, A. Kyker, and P. Roussel, "The Microarchitecture of the Pentium 4 Processor," *Intel Technology Journal*, Issue 1, 2001, article 2.
- [7] M. Bekerman, A. Yoaz, F. Gabbay, S. Jourdan, M. Kalaei, and R. Ronen, "Early Load Address Resolution via Register Tracking," in *Proceedings*

of the 27th International Symposium on Computer Architecture, June 2000, pp. 306-315.

- [8] Mobile Pentium® III processors—Intel SpeedStep® Technology
<http://www.intel.com/support/processors/mobile/pentiumiii/ss.htm>.
- [9] Intel® Centrino™ Mobile Technology Performance Brief
<http://www.intel.com/performance/resources/mobiletechnology/>
- [10] SPEC CPU2000 V1.2
<http://www.specbench.org/osg/cpu2000/>.

AUTHORS' BIOGRAPHIES

Simcha Gochman is a principal engineer/architect with Intel's Mobile Platforms Group in Haifa, Israel, leading the microarchitecture definition activities of the Intel Pentium M processor and its future derivatives. Simcha has been with Intel for 18 years. Earlier in Intel he led the microarchitecture definition of the Intel Pentium Processor with MMX™ technology and was involved with the design of the 80860 processor and the 80387 numeric coprocessor. Simcha received his M.Sc. degree from the Technion, Israel Institute of Technology in 1984. His e-mail is simcha.gochman@intel.com.

Ronny Ronen is a principal engineer/researcher and director of Intel's Microprocessor Research Lab in Haifa, Israel, focusing on microarchitecture research. Ronny was heavily involved in the definition stages of the Intel Pentium M processor. He has been with Intel for 22 years. Earlier in Intel he led the compiler and performance simulation activities in the Intel Israel Software department. Ronny received his M.Sc. degree from the Technion, Israel Institute of Technology in 1979. His e-mail is ronny.ronen@intel.com.

Ittai Anati is a senior architect with Intel's Mobile Platforms Group in Haifa, Israel, focusing on the microarchitecture of the Intel Pentium M processor family. Ittai has been with Intel for 14 years. Earlier in Intel, Ittai was part of the i860XP, Pentium Overdrive, and Pentium with MMX technology design teams. He received his B.Sc. degree from the Technion, Israel Institute of Technology in 1990. His e-mail is ittai.anati@intel.com.

Ariel Berkovits is a senior architect with Intel's Mobile Platforms Group in Haifa, Israel, leading the Intel Pentium M processor performance analysis, projection and performance validation activities. Ariel has been with Intel for 11 years. Earlier in Intel he led the

™ MMX is a trademark of Intel Corporation or its subsidiaries in the United States and other countries

development of performance simulators and was involved in the definition of graphic chipsets. He holds a B.Sc. degree from the Hebrew University in Jerusalem. His e-mail is ariel.berkovits@intel.com.

Tsvika Kurts is a senior architect with Intel's Mobile Platforms Group in Haifa, Israel, focusing on the Intel Pentium M processor bus, platform, and debug hooks. Tsvika has been with Intel for 18 years. Earlier in Intel Tsvika was part of the Intel Pentium® Pro processor bus architecture and system validation team and was involved in the Intel Pentium 4 processor bus protocol development. He received his B.Sc. degree in 1984 and his M.Sc. degree in 1992, both from the Technion, Israel Institute of Technology. His e-mail is tsvika.kurts@intel.com.

Alon Naveh is a senior architect with Intel's Mobile Platforms Group in Haifa, Israel, focusing on processor and platform power management. He received his B.Sc. degree from the Technion, Israel Institute of Technology in 1983, and holds an MBA degree from San Jose State University. Alon was involved in the definition of the Intel Pentium M processor power management, PCI-E, and the Odem chipset. Prior to Intel, Alon worked in Motorola Semiconductor and in National Semiconductor. His e-mail is alon.naveh@intel.com.

Ali Saeed is a senior systems architect in Intel's Mobile Platforms Group. His primary research area includes performance and power profiling of processors and chipsets with an emphasis on performance, power, and battery life estimations. Prior to joining Intel, Ali was manager of performance analysis in Compaq's Commercial PC Group. He holds a Bachelors degree in Electrical Engineering from the University of Houston. His e-mail is ali.saeed@intel.com.

Zeev Sperber is a senior architect with Intel's Mobile Platforms Group in Haifa, Israel. He received his B.Sc. degree from the Technion, Israel Institute of Technology in 1984. Zeev has been with Intel for 19 years. He was involved in the definition of the Intel Pentium M processor out-of-order and execution sections. Earlier in Intel, Zeev was part of the Intel Pentium processor with the MMX technology team. His e-mail is zeev.sperber@intel.com.

Robert C. Valentine is a senior architect with Intel's Mobile Platforms Group in Haifa, Israel, focusing on the Intel Pentium M processor's fetch and decode cluster. Robert has been with Intel for 13 years and has worked both as a microarchitect in the Haifa's Microprocessor Platforms Group and as a member of Haifa's Research team. Prior to Intel, Robert worked in Prime Computer and Honeywell Information Systems. He holds an M.Sc.

degree in Computer Engineering from Boston University. His e-mail is bob.valentine@intel.com.

Copyright © Intel Corporation 2003. This publication was downloaded from <http://developer.intel.com/>.

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

For further information visit:

developer.intel.com/technology/itj/index.htm



Intel® Technology Journal

Intel® Centrino™ Mobile Technology

System-Level Validation of the Intel® Pentium® M Processor

System-Level Validation of the Intel® Pentium® M Processor

Isic Silas, Mobile Platforms Group, Intel Corporation
Igor Frumkin, Mobile Platforms Group, Intel Corporation
Eilon Hazan, Mobile Platforms Group, Intel Corporation
Ehud Mor, Mobile Platforms Group, Intel Corporation
Genadiy Zobin, Mobile Platforms Group, Intel Corporation

Index words: Pentium® M processor, System Validation, Periodic SMI, Bus Marginality Testing

ABSTRACT

The Intel® Pentium® M processor is the first Intel microprocessor with major architectural and micro-architectural changes in performance and power and frequency optimizations, which directly entered the mobile market, i.e., without a previously validated desktop version. A very tight Intel Pentium M processor post-silicon period from first silicon to launch and lack of IA-32 testing and debug expertise by the Mobile OEMs, dictated a different approach to validate the CPU, at both the Intel level and the OEMs' level. The main goal was to ramp up the Pentium M validation capabilities very fast and uncover all silicon bugs before they were reported by OEMs (to prevent the time and effort it would take to debug sightings in the OEMs' environment), while maintaining the tight OEM development cycle from samples to launch.

This paper describes a novel complex post-silicon system validation methodology, that enabled us to deliver a healthy Intel Pentium M processor for the launch of the Intel® Centrino™ mobile technology, despite the tight validation schedule. This was achieved by Intel's uncovering all logic and circuit issues so no silicon-related bugs were found by customers, and by reducing dramatically the system-level-failures' debug time.

INTRODUCTION

The Intel Pentium M microprocessor, a key element of Intel Centrino mobile technology, is the first Intel

processor with major architectural and micro-architectural changes in performance and power and frequency optimizations, which directly entered the mobile market, i.e., without a previously validated desktop version. Features of the Pentium M processor such as Enhanced Intel SpeedStep® technology, new branch prediction algorithms, and micro-ops fusion could potentially be a source of severe logic issues, while timing enhancements and low voltage working points (to reduce the CPU's power consumption) could result in marginality problems in the CPU's core and the I/O circuit.

The traditional CPU post-silicon platform-level validation cycle, at Intel, consists of parallel development and testing of system validation (SV), compatibility verification (CV), and OEMs programs.

The objectives of CV are to ensure that the microprocessor functions correctly in the standard system, with existing operational systems and commercial applications, and that it runs on customer reference boards equipped with commercially available (for the moment) hardware configurations. The drawbacks of using CV are insufficient visibility and controllability on CPU features usage and testing coverage, and difficulties in debugging of failures that in many cases are platform hardware or software rather than CPU-related.

The OEMs' primary goal is to maintain the tight customer development cycle from samples to launch. Their validation process is very similar to that of CV, with a focus on overall system functionality.

SV is based on testing the microprocessor functionality in specially designed, SV-hardware configurations, by running unique SV random tests that typically have nothing in common with any commercial software. The

®Intel, Pentium, Intel SpeedStep, and Intel Centrino are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

special SV hardware aims to simulate all commercial platform configurations including all the add-in cards options. Similarly, the SV random testing aims to simulate all the existing, as well as yet to be developed, commercial software code combinations. These two factors provide very good controllability on validation coverage, since they allow for the accurate generation of any desired event and IA-32 scenario. The SV environment also assumes an efficient failure localization, which guarantees fast issue debugging and identification.

The very tight Intel Centrino mobile technology program schedule was mainly due to the short period between the Intel Pentium M processor first samples and launch. By not having debug hooks and expertise from the mobile OEMs, we had to approach the validation cycle in a different way at both the Intel and OEM level. The goal of SV was to ramp up the Pentium M processor validation capabilities in the fastest way and uncover all silicon logic and circuit marginality issues before they were reported by OEMs. Since OEMs got early Pentium M processor engineering samples (on average, approximately two weeks after SV got units) before validation was completed, uncovering issues before they were uncovered by the OEMs was a huge challenge. At the same time, an early enabling program was developed to allow OEMs to start their platform integration and validation; to test chipset, memory, and graphics subsystems; as well as other platform components, even before the Pentium M samples became available.

This paper describes different aspects of the Intel Pentium M processor's complex post-silicon validation program and details the major results.

THE INTEL PENTIUM M PROCESSOR SYSTEM VALIDATION

The Intel Pentium M processor system validation (SV) methodology is based on the following major components:

- A unique SV platform that differs from a standard motherboard because of specially developed hardware agents connected to the CPU front-side bus and other pins.
- A package of random instruction-based SV testing software, that ensures excellent coverage for logic and circuit marginality aspects.
- A Periodic SMI (PSMI) methodology for automatic and very fast reproduction of system-level failures on the Pentium M processor software model (RTL) or the Debug Tester.

Intel Pentium M Processor System Validation Platform

“Golan”: Intel Pentium M processor/Intel® 855PM chipset SV Platform

The design of the Pentium M SV platform focuses on enhancing the following capabilities that are a must for qualitative and efficient validation of modern microprocessors:

- *Coverage/controllability.* This enables the most complex scenarios on CPU busses and other pins, to be easily programmed via testing software. Thus the processor bus behavior can be simulated, when used with various chipset types and add-in cards, by using only one type of chipset and without any add-in cards.
- *Determinism.* This ensures that you can immediately reproduce system failures, clock by clock.
- *Automation.* Automation allows for remote execution of SV tests and for remote/programmable modification (shmoo) of CPU operating parameters (e.g., voltage, frequency, temperature).




Figure 1: Block diagram of Golan, an Intel® Pentium® M processor/Intel® 855PM chipset system validation platform

A block-diagram of the Intel Pentium M processor configured with an Intel® 855PM chipset SV platform, named Golan, is shown in Figure 1.

The “brain” of this platform, the Mobile Agent, is implemented in a Field Programmable Gate Array (FPGA) and controls a wide set of features to support all SV requirements. The essential part of the Mobile Agent is the Programmable Signal Generator (PSG). This device

was designed to enable simultaneous injection of signals during CPU code execution. Those signals, which are activated in different modes (e.g., single or burst, level or pulse), include legacy signals like NMI, SMI, and INTR, which are connected directly to the processor pins, and various break events, which are sent from the IO Control Hub (ICH). Each PSG signal duration and timing parameter can be programmed via software. In addition, the Mobile Agent can identify different power management events and use them as trigger sources for PSG signals. It guarantees good coverage of various border scenarios, essential for mobile platforms, that previously were subject to real CPU problems or other platform component issues. It should be emphasized that PSG completely eliminates the need for commercial Peripheral Component Interconnect (PCI) or graphics cards for the SV board, ensuring at least the same level of CPU bus stressing and significantly better reproducibility.

The Golan SV board also includes a power management Rule Checker/Tracker designed to monitor the CPU's and the system's power management modes transitions, and the Enhanced Intel SpeedStep rules.

Very accurate, automated control over the system clock scheme enables the CPU Processor System Bus (PSB) frequency modification (shmoo), which is widely used by SV for Pentium M circuit marginality testing, together with the Core voltage control (shmoo) capability, which is also implemented in the Golan board.

Among other Golan SV board features, are power consumption (V,I) measurements and history recording, reset and straps, and a Joint Test Action Group (JTAG) standard controller.

Fork Transactions Generator: SV Processor System Bus (PSB) Agent

Another leading methodology for the Pentium M processor SV platform is the usage of active bus agents to provide bus activation right on the spot. The biggest advantage of these agents over end point agents (such as PCI cards) is their capability to inject aimed transactions, triggered by normal bus events, and their high success rate in failure reproduction. The most important active bus agent used for the Pentium M processor validation was the Fork Transactions Generator (Fork TG), residing on the processor system bus (PSB).

Because it was designed for mobile systems, the Pentium M processor bus supports uni-processor mode only, thus, preventing the usage of another CPU for its caches and bus unit stressing. Therefore, Fork TG, implemented as a re-programmable FPGA, demonstrates a new concept in the architecture of a uni-processor validation platform. Specially designed to be placed between the processor and the Memory Control Hub (MCH), Fork TG is acting in a

true uni-processor environment. Nevertheless, neither the processor nor the MCH are aware, logically, of its presence. When injecting PSB transactions, Fork TG manipulates each of the two chips to act as if the counter-chip is the source for those transactions. This innovative solution brings plenty of new features and benefits to the system validation environment.

Fork TG activity is based on controllable triggering and inter-operating injection streams. The notion of injection implies placing transactions on the processor system-bus and thus creating traffic. The density of this traffic as well as its duration are regulated by Fork TG configuration parameters programmed via validation software as part of the SV tests.

Three modes have been defined in order to regulate Fork TG injection streams: pre-defined, shadow, and self-generated. The pre-defined mode is mainly used to exercise the cache cluster (L1 and L2) and the Memory Ordering Buffer (MOB) in the processor. In this mode, Fork TG monitors the bus, looking for an event (e.g., a bus transaction) that triggers it to stream the pre-defined package of transactions, which were programmed beforehand to Fork TG's external RAM.

The shadow mode was conceived for exercising the cache cluster of the CPU, along with the processor bus unit. In this mode the transaction sampled on the PSB undergoes some updates in Fork TG, and the newly created transaction is placed back on the bus as a part of the Fork TG injection stream. In other words, the original transaction is 'shadowed' by the newly generated one, altered according to some rule. This is the essence of this operation. Possible alteration rules, chosen during SV test generation, include replacement of the request type (e.g., a Read Invalidate operation instead of a Read Data operation), replacement of the memory type, a change in the length of the request, or an update of the transaction address within the page (4K) boundary.

In Self-Generated mode, the PSB is stressed by bursting out injections irrespective of current bus events. Three parameters, start address, end address, and address step, along with the number of loops, set the duration of the self-generated test.

The Fork TG has been heavily used throughout the Pentium M processor post-silicon validation period, mostly for caches and the PSB stressing. This resulted in the detection of two real silicon issues.

Random Instruction Testing: Pentium M processor System Validation Methodology

There are several different methodologies for processor system-level validation. The Pentium M processor system validation was principally based on the Random

Instruction Testing (RIT) concept. This means that CPU correct functionality, logic and circuit wise, is being checked by a tremendous amount of discrete tests aimed at covering all possible architectural and micro-architectural scenarios enabled by Spec. Each random test (seed), typically runs from reset, emulates some artificial operational system with all its attributes (e.g., descriptor and interrupt tables, exception handlers, paging hierarchy), and also contains several thousand instructions of “user’s code,” running under this pseudo-Operating System (OS). Both OS attributes and user’s code are maximally random, in terms of instruction sets, memory allocation for OS level structures, code and data, caching policy, etc. Random programming of SV platform residing agents, namely, Mobile Agents (signals, events) and Fork TG (processor system bus transactions) are also under seed responsibility.

Random test generation is being done in a software environment, specially developed at Intel. It includes a kernel, which is applicable to any current and future microprocessor, and modules that define processor-specific test content. These modules, which, by themselves, are very complex software projects, are responsible for the coverage of microprocessor health. Test generation rules implemented in these modules result in the success of silicon in mass production, e.g., absence of escaped bugs.

The Pentium M SV test content has been developed, taking into account these processor multiple architectural and micro-architectural enhancements, as compared to previous generations of CPUs, and the unique mobile features of the Pentium M processor. SV engineers spent much time with architects and designers in order to identify the most risky areas of this processor, before implementing their coverage in testing software. Among the areas that required special attention were the following:

- Enhanced Intel SpeedStep technology assuming processor frequency transitions during regular code execution. CPU response to different events, like power management or external interrupts randomly occurring during any of the transitions, was a special focus of SV
- Various scenarios related to power management transitions of the CPU, e.g., to/from sleep and deep sleep states, with heavy involvement in break events, interrupts, APIC messages, and snooping
- Novel micro-architectural features, e.g., micro-ops fusion
- Marginal data paths capable of causing system failures, specifically at high frequencies

Random Instruction Testing methodology assumes the execution of many millions of seeds before the required coverage of all CPU features is achieved. Therefore, the SV Lab infrastructure includes up to 100 Golan platforms and several hundred test generators able to produce seeds 24-hours a day, executing approximately 1 trillion instructions per week. Enormous effort was put into enhancing generation and execution throughput, resolving Lab network problems and achieving automatic control over lab resources. As a result, around 10 billion random seeds, each containing several thousand random instructions, were run by the Pentium M processor throughout its validation period from first silicon arrival to launch.

One more important thing that influences the post-silicon validation success, especially during the first critical weeks, is the cleanliness of testing software. To ensure the software is clean, the Pentium M processor test content has passed massive dry runs on Intel Pentium® 4 platforms, on Golan SV boards equipped with the Erez interposer, described below, and on Pentium M models. As a result, when silicon arrived, the system validation team faced few false alarms and could concentrate just on debugging real CPU issues.




Figure 2: Cumulative bug detection rate by Intel Pentium® M processor SV from first silicon to launch

All factors described in this section have resulted in very fast bug rate detection once Pentium M processor samples arrived at the SV Lab. Figure 2 shows the cumulative bug detection rate from first silicon to launch, and it clearly shows that 50% of the logic issues found in the Pentium M processor were uncovered during the first two to three weeks of validation. Also, the latest bug that needed a silicon fix was uncovered after 20 weeks of validation, while all further issues were of low severity and were closed via errata.

Periodic SMI (PSMI) Failures Reproduction Methodology

An essential component of any silicon validation program is the capability to quickly and efficiently reproduce system failures on the RTL model of the chip under test, in the case of a logic issue, or on the Tester, in the case of a circuit issue. The major problem of reproduction is to ensure that silicon behavior on the system is absolutely synchronized with that of the RTL/tester. This is because most failures occur far beyond the reset point, while all processor arrays and registers, e.g., caches, TLBs, branch predictor, contain some history, which is not visible externally. The Periodic System Management Interrupt (PSMI) methodology, developed at Intel approximately five years ago, is aimed at resolving this problem and providing full synchronization between microprocessor behavior on the system and simulation environment during the failure reproduction phase. This methodology assumes an injection of periodic signals on the SMI pin of the CPU, while the latter executes a failing test or application, forcing it to enter periodically into a specially developed SMI handler, as shown in Figure 3.

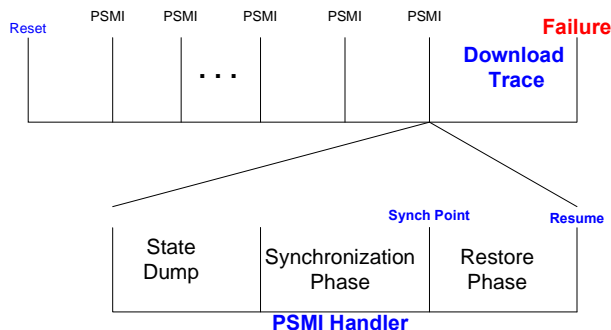


Figure 3: Schematic description of periodic SMI methodology

The Logic Analyzer trace, downloaded for further reproduction, should reflect full CPU external activity (taken from all its pins) starting from the SMI closest to the failing point. The most critical parts of the sophisticated PSMI handler are the Dump and Synchronization Phases as shown in Figure 3. They provide a dump of the content of the CPU internal state to memory, making it visible on the processor system bus, and an initialization/flush of some arrays and state machines. Therefore, the CPU behavior, after the Sync point of the handler, is absolutely deterministic, and the trace is ready for conversion and reproduction on the RTL model or Tester.

PSMI technology has been successfully used for system-level validation of Intel Pentium® MMX™, Pentium® III

and Pentium® 4 processors. The mobile-oriented features of the Intel Pentium M processor brought new challenges derived from the following facts:

- SMI is heavily used in mobile platforms; therefore, the potential conflict between regular SMI and PSMI should be eliminated. This has been done by separating SMI and PSMI both in the internal CPU design, and at the system memory space level.
- Enhanced Intel SpeedStep technology, implemented in the Pentium M processor assumes internal clock frequency changes during regular CPU running. These frequency transitions have to be captured by the PSMI handler and transferred correctly to the simulation environment.
- Power management features that cause partial shutdown of Pentium M processor internal and external clocks are heavily used in the CPUs of mobile platforms. These also introduce serious reproduction issues.

In order to ensure full readiness of the PSMI flow in all corner cases, and to accommodate the above challenges, PSMI was validated extensively in the pre-silicon period. A comprehensive test plan was written, which included 60 focus tests and over 100 random tests. Each test emulated the PSMI flow—a trace being taken from a simulated system and driven into the target simulation environment. A simulation acceleration machine was heavily used to increase the execution throughput.

Despite the above increased complexity, the PSMI methodology has demonstrated impressive results during Pentium M processor system-level validation. The PSMI flow was ready for use and reproduced the first logic issue the first week after first silicon arrival. Overall, the post-silicon validation period, 100% (!) of all the logic issues brought from the system, have reproduced on the RTL model, with an average latency of only 1-2 days. As for circuit issues mostly related to internal speed paths, Intel's success rate is approaching 94% – much higher than in any previously validated microprocessors.

PROCESSOR SYSTEM BUS MARGINALITY VALIDATION

Due to the low-voltage and high-frequency Processor System Bus (PSB) of the Pentium M processor, a thorough signal integrity validation became critical to ensure the design robustness of customer reference boards and OEM's mobile systems. Typically, bus marginality is measured in two domains, namely, input buffer reference

voltage V_{ref} , and data/strobe timing Δt . A two-dimensional (V_{ref} , Δt) area of reliable system functionality forms an ellipse, called an Eye Diagram, as shown in Figure 4.

One can see that the boundaries of an Eye Diagram reflect conditions where the system starts failing, while its center provides an optimal design point for the V_{ref} and Δt parameters. A single Eye Diagram measurement requires several hundred iterations to be accomplished at different points, while measurements should be repeated under a wide spectrum of conditions and interconnection parameters, e.g., temperature, CPU or chipset skew units. This makes the bus marginality testing extremely time consuming (typically, about four to eight hours per “eye”).




Figure 4: An example of a bus marginality “Eye diagram”

To intensify the process of Pentium M processor analog validation, a novel tool, named the Bus Marginality Tester (BMT) was developed by the Pentium M Signal Integrity and Analog Validation team. The BMT is actually a package of hardware and software tools that is used to measure automatic signal integrity bus marginality. Its structure is hierarchical, as shown in Figure 5, consisting of three levels: Target, Host, and Server. Bus stress runs on the Target level (actually a board under test), Host controls the Eye Diagram building process and is connected to the Target via a PCI-to-PCI bridge. It regulates Eye Diagram parameters on the Target level via a specially designed margin card (V_{ref}) and BIOS settings (Δt). A single server manages several Target/Host blocks via a Local Area Network (LAN) and provides the user interface with an operator. It allows massive automatic parallel execution on several Target/Host blocks from one server. All Eye Diagram accumulated information coming to the server is stored in a special database and is subject to automatic sophisticated analysis.




Figure 5: Block diagram of Bus Marginality Tester (BMT)

BMT usage in the Pentium M processor project brought very impressive results. First of all, an average Eye Diagram measurement time was reduced from hours to ten to fifteen minutes. It allowed us to build about 40,000 functional “eyes,” providing wide coverage of bus interconnect validation, with a Shmoo of many parameters: package and board impedance, silicon skew material, temperature, IO voltage, various bus stress patterns, various memory combinations, etc.

BMT has uncovered and helped to resolve several serious system problems. Among them are two severe Small Outline Dual In-line Memory Module (SO-DIMM) memory issues, Double Data Rate (DDR) Logic Analyzer Interface issues, and DDR data/strobe de-centering problems. The BMT was effectively used for the fast comparison of signal integrity bus robustness design of different OEM motherboards. Thus, we were able to give OEMs early feedback on the robustness of their board design.

Erez Interposer: Early Enabling Program

A very tight Intel Pentium M processor post-silicon period from first silicon to launch has put significant pressure on OEMs’ platform testing and integration schedule. A major breakthrough in this area was early availability of a specially designed small form-factor, full speed, pin compatible interposer, named Erez. Equipped with a Pentium 4 processor, the Erez interposer was inserted into the CPU socket of the Pentium M system, and it successfully imitated the functionality. This program paved the way for platform designers of OEM systems, reference boards, and SV platforms to debug their hardware and ramp-up the critical mass of systems, months before Pentium M processor silicon arrived. In total, eight OEMs have participated in the early enabling program, using Erez interposers. Each OEM used it for the development of several mobile platforms. In parallel, it also enabled BIOS and software developers to test and clean their products.

RESULTS SUMMARY

An efficient System Validation program has been developed for the Intel Pentium M processor, a key element of Intel Centrino mobile technology. The objective of system validation (SV) is to ensure the delivery of a healthy Pentium M processor, from the point of view of logic and circuit marginality, in an extremely tight schedule, taking into account that it is the first time a brand new processor enters the mobile market, without any sophisticated IA-32 testing and debug capabilities.

Intel uncovered 100% of the Pentium M logic and circuit issues with no silicon-related bugs found by customers. In addition, 50% of the tens of the silicon issues were detected during the first two weeks of validation, showing the high quality of the SV test suite and the excellent throughput of moving failures from system to RTL and Tester, using the PSMT tool. A combination of a unique SV test environment, the PSMT tool, Erez early enabling program, and the BMT allowed Intel to deliver the Pentium M processor in time for the launch of Intel Centrino mobile technology.

DISCUSSION

The success of the Intel Pentium M processor system validation program was predetermined by the following factors:

- Validation engineers had extensive expertise in processor and platform-level architecture, processor testing, and validation platform design
- Validation platforms and testing software were defined and developed in a very cooperative atmosphere
- Innovative coverage and validation efficiency-oriented ideas (Fork Transactions Generator, Bus Marginality Tester, and Erez) were driven and implemented
- Very thorough preparations were made to guarantee the readiness of the validation platform, testing software, and debugging tools (PSMT) on the eve of silicon arrival

As modern microprocessors become more and more sophisticated, and as market competition tightens the post-silicon schedule, it is essential to combine efficiently all the above factors, as was done for the Intel Pentium M processor.

AUTHORS' BIOGRAPHIES

Isic Silas received his B.Sc. degree in Electrical Engineering from the Technion (Israel Institute of Technology) in 1988. He is the Intel Israel Mobile Group,

System Assurance department manager, responsible for developing hardware and software capabilities for components and platform-level testing, to ensure the delivery of healthy products and platforms. Isic joined Intel in 1988, working on many aspects of platform validation. He started as a hardware engineer developing validation platforms, and then moved on to validation software and validation and debug tools. He is now perceived as an Intel expert in the area of validation. His e-mail is isic.silas@intel.com

Igor Frumkin is a validation group manager in Intel's Israel Design Center, currently leading system-level validation of Intel microprocessors for mobile platforms. Igor joined Intel in 1995. He received his B.Sc. and M.Sc. degrees in Electrical Engineering from the Institute of Aviation Instruments Making in Leningrad (the former Soviet Union) in 1981, and his D.Sc. degree in Quality Assurance from the Technion (Israel Institute of Technology) in 1995. His e-mail is igor.frumkin@intel.com

Eilon Hazan is a debug tools team manager in Intel's Israel Design Center, leading the PSMT tool and other debug capabilities for system-level validation in the Mobile Platform Group. Eilon joined Intel in 1995. He received his B.Sc. degree in Electrical Engineering from the Technion (Israel Institute of Technology) in 1997. His e-mail is elion.hazan@intel.com

Ehud Mor is a Hardware Development group manager in Intel's Israel Design Center, responsible for the design of hardware solutions for system validation. Mor received his B.Sc. degree in Electrical Engineering from the Technion (Israel Institute of Technology) in 1979. Before joining Intel in 1994, Mor was a senior design engineer at Integrated Device Technology Inc., RAFAEL and Elbit Electronics. His e-mail is ehud.mor@intel.com

Genadiy Zobin is the Signal Integrity Engineering Group manager in the Intel Israel Development Center. He works on interconnect signal integrity design and validation of mobile platforms. Genadiy joined Intel in 1991. He has wide experience in platform design and validation, including logic, circuit, ASIC, and interconnect design. He received his M.Sc. degree in Electrical Engineering from Leningrad Polytechnic Institute in 1977. His e-mail is genadiy.zobin@intel.com

Copyright © Intel Corporation 2003. This publication was downloaded from <http://developer.intel.com/>.

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

For further information visit:

developer.intel.com/technology/itj/index.htm



Intel® Technology Journal

Intel® Centrino™ Mobile Technology

**Intel® Pentium® M Processor
Power Estimation, Budgeting,
Optimization, and Validation**

Intel® Pentium® M Processor Power Estimation, Budgeting, Optimization, and Validation

Dani Genossar, Israel Mobile Platforms Group, Intel Corporation
Nachum Shamir, Israel Mobile Platforms Group, Intel Corporation

Index words: power estimation, power budgeting, power reduction, Intel SpeedStep, thermal throttling.

ABSTRACT

This paper describes the approach taken by the design team to ensure that the Intel® Pentium® M processor will be a compelling microprocessor for the Intel® Centrino™ mobile technology-based platform. We discuss the power estimation flow and describe the power and thermal driven architecture and circuit enhancements of this architecture.

The Intel Pentium M processor is Intel's first CPU to provide an improved multi-gear low-overhead mechanism with Intel SpeedStep® technology and an advanced Thermal-Throttling-2 implementation. In normal conditions, the operating system can dynamically adjust the processor speed according to the performance requirements, allowing a power-on-demand operation. To protect the device from overheating during extreme power transitions, the Intel Pentium M processor uses a combined voltage and frequency control that provides efficient cooling with minimal impact on performance.

Tools and methodologies were developed by the design team to extract and analyze the power data for each of the basic functional blocks of the Intel Pentium M processor. The flow core component, the Stochastic Dynamic Power Estimator (SDPE), is a novel statistical power estimation tool. The power estimation activity provided the following:

1. Per block power estimation and break down to support the setting of the power Plan Of Record (POR).
2. Data to identify and plan power-reduction strategies.
3. On-going verification to ensure convergence towards the power POR.

To meet the aggressive POR, the design team used the data generated from the SDPE tool and invested heavily in power optimizations. Silicon-based measurements demonstrate the success of the power reduction work and show an excellent correlation with the power and thermal pre-silicon estimations.

The end result is that the Intel Pentium M processor is a compelling mobile product delivering high performance and improved battery life within a restricted mobile thermal envelope. The developed architecture, methodologies, and knowledge base pave the way for the design of Intel's next-generation power-efficient mobile products.

INTRODUCTION

The design of a mobile processor is guided by the thermal limitations of the platform. A thin form factor is desired, affecting the battery size and the efficiency of the cooling system. Reduction of active and idle power consumption provides longer battery life and allows operation at higher frequencies. The Intel Pentium M processor design team developed and applied pre-silicon analysis utilizing a statistical approach that enabled early estimation well before the design was stabilized. As a result we were able to identify power-saving opportunities and overheating-prone areas early in the design cycle. The Intel Pentium M processor power saving amounted to 35% of the maximal active power of the processor core. A novel multi-gear low-overhead Intel SpeedStep technology mechanism was implemented to provide frequency on

Pentium M is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Centrino is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Intel SpeedStep is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

demand, and the advanced Thermal-Throttling-2 significantly improved the utilization of the available cooling capabilities, contributing to the overall result of making the Intel Pentium M processor an attractive mobile processor with high performance and low power consumption.

To accomplish our reduction goals the design team needed to first set a POR for every functional unit in the processor. Once a POR was set, we needed to track adherence to this POR and provide the designer with the relevant feedback. The power design work included a wide range of activities. We started by developing our tools and methodologies. Analyses of previous-generation CPUs served as a baseline for determining the power POR for the Intel Pentium M processor and for identifying power-saving opportunities. When the processor RTL model became functional we started developing tests and power debug techniques and analyzed various high- and low-power applications. During the design phase we identified power bugs and additional power-saving opportunities. When the circuit database was stabilized, we analyzed the thermal stress and guided the setting of the temperature control logic.

Intel Pentium M processor silicon-based measurements show that the extensive power-reduction work was successful, allowing delivery of higher performance and higher frequency without increasing the power envelope. The high power vs. idle ratio was significantly improved, further reducing the average power consumption and extending the battery life. A review of the pre-silicon dynamic power estimates showed an excellent correlation with overall error rates in the order of 5%, and it also showed accurate identification of the device hot spots.

In this paper we describe the power design efforts including the estimation and analysis flows, architecture enhancements, and the actual reduction work. We conclude with a few examples demonstrating our thermal throttling efficiency and the accuracy of our power-simulation models.

DYNAMIC POWER ESTIMATION FLOW

The dynamic power is dissipated on charging the circuit parasitic capacitances and is linearly dependent on the number of signal toggles. To calculate the dynamic power we define the activity factor as the average number of zero-to-one transitions during a clock cycle. To extract the activity factors we performed logic simulation and counted toggles and state statistics.

For each node we calculated the dynamic power using $P = AF \cdot f \cdot C \cdot V_{CC}^2$ where AF is the activity factor,

f is the frequency, C is the lumped capacitance, and V_{CC} is the supply voltage.

The dynamic power analysis flow is described in Figure 1. The first stage is the development of the tests. For high-power tests, we maximize the command execution throughput, taking into account the processor parallel and out-of order capabilities. For low-power tests, we utilize the processor bottlenecks to achieve low execution throughput. The high-power tests are used for examining the thermal solution and power delivery efficiency. The low-power tests help to identify power bugs and savings opportunities.

We developed low-power tests that are based on execution bottlenecks. Use of slower commands such as division or square root cause the system to stall by filling the input command queue. The other types of idle tests are characterized by an empty command queue, for example, due to a second-level cache miss during code fetch. In high-power tests, we optimize the command flow according to the internal dependencies, execution ports configuration, and available dispatch and retire bandwidth. Some of the high-power tests maximize the overall power consumption, while others stress target units.

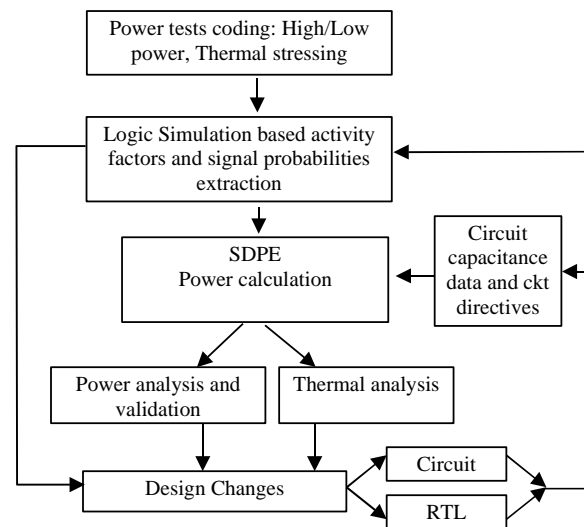


Figure 1: The power design cycle of the Intel Pentium M processor. Logic simulations were performed to collect activity statistics. The data were used by SDPE to calculate the power estimates. Analysis of the power results yielded design changes in the circuit and in the RTL.

To calculate the dynamic power we first extract the activity statistics by logic simulation of the various power

tests. The activity statistics of the Functional Unit Block (FUB) inputs were used by our power estimation flow, the Stochastic Dynamic Power Estimator (SDPE,) to generate vectors and calculate the FUB power. The SDPE estimates activity factors by transistor-level logic simulation. The SDPE flow uses the Monte-Carlo approach to generate the FUB input vectors according to given activity statistics. The input vectors are injected into the FUB and propagated by a unit delay transistor-level simulation [1, 2]. Thanks to the statistical implementation we estimated the power from the earliest design stages, even when mismatches between RTL and schematics existed on the FUB interface pins. The impact of missing data on accuracy was measured by comparing various default activity assignments.

The statistical approach drastically reduces the data storage requirements and saves simulation time. In some cases, the statistical approach affected the accuracy of the power estimation. For most FUBs we obtained reasonable results and successfully identified the main power consumers and power-saving opportunities. Excellent correlation was found between the overall pre-silicon power estimates and the results of the Intel Pentium M processor power measurements.

Power Analysis and Validation

The results of the RTL simulation and SDPE runs provide data for a variety of applications. Every few weeks we calculated the power roll-up and compared each unit to the POR commitment. The Max power test was chosen as the roll-up reference due to the wide coverage of the chip resources. Other high-power tests were used to cover FUBs that were not operated by the Max power test. Idle tests were used for power design validation and for identifying additional power-saving opportunities.

To monitor the clock-gating efficiency we analyzed the idle tests activity. Unlike the power-down state, during the idle tests the global clock is active and local logic gates the clock in the idle units. Observing the FUBs that are still active is an excellent way to identify gating candidates and power-design bugs. To implement the proposed gating one must justify the additional design effort. SDPE provides “what if” analysis for estimating the power-saving return on investment (ROI). The analysis is performed by overriding the clock enable signals.

We applied several techniques to identify power bugs and additional gating candidates. Wide vectors that are active during the idle tests were identified in RTL simulation and SDPE data. The root cause was mapped in most cases to one of the following categories: domino-driven gates, wide latches with a non-gated clock, or bugs in the clock-gating control logic.

Using schematics analysis we identified all gated clocks and compared the activity in various tests. We identified clocks that operated during idle tests more than during high-power tests and clocks that operated at all times.

During the power debug we encountered power-related bugs that did not have a functional impact, hence they were not detected by traditional validation methods. To detect such bugs, tests were modified to step between power modes. We verified that the power consumed was not affected by the history of the machine prior to the tested time window.

The estimation and automated analysis flows were developed and regularly run by a small group of three engineers. The total simulation run time is 13 hours for the core FUBs and an additional 30 hours for all the Level 2 cache arrays. Power analysis runs were performed every few weeks during the design phase of the Intel Pentium M processor.

Logic Optimizations

The most effective power-saving opportunities are to be found at the logic level. First priority was given to clock gating in order to prevent circuits from running when not used. The hierarchy in which this gating is to be implemented must be chosen carefully. Choosing a too high level hierarchy can significantly reduce the gating opportunities while a too low level hierarchy can end up with a control logic that consumes more power than is saved by the gated logic. By carefully structuring the microarchitecture, the machine is optimized to focus on the required activities and minimize the redundant ones.

New control logic was added to the first-level instruction cache and data cache units to detect and eliminate cases where accesses are being requested within the same page/line, thereby eliminating the tag lookup and minimizing the number of accessed banks. The power saved by eliminating these cycles outweighs the power consumed by the additional control logic, yielding a net savings for the overall design.

In the RAT (register renaming) unit, the register files were partitioned according to the data types (MMX™, Integer, Floating Point) instead of creating a worst-case combined data width, and accesses are performed based on the required data type. Thus the access of redundant data was eliminated up front, reducing the active power consumed by these circuits.

In the Branch Prediction Unit (BPU), which was designed from scratch for the Intel Pentium M processor, many

[™]MMX is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

power optimizations were made: Target Lookups take place only when needed, and Predictions are prevented for Unconditional Branches.

CIRCUIT OPTIMIZATIONS

Circuit techniques have a wide impact on the power budget required for implementing a given functionality. In the Intel Pentium M processor design, many domino-based circuits were replaced with static CMOS circuits eliminating the strong power effect domino circuits have on the input data polarity. Furthermore, by simplifying the timing restrictions on the domino inputs we were able to reach a favorable solution for both power and speed.

All arrays and register files were reviewed to ensure optimal banks partitioning such that the capacitive load that was toggled per access was minimized.

In the 1M second-level cache, leakage current becomes a significant factor hence the required performance (i.e., access latency) was balanced with the power constraints, yielding a design that reduced the second-level cache leakage by a factor of two, by using non-minimal channel lengths for almost the entire array.

To meet the Intel Pentium M processor speed goal we used a high-performance-Low-Vt (LVT) device that has a lower threshold voltage (V_t) at the expense of higher leakage. The need to carefully balance the utilization of LVT devices to improve the speed of circuits with the increase in leakage, led to the setting of an explicit LVT usage POR for every block in the design. The synthesis flow we used assigned the LVT devices at the cell level; therefore, the synthesis blocks had very high initial insertion levels. To reduce this insertion level, a device-level flow was utilized leaving LVT devices only on critical paths. This flow saved more than two-thirds of the original assignments. A similar flow was used on data path designs to identify redundant assignments. When a block exceeded the allocated budget, a review was held to ensure the design was optimal.

I/O Optimizations

The I/O power supply was separated from the core power supply to allow independent optimization of each of the power supply's voltage levels, as is evident from the Intel Pentium M processor data sheet. A special Dynamic On Die Termination (ODT) circuit was added to the output buffer design that enables disconnecting the on die termination when the CPU drives the bus low, thereby reducing by half the power consumed by the I/O. Data inversion support (first implemented in the Pentium® 4 processor) between the Intel® 855PM chipset and the Intel Pentium M processor further reduces the power dissipation, due to the line termination, by minimizing the number of bits driven low on the Gunning Transceiver Logic (GTL) bus. To minimize the time during which the bias current is on in the input buffers of the Processor Side Bus (PSB), a new signal (DPWR#) was added to the interface to indicate to the processor when to operate the input buffers. The overall impact of all these optimizations reduced the active power of the PSB interface by a factor of 2 and the average power by a factor of 10.

Power Management Optimization

The dynamic power component increases linearly with the frequency of the processor clock and by the square of the voltage, hence being able to dynamically adjust the voltage and frequency to the workload has an enormous impact on the average power. The Multi-gear Intel SpeedStep technology that was implemented in the Intel Pentium M processor allows the operating system to provide frequency on demand stepping through pre-defined voltage-frequency pairs, spanning a range of about 8x in the power and almost 3x in performance between the lowest and highest power-performance points [3].

Models generated from study of real applications usage in a mobile environment indicate that high performance is typically needed only for short bursts of time. Average power optimization therefore relies on the ability to switch the operating point frequently and in an efficient manner. To satisfy user interactive demands, the system response time must be kept low. Optimization of the Intel SpeedStep transition process implemented in the Pentium M processor reduced the switching time and allowed efficient use of the Multi-Gear Intel SpeedStep technology mechanism with minimal latency and performance degradation.

Pentium 4 is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Previous generations of Intel CPUs integrated a linear throttle mechanism, stopping the CPU for short periods of time and allowing it to cool. The power impact of controlling both core voltage and frequency is significantly higher than the linear yield of the standard frequency adaptation. The Pentium M processor is Intel's first CPU designed specifically for the mobile market implementing improved Thermal Throttle 2. Thermal Throttle 2 uses frequency and voltage scaling to control maximum and average CPU power. The use of combined voltage and frequency scaling results in a lower performance degradation compared to standard clock-throttling techniques within any given cooling solution.

THERMAL THROTTLING VALIDATION

The cooling intensity of mobile CPUs is adjusted by the system according to the processor temperature. Control is achieved by using variable fan speeds. The fan is activated only at high core temperatures to save the battery and reduce acoustic noise. The response of the cooling system is significantly slower than the processor self-heating process. As a result, the system cannot react on time to prevent rapid temperature increases affecting the device reliability, degrading maximal operating speed, and possibly even causing accidental shut-down due to temporary overheating.

To study the system dynamic thermal response, a special power-stepping test was written, providing repeated high-power pulses with a programmable duration.

To achieve extreme abnormal conditions the device was operated without a heat sink! The test alternates between high and low power segments at a duty cycle of 1 to 10 causing significant power and temperature transitions.

Thermal Throttle 2 provided excellent results, and the successful temperature clamping is shown in [Figure 2](#). It demonstrates that the Intel Pentium M processor can modify working conditions to ensure that the thermal envelope limit is not exceeded.




Figure 2: Bare die thermal throttling experiment demonstrating successful temperature clamping

PRE-SILICON ESTIMATES VALIDATION

A correlation study was performed using silicon-based Infra-Red Emission Microscopy (IREM) measurements of selected power tests. A simulated power density map and the corresponding IREM image are presented in [Figure 3](#) and [Figure 4](#). The color-coding represents the average power and local emission densities. The levels are black (lowest), red, orange, yellow and white (highest). Although the simulation results are presented as average power density per FUB, and the infrared emission is mainly obtained from non-stacked n-type devices, the correlation can clearly be seen.

Multiple measurements under various test modes enable us to measure and confirm the impact of the major power-saving features. The results confirm our pre-silicon estimates and sum up to 35% of the total active power.




Figure 3: Simulated power density




Figure 4: IREM measurement

SUMMARY

We described the dynamic power estimation work and its impact on the design of the Intel Pentium M processor. Low-power tests and various analysis techniques were used for identifying power-saving opportunities and for validating the power design. Most of the identified opportunities were implemented in the design.

The paper also described the architecture enhancements and various circuit, logic, and I/O optimizations that were implemented during the design of the Intel Pentium M processor.

Silicon measurements of the Intel Pentium M processor yielded excellent correlation to our pre-silicon estimates for a wide range of tests that are all within +/- 5% of our estimates. Furthermore, measurements also confirmed that the features implemented for the sole purpose of saving power account for a reduction of over 35% of the active power consumed.

CONCLUSION

The Intel Pentium M processor is a compelling mobile product delivering high performance within a restricted mobile thermal envelope with improved battery life. It provides significant advantages to the Intel Centrino mobile technology-based platform.

The developed architecture, methodologies, and knowledge base pave the way for the design of Intel's next-generation power-efficient mobile products.

ACKNOWLEDGMENTS

We thank the entire Intel Pentium M processor design team for being power aware, for supporting the power-estimation efforts, and for producing a power-conscious design. We also acknowledge the contribution of the Intel Strategic CAD Labs (SCL) for providing and supporting the logic simulation tools that are embedded in the Stochastic Dynamic Power Estimator (SDPE) power-estimation flow.

REFERENCES

- [1] M. D. Aagaard, R. B. Jones, and C.-J. H. Seger, "Formal verification using parametric representations of Boolean constraints," in *Proceedings of the 1999 IEEE Design Automation Conference*, pp. 402-7, 1999.
- [2] J. Yang and C.-J. H. Seger, "Introduction to generalized symbolic trajectory evaluation," in *Proceedings of the 2001 IEEE International Conference on Computer Design: VLSI in Computers and Processors*, pp. 360-5, 2001.

- [3] E. Rotem, A. Naveh, N. Shamir, O. Nathan, Internal publication, 2003.

AUTHORS' BIOGRAPHIES

Dani Genossar has been with Intel since he graduated with a B.Sc. in Electrical Engineering Cum Laude from the Technion, Israel Institute of Technology, in 1988. As a manager within the Mobile Platform Group Design Team he is also responsible for the power-estimation and optimization efforts. His e-mail is dani.genossar@intel.com

Nachum Shamir is the owner of iMPG power-estimation, analysis, and validation flows and leads the silicon power and thermal characterization efforts. He joined Intel in 1999. Nachum received a B.Sc. degree in Electrical Engineering from the Technion, Israel Institute of Technology, in 1988, an M.Sc. degree in Electrical Engineering from Tel Aviv University in 1994, and a Ph.D. degree in Electrical Engineering from the Technion in 1999. Nachum served as adjunct lecturer at the Technion Electrical Engineering faculty from 1998 to 2000. His e-mail is nachum.shamir@intel.com

Copyright © Intel Corporation 2003. This publication was downloaded from <http://developer.intel.com/>.

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

For further information visit:

developer.intel.com/technology/itj/index.htm



Intel® Technology Journal

Intel® Centrino™ Mobile Technology

Antenna Selection in Multicarrier Communication Systems

Antenna Selection in Multicarrier Communication Systems

Yuval Finkelstein, Wireless Networking Group, Haifa, Intel Corporation

Index words: communication, wireless, OFDM, fading channels

ABSTRACT

Imagine a packet of data transmitted by radio frequency through the air. This packet appears at two separate antenna ports of a radio receiver. Even though the same data appear at both antennas, a different path in space is associated with each antenna. Any path in space results in signal power reduction and usually in a power redistribution in the allocated frequency band. Both effects enhance the error rate and reduce the reliability of the decoded data. The receiver usually supports the wireless link only through a single antenna due to cost considerations. This paper focuses on the particular case where the signal at one antenna is strong, but does not efficiently utilize the available frequency band, and at the other antenna the signal is weak, but uniformly distributed over the band. The paper provides a general algorithm for selecting the optimal antenna and explains both quantitatively and qualitatively which antenna should be selected.

INTRODUCTION

Multipath fading is a major limitation to high-rate data transmission in wireless communication systems. Fading is a result of destructive interference between replicas of the signal arriving through different paths. Another aspect of the same effect is an intersymbol interference. This occurs when the replicas merge again at the antenna and successive data symbols are mixed together, forming a mismatched symbol.

Multicarrier communication methods overcome intersymbol interference by subdividing the allocated bandwidth into a few frequency sub-bands. At each carrier (sub-band), the data is transmitted using long symbol durations compared to the time delay between replicas. Thus, the impact of intersymbol interference is reduced. Transmitting the data simultaneously through all carriers results in an overall high rate of data transmission. However, destructive interference still distorts the data. Those carriers that are subjected to destructive interference have a low Signal-to-Noise Ratio (SNR).

Low SNR at defective carriers is partly overcome by applying spatial diversity (antenna diversity). Receiving the data simultaneously through multiple antennas increases the detection reliability. However, full antenna diversity is an expensive solution. A more cost-saving solution is to use semi-diversity. This is accomplished by selecting only a single antenna out of a given set. This antenna should offer the most reliable detection.

Selecting antenna is usually accomplished by considering total antenna power. This criterion is unsatisfactory because power is not the only factor that determines the performance of the receiver. In fact, the distribution of the power among carriers is critical as well. For example, take an antenna that has received a giant pulse of power. Most of it emerges in a single carrier. Selecting this antenna while the rest of the carriers suffer low SNR exhibits a high error rate.

This paper provides a simple algorithm for choosing the best antenna from among several. The algorithm improves the reliability of the data recovery process. The central pillar of this algorithm is an information-capacity-like parameter for each of the carriers. Using this parameter, instead of estimating power alone, yields more useful information about the contribution of each carrier. The performance of the receiver, while equipped with each antenna, is predicted by combining the appropriate contributions over all the carriers.

The antenna selection algorithm described in this paper was implemented inside Intel's chipset that handles the IEEE 802.11a standard. The chipset provides wireless connectivity to mobile PCs and also serves Intel® Centrino™ mobile technology.

The paper demonstrates the application of the algorithm by referring to a simplified system consisting of only two carriers. Full treatment of the IEEE standard, which also includes error-correction coding, is beyond the scope of this paper. However, the described principles are very

Intel Centrino is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

relevant to practical systems. An effort was made to make the contents clear to technical staff who are not directly involved with communication. Thus, terms such as constellation points, Signal-to-Noise ratio, equalization and gain control are demonstrated explicitly in this paper. After mathematically constituting the relevant glossary of terms, the suggested antenna-selection algorithm is introduced. Finally, the algorithm is proved to be successful in predicting the performance of a receiver equipped with an arbitrary characteristics antenna. This is established by numerical simulation of decoding noisy data.

RELEVANT GUIDELINES TO ORTHOGONAL FREQUENCY DIVISION MULTIPLEXING

As was previously mentioned, the main idea behind a multicarrier communication system is the division of a given frequency band into smaller frequency sub-bands. Orthogonal Frequency Division Multiplexing (OFDM) is a particular case of a multicarrier communication system, in which the frequency sub-bands overlap. This overlapping is not harmful due to the fact that the carriers are mathematically orthogonal. A good reason for using OFDM is its inherent immunity against narrowband interference. If a specific sub-band is severely disturbed due to multipath fading or interference, error-correction coding can still overcome the disturbance, thus preventing the entire link from failing, as is the case with a single carrier. Another reason for using OFDM concerns the complexity of the equalizer. Assuming that each sub-band is sufficiently narrow, the effect of the channel is taken into account by assigning only a single complex number to each carrier. In comparison to a single-carrier system, this simplifies the implementation of the equalizer to a great extent.

More extensive explanations about the principles of OFDM are documented elsewhere [1]. However, a few more guidelines are presented here to help clarify the information in the remaining part of this paper. Following IEEE standard 802.11a, all carriers that constitute a single data frame are subjected to the same modulation scheme: QAM constellation with a given order. The data is encoded by assigning a specific binary combination to each of the constellation points. This process is demonstrated in the next section. Thus, a single constellation point is determined for each frequency sub-band. Synthesizing the sequence of constellation points into a time domain signal is accomplished by Inverse Fourier Transformation. The inverse process takes place at the receiver. The time domain signal is transformed back to the frequency domain by means of Discrete Fourier Transform. Identifying which constellation point

was assigned by the transmitter to each frequency band constitutes the demodulation process. Recognizing the correct point despite the noise and the channel fading is vital for successfully receiving the data. Finally the data is extracted by decoding the constellation points back into binary.

DEFINITION OF A SIMPLIFIED SYSTEM

In order to clarify our strategy, let us assume a simplified multicarrier communication system, which uses only two carriers. That is, the allocated frequency band is only divided into two sub-bands. At each sub-band information is modulated by the well-known Quadrature Phase-Shift Keying (QPSK) method. Figure 1 presents the constellation plane for QPSK and our preferred encoding scheme.




Figure 1: Bits encoding scheme for QPSK constellation

Suppose that the sequence to be delivered is “1110.” For simplicity, we assume that no error-correction coding is applied. Encoding this sequence results in two complex numbers: $1+j$ and $1-j$, which comprise the baseband representation of this message. The transmitter combines these two complex numbers into a waveform in the time domain by using the Inverse Fourier Transformation. This waveform finally modulates the amplitude and phase of the RF carrier that is transmitted through the air.

Demodulation and Noise

At the receiver the reverse process is taking place, which results in measured points in the constellation plane. However, the constellation points measured by the receiver are corrupted by noise. This noise cannot completely be avoided, no matter what the quality of the receiver is. Once the message is transmitted the receiver measures two constellation points, (I_1, Q_1) and (I_2, Q_2) , in carriers “1” and “2,” respectively. Although the exact values $1+j$ and $1-j$ were sent by the transmitter, this would never be the case with the measured values, due to the noise. Since no error-correction coding was applied, demodulation is carried out just by assigning each

measured point to its nearest constellation point. That is, demodulation is performed by identifying the quarter in which the points (I_1, Q_1) and (I_2, Q_2) appear. If the noise power is more intensive than that of the signal power, demodulation becomes a random process. In this case successful decoding would happen with a high probability of 0.5. This of course is not satisfactory for constituting a communication link.

We then adopted an approach that makes use of an Ensemble average. This enables us to treat the noise in a statistical manner. Figure 2 represents measurements of (I_1, Q_1, I_2, Q_2) for an ensemble of receivers assuming additive white Gaussian noise.




Figure 2: Constellation planes of carriers “1” and “2”

We can see that all measurements are concentrated around the constellation point $1+j$ and $1-j$ according to the message that was sent. However, the noise shifts the measured values to different locations in the constellation plane. The figure makes visible the fact that each measured point can be described by the sum of two components: the signal component (either $1+j$ or $1-j$) and the noise component (n_1 and n_2 for carriers “1” and “2,” respectively). This can be mathematically expressed as follows:

$$P_1 = I_1 + jQ_1 = (1 + j) + n_1$$

$$P_2 = I_2 + jQ_2 = (1 - j) + n_2$$

The Ensemble average of the squared-magnitude of the noise components is equal for both carriers, as is evident from Figure 2. It is designated by σ^2 :

$$\langle |n_1|^2 \rangle = \langle |n_2|^2 \rangle = \sigma^2$$

It is very important to emphasize here that we do not restrict the statistical nature of the noise to any particular kind of noise. Although we used white Gaussian noise in our simulations, this kind of noise does not always have to be used. The mean of the noise, however, should be zero. The SNR is now defined as follows:

$$SNR(\sigma)_{[dB]} = 10 \log_{10} \left(\frac{\langle |P_1|^2 \rangle - \langle |n_1|^2 \rangle + \langle |P_2|^2 \rangle - \langle |n_2|^2 \rangle}{\langle |n_1|^2 \rangle + \langle |n_2|^2 \rangle} \right) =$$

$$= 10 \log_{10} \left(\frac{2}{\sigma^2} \right)$$

The value of σ for the Gaussian distribution in Figure 2 is 0.5, which results in an SNR of 9 dB.

CHANNEL EFFECT AND EQUALIZATION

Channel Effect

So far only the additive noise was considered as a distraction to demodulation and data extraction. However, any actual signal is disturbed to some extent by channel distortion. Replicas of the signal coming from walls and equipment interfere with the antenna. The result of this interference is the so-called “channel” effect. The total power that appears at the antenna port of the receiver is always smaller than the transmitted power, due to channel losses. However, the channel also distorts the original signal by changing its spectral content. In a multicarrier system, decomposing the received signal back to its spectral component is actually equivalent to measuring constellation points for each of the frequency sub-bands that constitute the system. Under the effect of the channel, these constellation points do not appear at their original location as in the transmitter. Instead, both their magnitude and phase are changed. That is, the channel moves each constellation point to an arbitrary new location, even if noise is completely absent.

Equalization

The channel distortion can be partly overcome by equalization. The data packets are usually preceded by a known training sequence. Thus, the channel distortion can be analyzed and can be taken into account throughout the process of demodulation. Assuming that the frequency sub-bands are narrow enough, the channel influence for each carrier can be expressed with a single complex number. The phase and magnitude of each number describe the transformation that the constellation points undergo in the relevant carrier. These complex numbers are called “channel coefficients.” The equalizer

coefficients are no more than the reciprocal of the channel coefficients. Multiplying the measured constellation points by the appropriate equalizer coefficients moves them back to their normal place. If the training sequence is long enough and the channel has static characteristics, the estimation error of the channel coefficients can be made as small as it needs to be. Therefore we neglect this noise component and assume ideal equalizer coefficients for the rest of this paper.

Nevertheless, equalization does not completely cancel the impairments of the channel. This is due to the fact that equalization cannot compensate for the intrinsically low SNR in carriers that are suppressed by the channel; in other words, carriers with a magnitude of channel coefficient that is much lower than one. The interplay between the total power of both carriers and the magnitude of their channel coefficients is most relevant to the topic of this paper, that is, antenna selection.

Gain Control

We continue with our double-carriers simplified system. This time we add the influence of a channel. We designate the channel coefficients by C_1 and C_2 for carriers "1" and "2," respectively. Thus, the constellation points measured at the receiver are redefined by

$$P_1 = I_1 + jQ_1 = (1 + j) \times C_1 + n_1$$

$$P_2 = I_2 + jQ_2 = (1 - j) \times C_2 + n_2$$

Next, we present the concept of Gain Control. While this function is accomplished both analogically and digitally in a rather complicated way in real systems, here we are just concerned with the essence of this feature. Our ideal gain controller verifies that

$$\langle |P_1|^2 \rangle + \langle |P_2|^2 \rangle = 4$$

The Ensemble average that appears in the equation above verifies that all the receivers in our ensemble have exactly the same gain. Substituting the explicit expression for the points P_1 and P_2 and recalling that the variance of the noise was defined as σ^2 we get

$$|C_1|^2 + |C_2|^2 + \sigma^2 = 2$$

The magnitudes of C_1 and C_2 can take values either higher or lower than one. An increase in the magnitude of one coefficient occurs at the expense of a decrease in the magnitude of the other. The reader should note that channel coefficients with a magnitude higher than one do not imply increased power in an absolute manner. The magnitudes of the channel coefficients only indicate how the measured power is distributed among the frequency

sub-bands and how the relative strength of the signal compares to the noise.

The expression for the SNR in the gain-controlled receiver is modified, since the total signal power is not 4 ($|1+j|^2 + |1-j|^2 = 4$) any more. Instead, it is the Ensemble average of the total power (both signal and noise) that equals 4. Thus, the total SNR, which refers to both carriers together, is

$$SNR_{total}(\sigma)_{[dB]} = 10 \log_{10} \left(\frac{2 - \sigma^2}{\sigma^2} \right)$$

Equalization and Modified Signal-to-Noise Ratios

After equalization, the original constellation points are restored with modified noise components:

$$P_1^{Eq} = (1 + j) + \frac{n_1}{C_1}$$

$$P_2^{Eq} = (1 - j) + \frac{n_2}{C_2}$$

The Ensemble averages of the squared-magnitudes of the modified noise components are

$$(\sigma_1)^2 = \left\langle \left| \frac{n_1}{C_1} \right|^2 \right\rangle = \frac{\sigma^2}{|C_1|^2}, \quad (\sigma_2)^2 = \left\langle \left| \frac{n_2}{C_2} \right|^2 \right\rangle = \frac{\sigma^2}{|C_2|^2}$$

Accordingly,

$$SNR_1[dB] = 10 \log_{10} \left(\frac{2}{|\sigma_1|^2} \right)$$

$$SNR_2[dB] = 10 \log_{10} \left(\frac{2}{|\sigma_2|^2} \right)$$

ANTENNA DIVERSITY

Suppose the receiver is equipped with two antennas: antenna *A* and antenna *B*. During signal detection the receiver switches back and forth between these two antennas. The gain controller sets the gain such that the condition $\langle |P_1|^2 \rangle + \langle |P_2|^2 \rangle = 4$ is fulfilled for the highly energized antenna. Thus, saturation of the receiver's amplifier is avoided no matter which antenna is attached. Once the gain is set, the receiver estimates two sets of channel coefficients. Each set refers to a different antenna and consists of two channel coefficients, a single coefficient for each carrier. As soon as estimations are completed, the receiver has to select either antenna *A* or antenna *B*, in order to receive the rest of the data. Which one should it select?

Power Considerations

Suppose that antenna *A* is the one with the higher total power. The gain controller sets the gain such that

$$|C_1^A|^2 + |C_2^A|^2 + \sigma^2 = 2$$

Indexes A and B are attached to the channel coefficients of antennas A and B , respectively. Assuming that the signal power in antenna B is reduced by a factor of X we get

$$X \left[|C_1^B|^2 + |C_2^B|^2 \right] + \sigma^2 = 2$$

Note that at this antenna, it is only the signal power that is reduced by X and not the noise power. Thus, the noise level σ^2 is the same no matter which antenna is selected. This statement is valid since the amplifier gain is set just before the beginning of the antenna selection process.

Symmetry and Asymmetry

As will be evident soon, the performance of the receiver is optimal when the signal power is homogeneously distributed between the two frequency sub-bands. We define an antenna whose channel coefficients are identical (the same for both carriers) as a symmetric antenna. In contrast, an asymmetric antenna is one with different channel coefficients. If the higher power antenna is also a symmetric one, then selecting the best antenna becomes a trivial task. However, the decision becomes more complicated if the high-power antenna is asymmetric, and the low-power antenna is symmetric. We focus our discussion on the latter. Since we defined antenna A to be the one with the higher power, we extend its definition and declare it asymmetric. In order to take into account the asymmetry property in a quantitative manner, we define the asymmetry parameter S as follows:

$$S = \frac{|C_2|}{|C_1|}$$

Thus, $0 \leq S < 1$ for antenna A whereas $S = 1$ for antenna B . When selecting the antenna only the parameters X and S should be taken into account; these describe power reduction in B and the asymmetry of the carriers in A , respectively. The values of these two parameters are accessible to the receiver since they can be ascertained once the channel coefficients are estimated.

Simulation of Bits Error Rate

In order to realize the effect parameters X and S have on receiver performance, we simulated the entire simplified system following the structure we described so far. We emphasize again that we do not restrict the statistical nature of the noise to any particular kind. Although we used white Gaussian noise in the simulations, this should not always be the case. If the statistical nature of the noise is known in advance, the Bits Error Rate (BER) can be predicted analytically. However, because we do not assume any specific statistical distribution we use a numerical attitude instead of an analytical one. Furthermore, in practical communication systems, interleaving and error-correction coding are probably

applied, making the BER analytical prediction too complicated. Thus, the numerical exploration of the simplified communication system, along with an empirical construction of the antenna selection algorithm, actually provides a complete methodology that is applicable to practical systems as well.




Figure 3: BER results for different values of the asymmetry-parameter, S

The plot at the top of Figure 3 presents BER results for our double-carriers system, which relies on QPSK modulation. Demodulation was carried out, as was explained earlier, by identifying to which quarter the measured constellation points belong.

The horizontal axis of the figure shows SNR_{total} in a logarithmic scale. Its definition is repeated here again for the convenience of the reader:

$$SNR_{total}(\sigma)_{[dB]} = 10 \log_{10} \left(\frac{2 - \sigma^2}{\sigma^2} \right)$$

Each colored line in the figure describes BER results for different values of the asymmetry-parameter, S . The furthest right is the case of $S=0.1$ with the worst performance, as expected. The furthest left is the case of $S=1$ with the best performance. In between, S displays intermediate values.

The empty circles at the figure denote the intersection of the curved lines with an invisible horizontal line, which

crosses the vertical axis at the BER value of 10^{-3} . These circles are plotted again at the lower part of the figure with a modified horizontal axis. The modification incorporates a shift of the zero to be just below the furthest left circle. Thus, the value that fits each circle in the new scale is just the distance from the furthest left circle. The reader should realize that the furthest left circle actually represents the symmetric antenna B . Therefore, the rest of the circles specify the amount of excess power that should be added to the asymmetric antenna A , in order to achieve the same BER result of 10^{-3} . This excess power depends on the value of the asymmetry-parameter, S , as the figure shows.

ANTENNA SELECTION

In constituting the antenna selection algorithm we follow these logical steps:

- Understand qualitatively why excess power is required for compensating for asymmetry.
- Based on this understanding, suggest an algorithm that predicts receiver performance for given values of X and S .
- Validate the algorithm and tune its free parameters. This is done by reconstructing the dependence of the excess power on the symmetry parameter S , as depicted in Figure 3.
- Verify proper antenna selection for any combination of asymmetry parameters S_A , S_B , and power difference X .

Therefore, why is excess power required for compensating for asymmetry? Increasing asymmetry while keeping total power constant causes Signal-to-Noise Ratio (SNR) degradation in one carrier and SNR improvement in the other. In other words, decreasing S gradually from 1 to lower values increases the spread of equalized points in one constellation and decreases the spread in the other. Thus, more and more points move to a wrong quarter in the constellation plane of one of the carriers while at the same time spreading is reduced in the other carrier. However, if spreading is initially small and a good communication link can be established for $S=1$, reducing the spreading further has no benefit at all. In contrast, the SNR degradation at one of the carriers will finally shift points to the wrong quarter and the communication link will fail. Applying more and more power while asymmetry increases keeps all constellation points inside the quarter they really belong to.

Based on this explanation we propose a grade for each carrier, based on its effective SNR. The grade is calculated based on the fact that the benefit from high signal power is saturated beyond a certain SNR value. In

the same way, decreasing the grade for defective carriers should be restricted if SNR goes below another value. The antenna is therefore selected by following these steps:

- Grade each carrier according to its SNR.
- Add up all grades to get a total score.
- Select the antenna with the highest score.

“Kapasity”

The missing building block for the algorithm above is the function by which grades are calculated, versus the SNR. We already indicated what should be the asymptotic behavior of that function for high and low values of SNR. This kind of asymptotic behavior is similar to the information-capacity of a channel [2]. Indeed, estimating channel capacity can be a useful tool for predicting receiver performance. However, we are not going to make any explicit use of information theory here. The function we are looking for is an empirical one, with free parameters for tuning the receiver performance. Despite all that, we named the total antenna score “Kapasity,” after the well-known phrase.

We define the Kapasity function $f(\sigma)$ as

$$f(\sigma) = \frac{1}{1 + \exp\left(\frac{\sigma - d}{q}\right)}$$

Note that f is defined in terms of σ instead of SNR, which should make the equations that follow easier to understand. However, the variance of the noise for each carrier (σ^2) and the SNR of each carrier are related, as was explained previously. It is easy to verify that f has the required asymptotic behavior, as σ approaches either zero (i.e., high SNR) or infinity (i.e., low SNR). The explicit form of the function f can be chosen in many other ways, if the correct asymptotic behavior is kept. For hardware implementation, any piecewise linear approximation of the function f can fit as well. The parameters d and q are free parameters that should be tuned in order to optimize the antenna selection process. Factors that affect the optimized values of d and q are the type of error-correction coding and the available hardware resources for Kapasity calculations. However, it is easy to guess what should be the value of the parameter d . Notice that $f=0.5$ when $\sigma=d$. At this point the gradient of f is maximal, and minor changes of σ result in major changes in the Kapasity. This tendency also exists with the demodulation error rate, as the noise standard deviation approaches half of the distance between two adjacent constellation points. Thus, the likely assumption is that the optimal value of d for our QPSK constellations should be around 1. We will see soon that this is really the case. The value of q determines the slope of f around $\sigma=d$. The

higher the value of q the lower the slope of f is at that point.

VALIDATION OF THE ALGORITHM

Figure 3 describes the excess power required for an asymmetric antenna compared to a symmetric one, in order to keep both antennas performing at the same level. This illustrates one scenario for antenna selection where $S_B=1$ and the Kapsity is equal for both antennas. We follow this trend and speculate what should be X for each value of S_A . Formulating this mathematically in terms of Kapsity we get

$$f(\sigma_1^A) + f(\sigma_2^A) = 2f(\sigma^B)$$

The superscripts A and B indicate which antenna σ belongs to, and the subscripts 1 and 2 relate to carriers 1 and 2 , respectively. Because we assume that antenna B is symmetric, both its carriers have the same value of σ . Therefore, from now on we omit the subscripts for antenna B .

As soon as we express σ_1^A, σ_2^A and σ^B in terms of X and S , we can insert them into the Kapsity equation above. We go back to the mathematical model of the simplified system. Recall that after equalization the noise component in each carrier can be characterized by an effective variance, which depends on the relevant channel coefficients

$$(*) \quad (\sigma_1^A)^2 = \frac{\sigma^2}{|C_1^A|^2}, \quad (\sigma_2^A)^2 = \frac{\sigma^2}{|C_2^A|^2}, \quad (\sigma^B)^2 = \frac{\sigma^2}{|C^B|^2}$$

Extracting C_1^A and C_2^A (in terms of S and σ) from the equations that define gain and asymmetry,

$$|C_1^A|^2 + |C_2^A|^2 + \sigma^2 = 2, \quad S = \frac{|C_2^A|^2}{|C_1^A|^2}$$

and inserting them into (*), we get

$$\sigma_1^A = \sqrt{1+S^2} \sqrt{\frac{\sigma^2}{2-\sigma^2}}$$

$$\sigma_2^A = \sqrt{\frac{1+S^2}{S^2}} \sqrt{\frac{\sigma^2}{2-\sigma^2}}$$

Extracting C^B (in terms of X and σ) from the equation that defines X ,

$$|C_1^B|^2 + |C_2^B|^2 = 2|C^B|^2 = \frac{2-\sigma^2}{X}$$

and inserting it into (*), we get

$$\sigma^B = \sqrt{2X} \sqrt{\frac{\sigma^2}{2-\sigma^2}}$$

From Figure 3 we derive σ and S for each circle and insert both of them into σ_1^A and σ_2^A above. Thus, the left-hand side of the equal-Kapsity equation is calculated for each empty circle in Figure 3. All that is left is to extract X , which only appears at the right-hand side of the equation.




Figure 4: Excess Power calculated by using the Kapsity function, versus the asymmetry-parameter, S

The empty circles at the top plot in Figure 4 describe the Excess Power versus S as calculated by the Kapsity model. Red and blue circles illustrate the results while using $(d=1.2, q=0.04)$ and $(d=1.2, q=0.16)$, respectively. The shape that the Kapsity function takes using each set of parameters is depicted at the lower part of Figure 4, keeping the same convention of colors. The red continuous line in Figure 4 represents the Excess Power versus S derived from the open circles in Figure 3. It is clear that using the Kapsity model with $(d=1.2, q=0.04)$ accurately predicts the results of the numerical experiment. We consider this to be proof of the validity of the antenna selection algorithm. It is also evident that the initial guess of $d \approx 1$ is very close to the actual value. The other set of d and q , which isn't a good fit in Figure 4, is discussed in the last section of this paper.

FINAL VALIDATION OF THE ALGORITHM

The algorithm is formulated following these steps:

- Calculate the total Kapsity for each of the relevant antennas (might be more than two) by summing up Kapsity contributions over all carriers (probably much more than two):

$$Kapsity = \sum_n f(\sigma_n)$$

- Chose the antenna with the highest Kapsity.

Although the excellent fit between the red circles and the continuous line in Figure 4 validates the algorithm, we would like to demonstrate its applicability in a broader sense. In the previous section, we focused on the case where $S_B=1$. In this section we consider the general case in which both S_B and S_A have arbitrary values between 0 and 1.

Let us focus again on the results of the numerical experiment presented in Figure 3. We depicted the performance of our double-carriers system in terms of BER versus $SNR_{total}(\sigma)$ and the asymmetry parameter S . This time we consider the entire set of continuous lines in Figure 3 (instead of only considering the set of open circles). Recall that each line shows how the BER is changed with $SNR_{total}(\sigma)$, for each definite value of S . Thus, for final validation of the algorithm, we should reconstruct the complete BER diagram of Figure 3 in terms of Kapsity. Figure 5 presents exactly that.

The left-hand plot shows the value of $[2-Kapsity]$ versus $SNR(\sigma)$ and S using $(d=1.2, q=0.04)$. The right-hand plot shows $[2-Kapsity]$ versus $SNR(\sigma)$ and S , using $(d=1.2,$

$q=0.16)$. Replacing the Kapsity with its complement of 2 is just for producing a monotonic decreasing function instead of an increasing one. This makes easier the comparison between Kapsity plots and BER plots. Each colored line in Figure 3 is transformed into a line with the same color in Figure 5. The open circles show how points with a definite BER value in Figure 3 appear in the Kapsity diagram. It is clear that an absolutely horizontal line exists, which separates circles of different colors. Thus, points with the same BER in Figure 3 become points with the same Kapsity in Figure 5. This finally completes the validation of the proposed antenna selection algorithm.

It is evident from the left-hand side plot of Figure 5 that hardware implementation of Kapsity calculations demands an extremely large number of bits. Otherwise the antenna-selector will not be able to choose between the antennas. However, the parameters $d=1.2, q=0.16$ enable binary representation of the Kapsity using less bits, as shown in the right-hand side plot in Figure 5. This results, however, in a reduction in the efficiency of antenna selection. It is evident from the fact that the separating line between same color circles is not a straight horizontal line anymore. Thus, there is a tradeoff between antenna selection efficiency and hardware cost.




Figure 5: Kapsity versus $SNR(\sigma)$ and S

CONCLUSION

We reviewed the basic principles of multicarrier communication systems and explained how channels can modify power distribution among carriers and how equalization might affect the noise component in each carrier. We described an antenna selection algorithm developed in Intel. Although the algorithm was applied to a simplified communication system, the described methodology was found to be applicable for practical systems as well. Hence it was implemented inside Intel's Orthogonal Frequency Division Multiplexing (OFDM) modem that provides wireless connectivity for mobile PCs.

REFERENCES

- [1] R. Van Nee and R. Prasad, *OFDM For Wireless Multimedia Communication*, Artech House, Boston-London, 2000.
- [2] J. G. Proakis, *Digital Communications*, McGraw-Hill, 4th ed., Boston, 2001.

AUTHOR'S BIOGRAPHY

Yuval Finkelstein joined the wireless Digital Signal Processing team in the Intel Development Center in Haifa in 2000. He has prior industrial experience in the areas of electro optics and solid-state physics. He received his Ph.D. degree in physics from the Israel Institute of Technology, the Technion, in 1997. His e-mail is yuval.finkelstein@intel.com

Copyright © Intel Corporation 2003. This publication was downloaded from <http://developer.intel.com/>.

Legal notices at <http://developer.intel.com/sites/corporate/tradmarx.htm>.

For further information visit:

developer.intel.com/technology/itj/index.htm