Volume 07 Issue 02 Published, May 21, 2003 ISSN 1535-864X

Intel® Technology Journal

Intel® Centrino™ Mobile Technology

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

A compiled version of all papers from this issue of the Intel Technology Journal can be found at:
http://developer.intel.com/technology/itj/index.htm

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

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.

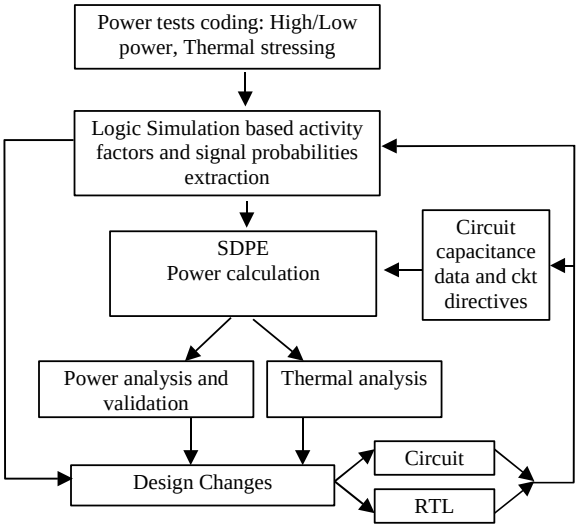

Figure 1: The power design cycle of the Intel Pentium M processor. Logic simulations were performed to collect activity statistics. The data were used by SDPE to calculate the power estimates. Analysis of the power results yielded design changes in the circuit and in the RTL.

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

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 (Vt) 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.

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.

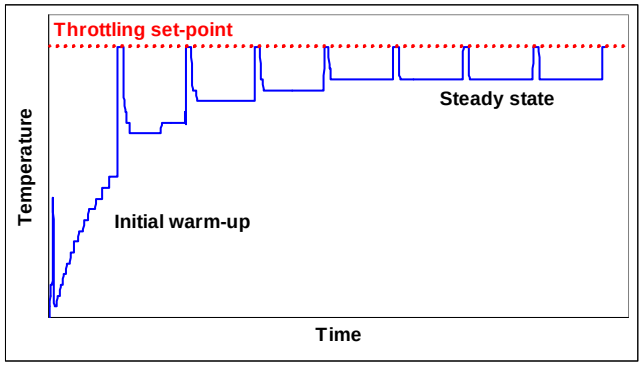

Figure 2: Bare die thermal throttling experiment demonstrating successful temperature clamping

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.

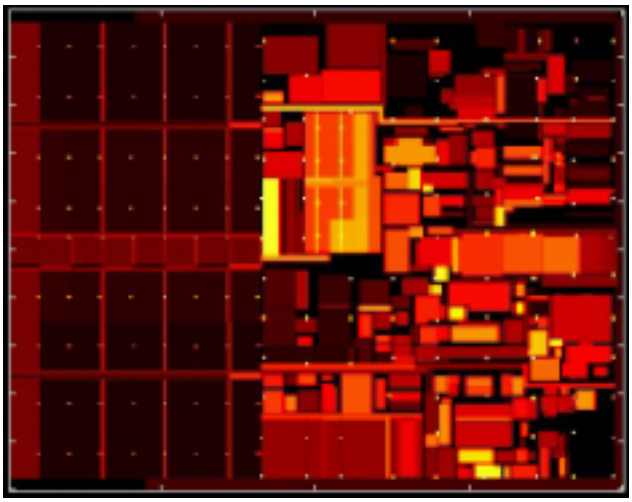

Figure 3: Simulated power density

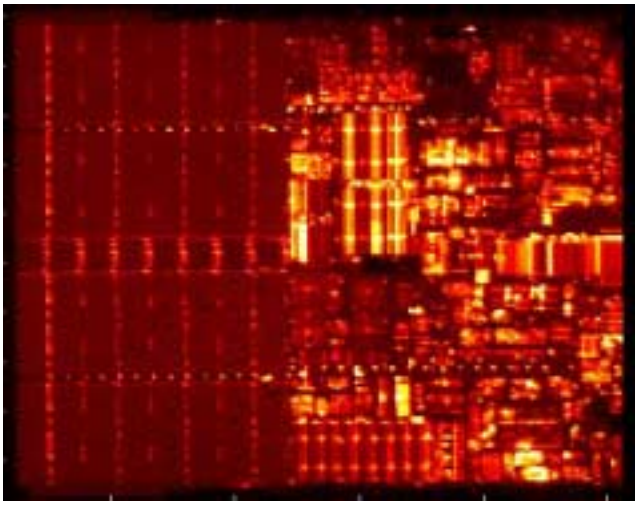

Figure 4: IREM measurement

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

For further information visit:

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