

Introduction to Intel® Core™ Duo Processor Architecture

Simcha Gochman, Mobility Group, Intel Corporation
Avi Mendelson, Mobility Group, Intel Corporation
Alon Naveh, Mobility Group, Intel Corporation
Efraim Rotem, Mobility Group, Intel Corporation

Index words: Intel Core Duo, low power, CMP, multi-threading

ABSTRACT

The Intel® Core™ Duo processor is a new member of the Intel® mobile processor product line. It is the first Intel® mobile microarchitecture that uses CMP (Core Multi-Processor; i.e., multi cores on die) technology. Targeted to the market of general-purpose mobile systems, the Intel Core Duo core was built to achieve high performance, while consuming low power and fitting into different thermal envelopes.

In order to achieve the required performance, a CMP-based microarchitecture was designed to achieve power-efficient architecture, each performance improvement was evaluated against the power cost, and only the power-efficient performance features were implemented.

On top of that, special hardware mechanisms were added to better control the static and the dynamic power consumption. As a result, the Intel Core Duo processor provides higher performance in the same form factors without needing to increase the cooling capability.

INTRODUCTION

The Intel Core Duo processor is a new member of the Intel mobile processor product line. It is the first Intel mobile microarchitecture that uses CMP (multi cores on die) technology. Building a general-purpose mobile core is a challenging task since, on the one hand, the system needs to maintain the highest level of performance, while on the other hand, the system must fit into different thermal envelopes, as illustrated by Figure 1, and improve power efficiency.

Intel Core Duo is based on Pentium M processor 755/745 core microarchitecture with few performance improvements at the level of each single core. The major performance boost is achieved from the integration of dual cores on the die (CMP architecture). This agrees with our assessment that continuing to improve single

thread performance is rather costly in terms of power and may achieve diminishing returns in terms of efficiency, if major microarchitecture enhancements are not made. The big potential for improved performance is through exploring parallelism between threads. However, the CMP architecture presents many challenges for power and thermal control to still fit into the mobility constraints.



Figure 1: Products using different thermal envelopes*

In this paper we present the new Intel Core Duo microarchitecture and show how the need to target power-efficient general-purpose processors has affected many of our decisions. We provide a general overview of the different ingredients of the Intel Core Duo system, while the other papers in this issue of the *Intel Technology Journal* focus on more specific aspects of the system such as the CMP microarchitecture and the power and thermal control methods.




Figure 2: Intel Core Duo processor floor plan

As Figure 2 shows, Intel Core Duo technology is based on two enhanced Pentium M cores that were integrated and use a shared L2 cache. The way we integrated the dual core in the system had a major impact on our design and implementation process. In order to meet the performance and power targets we aimed to do the following:

- Keep the performance similar to or better than that of single thread performance processors in the previous generation of the Pentium M family (that use the same-size L2 cache).
- Significantly improve the performance for multithreaded and multi-processes software environments.
- Keep the average power consumption of the dual core the same as previous generations of mobile processors (that use a single core).
- Ensure that this processor fits in all the different thermal envelopes the processor is targeted to.

In this paper we provide a high-level description of the main Intel Core Duo features and discuss how each feature fits into the targets of the various projects.

THE IMPROVED PENTIUM M PROCESSOR-BASED CORES

The core of the Intel Core Duo processor-based technology is an enhanced Pentium M processor 755/745¹ core converted to 65nm process technology.

¹ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See

http://www.intel.com/products/processor_number for details.

The main focus of the core enhancements was to do the following:

- Support virtualization (Virtualization Technology²) [3].
- Support the new Streaming SIMD Extension (SSE3) [4].
- Address performance inefficiencies mainly in the handling of SSE/SSE2, FP (x87) and some long latency integer instructions.

Intel Core Duo Processor-based Technology Core Performance Improvements

Intel Core Duo processor-based technology introduces performance improvements in the following areas:

- Streaming SIMD Extensions (SSE/2/3)
- Floating Point (x87)
- Integer

The main difficulty with SSE implementation in Pentium M is caused by the fact that SSE/2/3 is a 128-bit wide microarchitecture while the Pentium M execution core is 64-bits wide (in order to meet power and energy constraints). Making the machine twice as wide may produce more heat and so will have a significant impact on the Thermal Design Point (TDP) of the system as well as some impact on battery life. Since the Pentium M was primarily designed for mobility we preferred to make it relatively narrow and cope with the SSE performance issues. The by-product of this tradeoff is that each SSE vector operation is “broken” into 64-bit wide micro-operation (uOp) pairs. Such instructions suffer from several performance bottlenecks in the Pentium M pipeline, mainly in the Front End (FE) of the pipeline. For example, the Instruction Decoder in the Pentium M processor can potentially handle three instructions per cycle but only the first decoder in a row is capable of handling complex instructions. The other two decoders are limited to single uOp instructions only. This works fine in most cases since the most frequent instructions are single uOp. However, this is not the case with SSE instructions: only scalar SSE operations are single uOps while the vector operations are typically 2-4 uOps. This results in several potential bottlenecks in the

² Intel® Virtualization Technology requires a computer system with a processor, chipset, BIOS, virtual machine monitor (VMM) and applications enabled for virtualization technology. Functionality, performance or other virtualization technology benefits will vary depending on hardware and software configurations. Virtualization technology-enabled BIOS and VMM applications are currently in development.

FE: the Instruction Decoder in the Pentium M can only handle one SSE vector operation per cycle, causing starvation in the rest of the machine. This bottleneck was addressed in the Intel Core Duo core: a new mechanism was introduced that allows lamination of pairs of similar uOps. This mechanism along with enhanced uOp fusion allows handling of the SSE/2/3 vector operation by a single laminated uOp. The instruction decoders were modified to handle three such instructions per cycle, increasing significantly the decode bandwidth of SSE vector operations. The laminated uOps streaming down the pipe are at a certain point un-laminated, reproducing again the 64-bit wide uOp pairs to feed the machine. These changes not only improve performance of vector operations but also save some energy since the FE, no more a bottleneck, can be clock gated whenever its uOp buffer is filled beyond a certain watermark.

Another bottleneck that was discovered was the handling of the floating point (FP) Control Word (CW). The FP CW is part of the x87 state and was usually viewed as “constant”; namely it is loaded once at the beginning and stays constant throughout the program. This is indeed the way the FP CW is used by most of the programs. However there are some FP applications that manipulate the “rounding control” which is located in this register: the default rounding mode is “rounding to nearest even” but before converting results to fixed point, some applications change the round control to “chop” (this is the rule with C programs for example). Such behavior was treated rather inefficiently by the Pentium M core: each manipulation of the FP CW was effectively stalling the pipeline until its completion. The Intel Core Duo core introduced a new renaming mechanism for the FP CW so that four different versions of this register can coexist on the fly without stalling the machine.

Intel Core Duo also improved the latency of some long latency integer operations such as Integer Divide (IDIV). Although these instructions are not very frequent, because of their extremely long latencies, their accumulative affect on integer benchmark scores have shown to be very significant. The basic Divide algorithm has remained unchanged; however, Intel Core Duo Divide logic exploits opportunities for “early exit.” The Divide logic calculates in advance the number of iterations that are required to accomplish the operation. This is indeed data dependent; however, it is often significantly smaller relative to the maximal number of iterations. Once the required number of iterations is accomplished the divider wraps up the results. This does not impact the maximal Integer Divide latency; however, on average it is much faster.

Another enhancement that benefits different kinds of applications is the introduction of a new mechanism of

H/W prefetcher. This mechanism identifies streaming loads at a very early stage in the machine and speculatively predicts the future incarnation of these loads. These speculative requests are looked up in the shared L2 cache and if miss, they’re speculatively prefetched from the external memory. This mechanism is dynamically deactivated whenever there are many demand requests pending (a watermark mechanism). The benefit of this change is an average reduction in load latency.

The performance implication of these enhancements on single-threaded (ST) applications as well as on multithreaded (MT) applications are discussed in [1]

CMP-GENERAL STRUCTURE

Intel Core Duo processor-based technology implements shared cache-based CMP microarchitecture in order to maximize the performance of both ST and MT applications (assuming the same L2 cache size). Figure 3 describes the general structure of our implementation. The figure shows the following:

- Each core is assumed to have an independent APIC unit to be presented to the OS as a “separate logical processor.”
- From an external point of view the system behaves like a Dual Processor (DP) system.
- From the software point of view, it is fully compatible with Intel Pentium 4 processors with Hyper-Threading (HT) Technology³ [6], and DP-based systems. However, special optimizations could be applied to improve the performance of the share-based cache organization.
- Each core has an independent thermal control unit (discussed later in this paper and also covered in [2]).
- The system combines per-core power state together with package-level power state.

The paper *CMP Implementation in Intel Core Duo Systems* [1] extends the discussion on the CMP implementation and compares its performance with other configurations such as the use of split cache architecture. The results shown there indicate that the new proposed

³ Hyper-Threading Technology requires a computer system with an Intel® Pentium® 4 processor supporting HT Technology and a HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See www.intel.com/products/ht/Hyperthreading_more.htm for additional information.

microarchitecture maximizes the performance benefits of both ST and MT execution at a given cache size. The enhancements we implemented in each of the cores allow us to improve both the ST performance (in specific cases) as well as the MT execution. It also allows us to improve the power and the thermal control of the system, and to achieve similar average power consumption, as was the case in the single-core Pentium M processor.

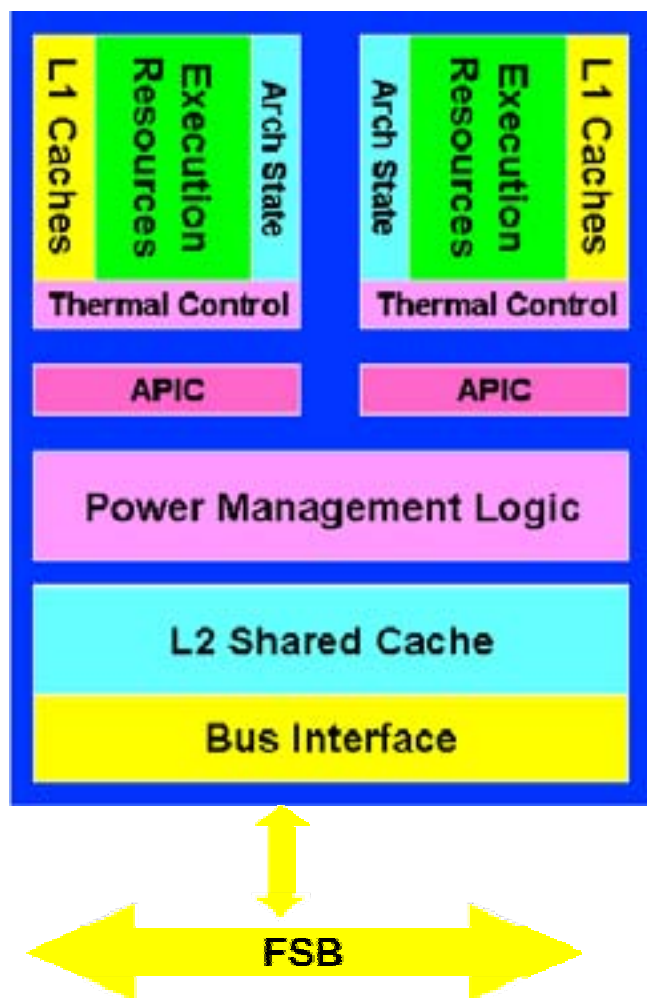


Figure 3: The general structure of the Intel Core Duo implementation

POWER CONTROL

Extending the battery life, while improving the performance, was one of the main goals in designing the Intel Core Duo processor. Battery life is affected by dynamic power, caused when the processor is active, and by static power, which is the power wasted when a unit or the entire processor is not active. Intel Core Duo microarchitecture saves both types of power.

Figure 4 describes the general process we followed in order to reduce the power during the development cycle

of the Intel Core Duo processor. As can be seen, the average power consumption was reduced by handling the problem at all different levels of the design, starting with adjusting the process technology through all the design stages of production.


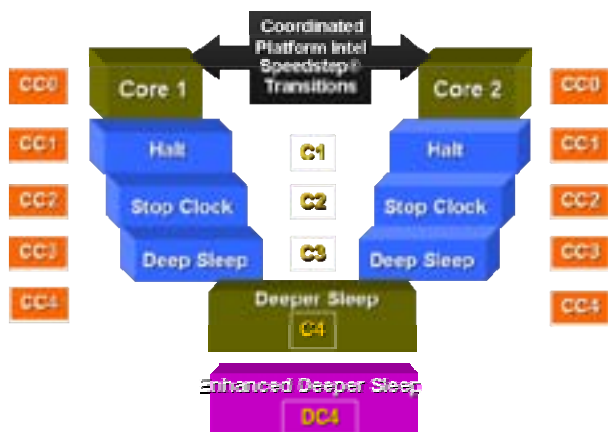


Figure 4: Low-power processor-design process

In order to save leakage power, the Intel Core Duo system uses mainly two techniques: enhanced sleep states control and Dynamic Intel® Smart cache sizing. In order to control the active power consumption, Intel Core Duo technology uses a technique based on Intel SpeedStep® technology.

The traditional way to control the power and the thermal of the system is via a software/hardware interface. One of the most common schemes to achieve this is called ACPI [5], where the system defines different levels of sleep modes, and each of the states represents a more efficient way to save power, at the expense of a longer time to bring the system back into operational mode. (For more details on this method, please see [2]). The challenge of adding a second core on die while improving the overall power-consumption demands an improvement to the power states of the system in order to avoid power being wasted whenever a core is not active. We face two main problems: (1) since only a single power plane is used, it forces us to run all cores with the same voltage and frequency, and (2) the chipset and the OS see both cores as a single entity that has the same state at the same time. Thus, the Intel Core Duo processor presents two separate views on the power state of the system; internally we manage the states of each core independently (we call it per-core power state) and externally we view the system as having a single, synchronized power state. Figure 5 provides an overview of this approach.



- CPU/package sleep states:
- **C0 – Active** CPU is on
- **C1 – Auto Halt** Core clock is off
- **C2 – Stop clock** Core and bus clock are off
- **C3 – Deep sleep** Clock generator is OFF
- **C4 – Deeper sleep** Reduced VCC
- **DC4 –Deeper C4** Further reduced VCC

Figure 5: Power states of the Intel Core Duo processor

As we can see the Intel Core Duo processor defines five different sleep states of the system. The first three states allow local power-saving measures to be activated individually per core, while the last two states require a coordination of the entire package for the power-saving measures to be activated.

A core which is in C₀ power state is assumed to be in running mode. When the core has nothing to do, the OS issues a halt command that moves it to CC1, where execution is halted and clocks are stopped. When it detects even lower levels of activity (via the ACPI mechanisms [2]), the OS will further promote the idle state of each of the cores beyond CC1 to CC2, CC3, or CC4 states, based on the core activity history. In the CC2 and CC3 states, additional core-level power-saving measures can be activated, achieving a lower average power consumption. Starting from C4 state, core voltage reduction is applied to further increase average power savings. Since the cores are connected to the same power plane, this must be done in coordination between the two cores, and this is known as package-level C4 and package-level DC4.

While being in a sleep state, the system still consumes static power (leakage). In Intel Core Duo technology, we implement an advanced algorithm that tries to anticipate the effective cache memory footprint that the system needs when moving from a deep sleep state to an active

mode. The new mechanism keeps only the minimum cache memory size needed active, and it uses special circuit techniques to keep the rest of the cache memory in a state that consumes only a minimal amount of leakage power.

In order to control the active power consumption, Intel Core Duo technology uses Intel SpeedStep technology. When a set of working points is defined, each one has a different frequency and voltage and so different power consumption. The system can define at what working point it works in order to strike a balance between the performance needs and the dynamic power consumption. This is usually done via the OS, using the ACPIs.




Figure 6: Changing working point in Intel Core Duo processor

The way the system moves from one working point to another is described in Figure 6. As illustrated, in order to move from a “high” working point to a lower one, the system can switch the frequency almost immediately, but it will take the system some time to lower the voltage. When moving from a low working point to a higher one, we need to increase the voltage first (slow operation) and only then can we increase the frequency.

By extending the hardware mechanisms to better support advanced power states and sleep states the Intel Core Duo processor achieves improved power performance efficiency. The power-efficiency improvement over processor generations is shown in Figure 7. As a result, the Intel Core Duo processor provides higher performance in the same form factor without needing to increase the cooling capability.




Figure 7: Power performance efficiency

THERMAL DESIGN POINT

Thermal management is another fundamental capability of all mobile platforms. Managing the platform thermals enables us to maximize CPU and platform performance within thermal constraints. Thermal management also improves ergonomics with a cooler system and lower fan acoustic noise.

In order to better control the thermal conditions of the system, the Intel Core Duo processor presents two new concepts: the use of digital sensors for high accuracy die temperature measurements and dual-core multiple-level thermal control.

In the previous Pentium M processor, a single analog thermal diode was used to measure die temperature. Thermal diode cannot be located at the hottest spot of the die and therefore some offset was applied to keep the CPU within specifications. For these systems it was sufficient, since the die had a single hot spot. In the Intel Core Duo processor, there are several hot spots that change position as a function of the combined workload of both cores. Figure 8 shows the differences between the use of the traditional analog sensor and the use of the new digital sensors.




Figure 8: Analog vs. digital sensors in Intel Core Duo processors

As we can see the use of multiple sensing points provides high accuracy and close proximity to the hot spot at any time. An analog thermal diode is still available on the Intel Core Duo processor. The use of a digital thermometer allows tighter thermal control functions, allowing higher performance in the same form factor. The improved capability also allows us to achieve better ergonomic systems that do not get too hot, can operate more quietly, and are more reliable. Unlike diode-based thermal management algorithms that require some temperature guard band (or activating the self throttle mechanism as a safety-net), the digital thermometer is tested and calibrated against specifications. Full functionality and reliability of the processor are guaranteed, as long as the reported temperature is equal to or below the maximum specified temperature. Any inaccuracy or offset are programmed into the device and already accounted for.

The thermal measurement function provides interfaces to power-management software such as the industry-standard ACPI. Each core can be defined as an independent thermal zone, or a single thermal zone for the entire chip. The maximum temperature for each thermal zone is reported separately via dedicated registers that can be polled by the software.

In addition to the polling capability, the digital thermometer implements event-based reporting. Control software programs temperature thresholds that require actions. Such actions can be fan activation or passive control policy such as dynamic voltage and frequency scaling. Upon temperature crossing of the threshold, an APIC-defined interrupt is generated and it initiates the requested action.

Intel Core Duo technology implemented a dual-core power monitor capability. Power monitor functionality is provided in order to prevent thermal exceptions, and it can throttle the CPU once the temperature exceeds specifications. The overview of the power monitoring logic is described in Figure 9.

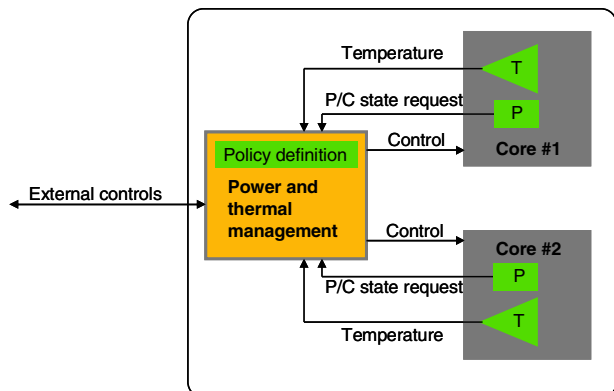


Figure 9: Thermal control overview

The power monitor continuously tracks the die temperature. If the temperature reaches the maximum allowed value, a throttle mechanism is initiated. A multi-level tracking algorithm is implemented. Throttling starts with the more efficient dynamic voltage scaling policy and if not sufficient, the power monitor algorithm continues lowering the frequency. If an extreme cooling malfunction occurs, an Out of Spec notification will be initiated, requesting controlled shutdown. Lastly, the CPU can initiate a thermal shutdown and turn off the system.

Power and thermal management activities in notebook computers are usually performed by the OS and platform control functions. These thermal management features are designed to best serve user preferences under notebook constraint conditions. Thermal monitor function is not expected to be activated under these normal operation conditions. The thermal monitor mechanism ensures that the CPU will never exceed the CPU-specified parameters and guarantees functionality and reliability at any time.

The use of high accuracy temperature reading together with thermal monitoring protection enables high performance in thermally limited form factors, while allowing improved ergonomics and high reliability.

PLATFORM POWER MANAGEMENT

Intel Core Duo processor technology closely interacts with other components on the platform. One such component is the Voltage Regulator (VR). VR power losses at low CPU utilization may get as high as the CPU power. The losses of the VR are due to the need to

deliver high current at quick response times. Intel Core Duo processors implemented a feedback mechanism to the VR. The CPU tracks its activity at any time. If utilization goes down, the CPU communicates a signal to the VR, allowing it to switch to a lower power consumption. A lower power state can be either a reduced number of phases or asynchronous operation. The communication is done using the voltage ID lines and PSI signal as described in Figure 10.

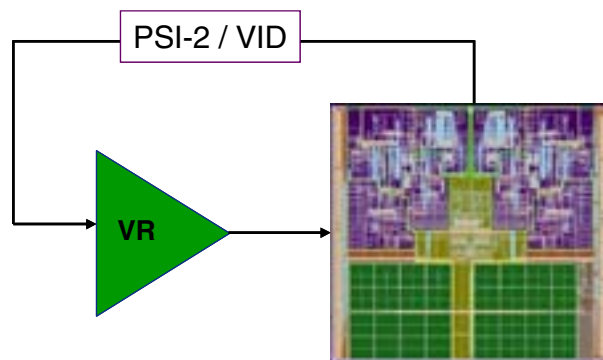


Figure 10: Voltage regulator interface

The CPU has internal knowledge of the activity demand and it communicates a request to go to higher power early enough for the VR to get ready for the increased demand.

Another power optimization is load line control. At low CPU activities, the voltage drop on the load line is smaller resulting in higher voltage and power to the CPU. At low workloads, the CPU reduces the voltage request, and early enough, before power consumption increases, a voltage increase request is sent to the VR.

Using utilization knowledge, available in the CPU, Intel Core Duo technology made it possible to reduce platform power, increase battery life, and improve form factor ergonomics.

INTEL® CORE™ SOLO PROCESSOR

In order to fit into very limited thermal constraints and power consumption, the Intel Core Duo processor has a derivative that contains a single core only. This can be achieved by either disabling one of the cores either at the OS level or as a BIOS option, or at the architecture level, where one core is disconnected from the power grid.

The first option is a user or OS decision. If you run a single-core OS on an Intel Core Duo system, it will keep the second core idle, at CC4 sleep state. Please note that due to the way the BIOS is set, each time an interrupt is received or a broadcast IPI is sent, this core may need to wake up and go immediately back to a sleep state, consuming small amounts of dynamic power.

The user can disable the second core via a BIOS option as well. In this case, the system does not recognize the other core and so it is kept in CC4 state all the time, consuming no dynamic power at all.

The disadvantage of the two methods described above is that the core still consumes static power. In order to avoid this and reduce the power consumption of the core even further, Intel introduces the single-core version of Intel Core Duo technology, called Intel® Core™ Solo processor, which disconnects the non-active core from the power grid, or saves the area and does not fabricate this part at all.

CONCLUSION

The Intel Core Duo processor is the first Intel processor that implements dual core on die. The processor addresses new challenges for providing the best performance under power and thermal constraints.

This paper described the main architectural features of the new processor focusing on the different performance, power, and thermal control features of the processor and of the system.

By applying punctual control between the performance, power and thermal features implemented in the Intel Core Duo system, we achieved a significant improvement in performance, at the same power consumption, and with improved thermal control mechanisms.

REFERENCES

- [1] Avi Mendelson, et al., "CMP Implementation in Intel® Core™ Duo Processor," *Intel Technology Journal*, Volume 10, Issue 2, 2006.
- [2] "Power and Thermal Management in the Intel® Core™ Duo Processor," *Intel Technology Journal*, Volume 10, Issue 2, 2006.
- [3] "Intel® Virtualization Technology Specification for the IA-32 Intel® Architecture" in <http://download.intel.com/technology/computing/vptech/C97063-002.pdf>
- [4] "IA-32 Intel® Architecture Software Developer's Manual Volume 1: Basic Architecture" in <http://download.intel.com/design/Pentium4/manuals/25366519.pdf> (chapter 13)
- [5] ACPI Specification at <http://www.acpi.info/spec.htm>*
- [6] G. Hinton, et al., "The microarchitecture of the Pentium 4 Processor," *Intel Technology Journal Q1*, 2001.

AUTHORS' BIOGRAPHIES

Simcha Gochman is a senior principal engineer with Intel's Mobile Platform Group in Haifa, Israel. Simcha has been with Intel for 21 years. Lately he has led the microarchitecture development of the Pentium M processors and of the Core Duo processor. Previous to this, he led the microarchitecture definition of the Pentium Processor with MMX technology and was involved in 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 at intel.com.

Avi Mendelson is a principal engineer in Intel's Mobile Platform Group in Haifa, Israel, and adjunct professor in the CS and EE departments, Technion, Israel Institute of Technology. He received his B.Sc. and M.S.c degrees from the Technion, Israel Institute of Technology and his Ph.D from the University of Massachusetts Amherst. Avi has been with Intel for 7 years. He started as senior researcher in Intel Labs, later he moved to the Microprocessor group where he serves as the CMP architect of Intel Core Duo processor. Avi's work and research interests are in computer architecture, low-power design, parallel systems, OS related issues and virtualization. His e-mail address is avi.mendelson at intel.com.

Alon Naveh is a senior architect with Intel's Mobile Platform 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 co-lead the power management definition of the Intel Core Duo architecture and was involved in the definition of the Intel Pentium M power management, PCI Express* and the Odem chipset. Prior to Intel, Alon worked in Motorola Semiconductor and in National Semiconductor. His e-mail is alon.naveh at intel.com.

Efi Rotem is a senior architect with Intel's Mobile Platform Group in Haifa, Israel, focusing on processor and platform power and thermal management. Efi joined Intel in 1995 and was involved in the definition and development of Pentium 4 and Pentium M processors. Previously at Intel, he led the Intel Pentium with MMX technology testing. He received a B.Sc. degree from the Technion, Israel Institute of Technology in 1986. His e-mail is efrain.rotem at intel.com.

Copyright © Intel Corporation 2006. All rights reserved. Intel, Core, Pentium, Intel SpeedStep, and MMX are trademarks or registered trademarks of Intel Corporation