

Dynamic Power Management

Market Trends – Internet Search Results

Edward Herbert
September 20, 2005
<http://eherbert.com>

The switched current power converter with switched charge circuits has very fast di/dt and dv/dt, which I believe offers significant opportunity for reducing processor power consumption and heat generation. I did some WWW searching, to learn what I could of initiatives in that direction.

Following is a list of WWW links. If this document is viewed on a computer that is "on line", the links should be active. (They are verified as active at the time this is being written). There are also a number of very selective copied and pasted quotes that I thought were relevant to dynamic power management (a small part of some of the articles). I have also added comments of my own.

I tried to find out more about **Intel's recently announced power reduction initiatives**. I could not find anything that was substantive.

<http://www.eetimes.com/news/semi/showArticle.jhtml?articleID=169600298>

"By the end of the decade, Intel aims to slash power use by 10 times, Otellini said. Ultimately, improved power performance would enable developers and manufacturers to develop new systems with new form factors and capabilities, he noted."

I was curious about the rumored Intel-Rim licensing arrangement:

Intel and Rim:

<http://www.technewsworld.com/story/45711.html>

"A pact between chip king Intel (Nasdaq: INTC) and Research in Motion (RIM), Canadian maker of the Blackberry mobile communication device, appears less imminent than originally reported."

<http://www.pcmag.com/article2/0,1895,1852503,00.asp>

"Intel and ADI [Analog Devices Inc.] jointly worked on [power management] technology that is in today's BlackBerrys," said Ken Dulaney, vice president of mobile computing at Gartner Inc."

[More about Intel and Rim is further on.]

Key quotes. Note, these are taken out of context and some may seem somewhat contradictory. Please review the entire websites at the links. Underscores were added for emphasis. [More complete quotes and the links are further on.]

"AMD Cool 'n' Quiet™ technology controls your system's level of processor performance automatically, dynamically adjusting the operating frequency and voltage up to 30 times per second, according to the task at hand."

Intel: "Foxton is said to automatically adjust the voltage in a processor in 12.5-miliwatt increments at 32 different settings. It can also modulate the frequency in a processor at 64 different increments."

"In a Montecito processor featuring Foxton Technology, an on-die voltmeter and ammeter measure power draw every 8 microseconds and adjust the processor frequency accordingly."

"By stepping voltage up and down in small increments separately from frequency changes, the processor is able to reduce periods of system unavailability (which occur during frequency change)."

"One of the major advantages that Enhanced Intel SpeedStep Technology brings is lower latency associated with P-state changes – on the order of 10mS. ----"

APCI: "It is necessary to move power management into the OS and to use an abstract interface (ACPI) between the OS and the hardware ---- "

eWeek.com: " ---- there are two operating systems that have support, with another on the way: SuSE Linux Version 9 for the desktop and server, and Microsoft Windows XP with Service Pack 2. Still to come is Windows 2003 Server, which will support PowerNow when Service Pack 1 arrives.

Conclusions:

Dynamic power management is being pursued aggressively.

Changing frequency seems to be particularly slow, except maybe with Intel's Foxton.

Present voltage control (dv/dt) seems to be very slow.

I found no mention of turning off one or more cores entirely, with fast recovery.

Very fast dv/dt should be very useful, particularly the ability to turn on rapidly from a true off state. There should be good opportunities for the switched current power converter with switched charge circuits.

<http://eherbert.com>.

Technologies:

I found references to a number of technologies that shed some light on the market trends in processor power management.

ARM Intelligent Energy Manager
National's PowerWise technology
AMD Cool'n'Quiet
AMD PowerNow!
Intel SpeedStep Power Management
HP Power Regulator
Intel Foxton
Enhanced Intel SpeedStep Technology and Demand-Based Switching
Advanced Configuration and Power Interface Specification, (HP, Intel, MicroSoft, Phoenix Tech, Toshiba)

ARM Intelligent Energy Manager National's PowerWise technology

<http://www.arm.com/news/5539.html>

"National's PowerWise technology and the ARM Intelligent Energy Manager technology can reduce power consumption of systems-on-chip by dynamically matching minimal voltage and frequency to the chip's workload,"

<http://www.us.design-reuse.com/news/news5140.html>

"Through the energy management scheme, ARM is developing a set of predictive algorithms that can determine the processor speed required to execute a particular task. Then, using information from these algorithms, ARM allows its processors to toggle processor speed in order to meet those needs, Henry said."

"National is complementing ARM's efforts with the development of the PowerWise controller. Through the PowerWise technology, National's controller can adjust voltage depending on an operating condition, Henry said. Additionally, the power controller can more easily adapt to variations in temperature and process technology."

<http://www.arm.com/news/2189.html>

"ARM is the industry's leading provider of 16/32-bit embedded RISC microprocessor solutions. The company licenses its high-performance, low-cost, power-efficient RISC processors, peripherals, and system-on-chip (SoC) designs to leading international electronics companies. ARM also provides comprehensive support required in developing a complete system. ARM's microprocessor cores are rapidly becoming a volume RISC standard in such

markets as portable communications, hand-held computing, multimedia digital consumer and embedded solutions. More information on ARM is available at www.arm.com."

AMD Cool'n'Quiet AMD PowerNow!

As far as I can tell, AMD Cool'n'Quiet and AMD PowerNow! are the same technology, the AMD Cool'n'Quiet being the name when it is applied to desk tops and PowerNow! Being the name when it is applied to mobile applications.

http://www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/33609.pdf

"The P-state transitions (Cool'n'Quiet and AMD PowerNow! technologies) allow a computer to dynamically switch between different processor performance states depending on the processor utilization. In mobile applications these transitions can help extend battery life while delivering maximum processor performance on demand. In desktop applications they can reduce heat dissipation of the processor, resulting in less use of the fan and less noise."

http://www.amd.com/us-en/0,,3715_12353,00.html

"AMD PowerNow!™ technology can reduce CPU power at idle up to 75%!"

"The AMD Opteron™ processor family now includes AMD PowerNow!™ technology with Optimized Power Management (OPM). AMD PowerNow! technology for the AMD Opteron processor can decrease overall power consumption for enterprise IT and workstation customers by optimizing performance-on-demand."

- "Provides performance-on-demand by dynamically adjusting performance based on CPU utilization ---- ."

http://www.amd.com/us-en/assets/content_type/DownloadableAssets/Power_Now2.pdf

"AMD PowerNow! technology controls your notebook's level of processor performance automatically, dynamically adjusting the operating frequency and voltage many times per second, according to the task at hand. When an application does not require full performance, significant amounts of power can be saved. However, the processor can "instantaneously" respond to increased workloads, allowing the system to deliver a responsive and rewarding computing experience."

http://www.amd.com/us-en/assets/content_type/DownloadableAssets/Cool_N_Quiet_Installation_Guide3.pdf

" ---- AMD Cool 'n' Quiet™ technology controls your system's level of processor performance automatically, dynamically adjusting the operating frequency and voltage up to 30 times per second, according to the task at hand. When an application does not require full performance, significant amounts of power can be saved."

http://www.amd.com/us-en/Corporate/VirtualPressRoom/0,,51_104_543~92637,00.html

"The integration of AMD PowerNow! technology with OPM can reduce power consumption in the datacenter by dynamically changing power states based on workload utilization."

<http://www.eweek.com/article2/0,1895,1737607,00.asp>

"AMD Joins Push for Cooler, More Efficient Server Room"

"While the PowerNow technology has been built into Opteron processors produced since May, machines with these processors need a BIOS update from AMD to make it work. And, of course their operating system needs to support PowerNow. To date, according to Lewis, there are two operating systems that have support, with another on the way: SuSE Linux Version 9 for the desktop and server, and Microsoft Windows XP with Service Pack 2. Still to come is Windows 2003 Server, which will support PowerNow when Service Pack 1 arrives. Lewis said that support for PowerNow is available in the Linux kernel, but most distributions have yet to include it. She said that Red Hat Linux should include the capability soon.

Intel SpeedStep Power Management

The links just below are all for the Intel PXA27x family, a low power processor used in PDAs and the like, but they may give some insight into what Intel would like to do at higher power levels.

<ftp://download.intel.com/design/pca/applicationsprocessors/whitepapers/30057701.pdf>

This link is to a "White Paper" on SpeedStep Power Management.

"First available in the Intel PXA27x processor family, Wireless Intel SpeedStep® Technology provides the ability to dynamically adjust the power and performance of the processor based on CPU demand. ---- "

"4.4 Programmable Frequency Change Management (Intel DFM) and Programmable Voltage Change Management (Intel DVM)

"4.4.1 DFM

"The Intel PXA27x processor's core and peripheral clocks are derived from PLLs. The Intel PXA27x processor implements Intel DFM by allowing the core clock to be configured dynamically by software. ---- "

"4.4.2 DVM

"The Intel PXA27x Processor implements Intel DVM through its Voltage Manager. The Voltage Manager provides voltage management through use of an I2C unit (PWR_I2C) that is dedicated to communication with an external PMIC regulator, and through use of a Voltage Change Sequencer."

<http://www.intel.com/cd/ids/developer/asmo-na/eng/dc/pca/power/176608.htm?page=1>

An article on Power Management for the PXA27X processor.

<http://www.intel.com/design/pca/applicationsprocessors/manuals/280000.htm>

Intel® PXA27x Processor Family Developer's Manual, October 2004. This link has much more information: It is 1242 pages! The good stuff is in section 3. It takes 150 us to change frequency. The delay to change voltage is programmable, I think it's 30.5 us increments, and the default is 125 ms. In 3.7.6.5, there seems to be a provision to accept an interrupt from the external regulator when the transfer is complete.

<http://www.intel.com/design/pca/applicationsprocessors/datashts/280003.htm>

Intel® PXA27x Processor Family Electrical, Mechanical, and Thermal Specification. More information about the PXA27x processor family.

HP Power Regulator

<http://h18004.www1.hp.com/products/servers/management/ilo/power-regulator.html>

"This new power management technology enables dynamic or static changes in CPU performance and power states. In dynamic mode, Power Regulator automatically adjusts the server's processor power usage and performance to match CPU application activity. This improves server energy efficiency by giving CPUs full power for applications when they need it and enabling power savings without performance degradation when application activity is reduced. Power Regulator effectively executes automated policy-based power management at the

individual server level. In addition, a unique static low power mode allows servers to run continuously in a system's lowest power state."

Intel Foxton

Intel Enhanced Intel SpeedStep Technology and Demand-Based Switching

<http://www.powermanagementdesignline.com/news/60401054>

"The low-power technology, dubbed Foxton, is initially targeted for Intel's dual-core, 64-bit Itanium processor line, but it will be eventually incorporated within its 32- and 64-bit Xeon chips."

" ---- Foxton is said to automatically adjust the voltage in a processor in 12.5-miliwatt increments at 32 different settings. It can also modulate the frequency in a processor at 64 different increments."

<http://www.intel.com/technology/magazine/computing/foxton-technology-0905.htm>

A new Intel technology code named "Foxton" provides a mechanism for select Intel® Itanium® 2 processors to adjust core frequency during operation to boost application performance and make the most of available power. This capability will be available for the first time on the forthcoming dual-core Intel Itanium processor code named "Montecito," and will fundamentally change the way that developers consider frequency as a factor in software and hardware design.

In a Montecito processor featuring Foxton Technology, an on-die voltmeter and ammeter measure power draw every 8 microseconds and adjust the processor frequency accordingly.

<http://www.eweek.com/article2/0,1895,1760383,00.asp>

Writing about Montecito:

"The chip will offer two cores on a single die, with each core running multiple threads simultaneously. A power meter on the chip lets Foxton dynamically adjust processor voltage and frequency according to demand, resulting in a performance boost while maintaining power consumption. Foxton will be brought into Intel's Xeon processors at a later date, Modi said."

"In addition, Montecito will feature demand-based switching capabilities currently found in Intel's Xeon and Pentium processors. The technology works with the server's operating system to reduce energy consumption at times of low utilization, Modi said."

<http://www.intel.com/cd/ids/developer/asmo-na/eng/195910.htm>

"Enhanced Intel SpeedStep Technology and Demand-Based Switching on Linux:"

"Separation between Voltage and Frequency Changes. By stepping voltage up and down in small increments separately from frequency changes, the processor is able to reduce periods of system unavailability (which occur during frequency change)."

"Clock Partitioning and Recovery. The bus clock continues running during state transition, even when the core clock and Phase-Locked Loop are stopped, which allows logic to remain active."

Page 4:

"One of the major advantages that Enhanced Intel SpeedStep Technology brings is lower latency associated with P-state changes – on the order of 10mS. ----"

http://news.com.com/Itanium+not+meeting+Intels+goals/2100-1010_3-5349588.html

"Montecito will arrive at the "tail end" of 2005 ---- and it will use the "Foxton" technology, which speeds up the processor 10 percent to 15 percent when it's running cool enough."

Advanced Configuration and Power Interface Specification, (HP, Intel, MicroSoft, Phoenix Tech, Toshiba)

<http://www.acpi.info/DOWNLOADS/ACPIspec30.pdf>

Advanced Configuration and Power Interface Specification, (HP, Intel, MicroSoft, Phoenix Tech, Toshiba) 618 pages. See pages 18 and 41.

"1.2 Power Management Rationale

"It is necessary to move power management into the OS and to use an abstract interface (ACPI) between the OS and the hardware to achieve the principal goals set forth above.

- "Minimal support for power management inhibits application vendors from supporting or exploiting it.
- "Moving power management functionality into the OS makes it available on every machine on which the OS is installed. The level of functionality (power savings, and so on) varies from machine to machine, but users and applications will see the same power interfaces and semantics on all OSPM machines."

"3.1 System Power Management

"Under OSPM, the OS directs all system and device power state transitions. Employing user preferences and knowledge of how devices are being used by applications, the OS puts devices in and out of low-power states. Devices that are

not being used can be turned off. Similarly, the OS uses information from applications and user settings to put the system as a whole into a low- power state. The OS uses ACPI to control power state transitions in hardware."

Intel and Rim

More about the rumored Intel and Rim collaboration (dated 08/25/05):

<http://www.crmbuyer.com/story/45711.html>

"Dual Core Difficulties

"According to a research report released this week and written by Peter Misek and Dushan Batrovic, two investment analysts for Canaccord Capital, of Vancouver, B.C., Intel is having trouble making a mobile version of its dual core chipset.

"Our checks originally lead us to believe that Intel was encountering considerable difficulty in developing a battery, bandwidth and heat-efficient Centrino dual core chipset," the pair wrote in their report, a copy of which was obtained by TechNewsWorld.

"According to the duo, RIM has technology that could address those three problems. "RIM's technology presents considerable advantages to Intel's original dual core design" they wrote. "[W]e have heard that RIM's combined solution would improve battery life three-fold and significantly reduce the Centrino's bandwidth and heat generation."

"We believe this partnership would make sense for Intel, they added, "additionally because the company is expected to be looking beyond laptops and PCs towards the cell phone market, which represents an 800 million-unit per year opportunity."

Miscellaneous links:

[Enterprise Platform Power Efficiency for Future Server Systems](#)

Tomm Aldritde, Manager – Enterprise Power & Thermal Technology Labs
Intel Corporation — IBM Platform Technology Symposium September 14-15, 2004
Kahler Grand Hotel, Rochester, MN

[Taurus: A Taxonomy of the Actual Utilization of Real UNIX and Windows Servers](#),
David G Heap, Principal IT Consultant, IBM Enterprise Server Group, Somers, NY,
USA