Intel Corporation
2200 Mission College Blvd.
P.O. Box 58119
Santa Clara, CA 95052-8119

News Fact Sheet

CONTACT: Megan Langer
503-712-4305
megan.e.langer@intel.com

INTEL TECHNICAL DISCLOSURES AT ISSCC

Feb. 3, 2008 -- Intel Corporation is presenting 15 technical papers at the International Solid State Circuits Conference (ISSCC), Feb. 3-8 in San Francisco.

Under the theme of wireless mobility and the company's vision to deliver a complete Internet experience ranging from your pocket to other CE devices, Intel will detail its forthcoming 45 nanometer high-k metal gate, low-power processor architecture, codenamed "Silverthorne," for ultra-mobile and mobile Internet devices.

Intel researchers will also present significant milestones in developing low-cost, digital multi-radios that in the future could allow a variety of small devices to handle a mixture of wireless radio technology standards from just one chip that consumes less power than today's bulky analog versions.

In addition, Intel will disclose more information in the following areas: the company's ongoing Terascale project and effort to deliver everyday processing that exceeds a trillion operations per second or TeraFLOPS speed; more detail on the company's 45nm high-k metal gate process; its next-generation, first 2 billion transistor chip -- the Intel® Itanium® processor codenamed "Tukwila"; and, the company's phase change memory advances (as part of the pending Numonyx company).

◆ ◆ ◆ ◆ ◆

PROCESSOR TECHNOLOGY

"A Sub-1W to 2W Low-Power IA Processor for Mobile Internet Devices in 45nm High-K Metal-Gate CMOS"
Session 13.1, Feb. 5, 8:30 a.m.

-- more --

Intel disclosed details about its new low power IA microarchitecture, which is the basis of the 45nm high-k metal gate Silverthorne processor being targeted for first-generation Mobile Internet Devices.

- This microarchitecture will be fully compatible with the Core 2 Duo instruction set, is based on dual-code, dual-issue, in-order execution and implements a 16-stage processor pipeline. The microarchitecture will implement ground-breaking power management techniques such as Deep Power Down (C6) state, non-grid clock distribution, power-optimized register-file, clock gating, CMOS bus mode and split IO power supply to aggressively reduce dynamic and leakage power.
- As a result of these innovative power management techniques, the 45nm high-k metal gate Silverthorne processor is expected to achieve a 10x lower thermal power level (compared to the Ultra Low Voltage single-core Intel processors in 2006) while delivering high performance to run the full Internet and breadth of software applications.
- The microarchitecture, which was built from the ground up, will also deliver high performance at sub-watt power levels.

◆ ◆ ◆ ◆

"A 65nm 2-Billion-Transistor Quad-Core ItaniumR Processor"
Session 4.6, Feb. 4, 4:15 p.m.

Intel describes the world's first 2 billion transistor microprocessor, a quad-core Itanium microprocessor codenamed Tukwila, which enables a leap in performance and capabilities, higher system integration, advanced RAS features and larger caches.

- Tukwila features 30MB of total on-die cache, representing a more than 10 percent increase over the current generation.
- The new high-speed QuickPath interconnect and dual integrated memory controllers are coupled with leading-edge interconnect RAS. The quad-core chip is coupled with higher bandwidths and large caches to enable a doubling in performance of Tukwila over the current Intel® Itanium® 9100 series processor.
- Accomplishing a more than 2x* performance increase despite a high level of system integration (Quad-Core, QuickPath interconnects, integrated memory controllers, advanced RAS, large cache, etc.) on die is a huge achievement for a processor used in this mission critical segment.
- Tukwila features improved Soft Error Rate (SER) immunity. This means that soft error hardened circuits have been designed for advanced RAS with the goal of equivalent "Persocket SER" as the previous generation despite as much as a 3x* increase in logic circuits.
- The circuits have been designed to allow voltage and frequency management for optimal use of the power and thermal envelope of the processor. This allows balancing performance and power savings, as needed.
- The first version of Tukwila is scheduled to arrive towards the end of this year.

◆ ◆ ◆ ◆

WIRELESS COMMUNICATION TECHNOLOGY

"A 1x2 MIMO Multi-Band CMOS Transceiver with an Integrated Front End in 90nm CMOS for 802.11agn WLAN Applications"
Session 20.1, Feb 5, 1:30 p.m.

Since the advent of mobile laptop platforms into the wireless communication market, the usage of Wi-Fi in everyday life has seen an exponential increase. Further penetration into additional platforms and products, such as MIDs, handhelds and PDAs, requires a dramatic reduction in cost and smaller form factors of the wireless device. This can only be achieved by a higher level of component integration onto the silicon radio chips.

In this paper, Intel presents preliminary results of a unique radio chip design, implemented in standard 90nm CMOS process, which fully integrates the LNAs and high efficiency class-AB PAs (and their matching networks) in a 1x2 scheme for 802.11agn protocols. This design enables low power consumption, small form factor and low cost. Other key points in this paper include:

- Power efficient, full dual band TX with full power on chip class AB PAs + Digital Pre-distortion.
- Advanced Digital-Pre-Distortion Calibration for excellent performance and system stability.
- Dual band (2.4G and 5-6G) LNA integration.

◆ ◆ ◆ ◆ ◆

"A 28.6dBm, 65nm Class-E PA with Envelope Restoration by Pulse-Width and Pulse-Position Modulation"
Session 31.5, Feb 6, 3:45 p.m.

In order to support long-range communications (e.g. WiMAX), high power amplifiers (PA) in the vicinity of 1 watt are needed. In this paper, Intel researchers demonstrate a PA implemented in a modern 65nm process with no analog components. This process will allow integration of the PA with the rest of the transceiver which will reduce development costs. Additional circuit, layout and techniques and operation in switching mode were all implemented to maximize power efficiency.

Additionally, this paper describes a new approach to amplitude/power control where the pulse width of the input signal going to the switching PA is adjusted to give different amplitude/power levels. The proposed technique can move some of the burden of introducing amplitude information to the digital domain where it might be easier and more economical to solve rather than conventional approaches. Other key points in this paper include:

- Delivers close to 1 watt power for wide coverage.
- Uses novel technique to introduce complex modulation required for high data rates.
- Implemented in digital 65nm CMOS process for easy integration with other digital elements.

◆ ◆ ◆ ◆ ◆

"A 28mW Spectrum-Sensing Reconfigurable 20MHz 72dB-SNR 70dB-SNDR DT ΔΣ ADC for 802.11n/WiMax Receivers"

Session 27.5, Feb 6, 10:45 a.m.

As the wireless spectrum for Wi-Fi becomes increasingly crowded, the ability for a radio to automatically select the band with the least amount of interference is necessary.

In this paper, Intel researchers present the lowest-power 802.11n ADC and the first reconfigurable ADC for 802.11n and multi-radio applications. The analog to digital converter in this paper uses high sampling frequency that enables it to measure the power of every band in the entire Wi-Fi band. Furthermore, the increased speed of the ADC is exploited to increase the dynamic range for the ADC within the band of interest. This allows for a very digital friendly implementation of the radio as the analog filters of the radio can be replaced by digital filters. The power consumption of the ADC is lowest in its class and is a demonstration of the fact that increased process speed can be exploited for better performance.

- The 12-bit ADC allows analog to be replaced by digital circuits which enables manufacturing cost benefits.
- Senses interference from other radios in the same band and adjusts itself for optimal power/performance.
- Reduces power requirements when signal is strong.
- Provides optimal channel selection to maximize real-life throughput.
- Supports Wi-Fi/WiMAX bandwidths in a power efficient manner.

◆ ◆ ◆ ◆ ◆

"A 39.1-to-41.6GHz ΣΔ Fractional-N Frequency Synthesizer in 90nm CMOS"
Session 26.7, Feb 6, 11 a.m.

The unlicensed bandwidth available at 60GHz is making mm-wave technology attractive for multi-Gb/s, consumer-market applications. For example, a 2GHz channel would provide a 5Gb/s data-rate which would enable a consumer to wirelessly transfer a full HD movie from one device to another in under one minute (compared to 1.5 hours for legacy WLAN).

In this paper, Intel researchers, in collaboration with Georgia Tech, are demonstrating the first ever mm-wave CMOS synthesizer with <3kHz frequency resolution. Freq. synthesizers are used to generate the local oscillator signal that is used in radios to downconvert-upconvert the baseband signal. A fractional synthesizer can achieve a much finer resolution than traditional integer-N synthesizers. This finer resolution can be used to perform frequency correction and tracking using a much cheaper crystal. Other key points include:

- The mm-wave CMOS technology for multi-Gb/s wireless communication is a fundamental building block required in CMOS for integrated mm wave radios.
- This solution will reduce size and improve yield with built-in calibration.
- Low power operation due to injection locking divider by 4 with self-calibration.
- This solution offers Multi-Gb/s data-rate for (WPAN, wireless-HD etc).

◆ ◆ ◆ ◆ ◆

MEMORY TECHNOLOGY:

"A Multi-Level Cell Bipolar-Selected Phase Change Memory"

Session 23.5, Feb. 6, 10:45 a.m.

This paper will describe a breakthrough in Phase Change Memory (PCM) from the joint development program of Intel and ST Microelectronics. Together, the companies created the world's first demonstrable multi-level cell (MLC) device using PCM technology. PCM is based on changing the states of a chalcogenide material (Ge2Sb2Te5) also known as GST. It is a promising new memory technology that provides for very fast read and writes at lower power than conventional flash and allows for more stable data retention – possessing many of the best attributes of leading memory technologies. The move from single bit per cell to MLC also brings significantly higher density at lower cost per Mbyte – making the combined MLC on PCM a powerful development.

- By using a unique programming algorithm the authors have effectively created two additional states between amorphous and crystalline.
- Using H_20 states as an analogy a single level cell PCM device would be akin to observing H_20 in either a liquid (water) state or in a crystalline (ice) state. With the MLC technique the authors show they can now control the GST and set it in four states. Stretching the $_{H20}$ analogy we can now observe the following states: very amorphous (gas or "00"), amorphous (liquid or "01"), semi-crystalline (liquid w/ some ice cubes or "10"), and crystalline (solid block of ice or "11").
- The process features nine copper interconnect layers with extensive use of low-k interlayer dielectrics for improved power and performance integrated with lead-free packaging.
- Previously, Intel and ST Microelectronics demonstrated 4Mb memory arrays on 180nm and a 128Mbit memory device on 90nm using PCM. The 2008 ISSCC paper now shows data demonstrating a multi-level cell (MLC) device using PCM technology.

◆ ◆ ◆ ◆ ◆

"A 45nm Self-Aligned-Contact Process 1Gb NOR Flash with 5MB/s Program Speed"
Session 23.3, Feb. 6, 9:30 a.m.

The design presented in this paper encompasses the smallest reliable flash cell in the state of the art 45nm technology with 5 MB/sec program performance, the smallest periphery circuitry and a robust sense scheme. It is essential to transition rapidly to the next technology node to reduce production costs and at the same time deliver higher performance. To be successful in the marketplace, the new 45nm lithography technology must reduce cost per bit by 50 percent while delivering higher program performance. However, each lithography generation proves more difficult to build a reliable Flash Multi Level Cell (MLC), and reduce the die size impact of periphery circuitry around the array, especially for lower density products.

- The paper shows the Self-Aligned-Contact (SAC) process architecture, which allows cell size reduction and improves reliability of the flash cells.
- To achieve program performance of 5MB/sec, it required the development of a number of new circuit techniques. These include larger program bandwidth, faster verify mode, increased slew rates of high voltage nodes, maximized throughput of the program microcode and reduced control hardware delays.
- Innovative circuit techniques are used to overcome detrimental effects such as single cell charge loss/gain, induced charge loss, and random telegraphic noise.

- A new sensing scheme with lower input offset and fewer components increased the MLC sense margins Results show the input offset voltage for 1 sigma of Vt mismatch is less than 1mV and a sense amp offset (SAOS) reduction of 70 percent.
- An aggressive die size goal of 30mm2 for 1Gb required many periphery circuit improvements in row decoders, block redundancy scheme, charge pumps and logic circuits.

◆ ◆ ◆ ◆ ◆

"A 153Mb-SRAM Design with Dynamic Stability Enhancement and Leakage Reduction in 45nm High-K Metal-Gate CMOS Technology"
Session 21.1, Feb. 5, 1:30 p.m.

Intel developed a high-performance and low-power SRAM based on the industry's first 45nm high-k metal gate technology that provides a 2x density improvement, 10x leakage reduction, and 27 percent frequency increase over Intel's 65nm technology. Intel is already offering 32 products based on its 45nm Hi-k metal gate technology.

- Intel's 45nm SRAM takes full advantage of the significant benefits provided by the high-k metal gate technology in scaling, including power and performance, and it enabled a 50 percent larger on-die L2 (6MB) cache for rapid high-volume production in the second generation Intel® Core™ 2 Duo and Core 2 Quad microprocessors. Small SRAM cells allow for the integration of larger caches in processors, which helps increase performance.
- Intel's SRAM design features robust timing control schemes for high-volume manufacturing, along with efficient power management circuits, which gives circuits better tolerance for variations and helps achieve higher manufacturing yield.
- Intel has developed second generation dynamic sleep technology that further maximizes static power reduction in large caches under all process, voltage and temperature variation conditions. The enhanced design allows Intel to achieve a larger power reduction in on-die caches.
- Intel has also developed a new circuit technique called dynamic body biasing to further improve SRAM cell scalability for future scaling.

◆ ◆ ◆ ◆ ◆

"A 50nm 8Gb NAND Flash Memory with 100MB/s Program Throughput and 200MB/s DDR Interface"
Session 23.4, February 6, 10:15 a.m.

Recently introduced new technology from Intel and Micron will be presented.In the paper, the technologists will discuss a new high-speed NAND (HS-NAND) flash technology that can greatly enhance the access and transfer of data in devices that use silicon for storage. The new technology was developed jointly by Intel and Micron and manufactured by the companies' NAND flash joint venture, IM Flash Technologies (IMFT). Benefits of the technology include:

- Five times faster than conventional NAND allowing data to be transferred in a fraction of the time for computing, video, photography and other computing applications.
- Speeds up to 200 megabytes per second (MB/s) for reading data and 100 MB/s for writing data, achieved by leveraging the new ONFI 2.0 specification and a four plane architecture with higher clock speeds.

- In comparison, conventional single level cell NAND is limited to 40 MB/s for reading data and less than 20 MB/s for writing data.

◆ ◆ ◆ ◆ ◆

TERA-SCALE TECHNOLOGY:

"A 27Gb/s Forwarded-Clock I/O Receiver Using an Injection-Locked LC-DCO in 45nm CMOS"
Session 25.1, Feb 6, 8:30 a.m.

Tera-scale technology is Intel's vision of future platforms with 10-100s cores sharing connections to memory, other sockets and peripherals. To enable data-intensive emerging applications, I/O bandwidth must scale to >100Gbps, meaning each channel must be >10Gbps. Scaling the speed of I/O channels requires accurate clocks to time transmission and reception of data, which consumes large amounts of power, requires large area for filter components, and complex circuits to mitigate noise effects.

A 45nm test chip with new techniques for utilizing a forwarded clock signal (a clock transmitted on a separate channel with data) to time data at the receiver with fewer, simpler circuits. It removes large filter components yet filters high-frequency clock jitter (timing noise). Technically, we show that rather than a full PLL -- only the VCO portion of the PLL is truly needed -- and that relying on this simpler circuit saves energy while increasing performance. Results from our test chip show:

- Extremely high electrical speeds – enabling data links up to 27Gb/s per data line.
- The best energy efficiency of any I/O receiver over 20Gb/s: 1.6mW/Gb/s.

◆ ◆ ◆ ◆ ◆

"Energy-Efficient and Metastability-Immune Timing-Error Detection and Instruction-Replay-Based Recovery Circuits for Dynamic-Variation Tolerance"
Session 22.2, Feb. 7, 9 a.m.

In order to achieve tera-scale performance processors must be extremely efficient and have the ability to maximize performance/watt. The performance and energy efficiency of processors today are limited by the fact that architectures cannot tolerate any circuit-level timing errors. Although such errors may be temporary and very rare, to ensure correct operation, max speed must reduced and min voltage increased by a safety factor called a "guard band."

Intel researchers demonstrate a test chip with resilient circuits that detect and correct timing errors, eliminating the need for guard bands. Results from this test chip show that this is the lowest energy and fastest error-detection sequential circuit published to date. These techniques could lead to similar benefits for future processors. The test chip showed that these circuits allowed for:

- Up to a 32 percent performance gain by "overclocking" the chip (keeping voltage constant).
- Up to a 33 percent reduction in energy consumption by reducing the voltage (keeping performance constant).

- (Or) a combination of both by changing both settings.

◆ ◆ ◆ ◆ ◆

"A 320mV 56µW 411GOPS/W Ultra-Low Voltage Motion Estimation Accelerator in 65nm CMOS"
Session 16.6, Feb. 6, 4:15 p.m.

In order to provide the best performance/watt for both high performance and ultra-mobile applications, some of the transistors in future chips may be dedicated to accelerating such commonly used tasks as HD video processing. Such accelerators are task-specific, but can provide 5-10x better performance/watt. A key component of most video compression techniques is "motion estimation," which is used to identify redundant image data created when the same objects are moving from frame to frame. This task occupies 60-80 percent of the processing for video compression. Accelerating this could make HD compression possible on small devices and much faster on larger systems.

Intel researchers have demonstrated a video motion estimation accelerator with multiple performance improvements and the ability to operate at ultra-low voltages not typically attainable for most circuits. An array of such accelerators could enable ultra-low power video encoding for mobile devices. This test chip showed:

- Up to 10x better throughput than best reported accelerator (running a 3-step search algorithm).
- Ability tune voltage and performance to optimize energy efficiency to the task at hand.
- Operates below the normal minimum voltage (i.e. sub-threshold) down to 0.22 V.
- Ultra-low voltage of 0.3V yields 10x improvement in energy efficiency (411GOPS/W).

◆ ◆ ◆ ◆ ◆

"2GHz 2Mb 2T Gain-Cell Memory Macro with 128GB/s Bandwidth in a 65nm Logic Process"
Session 14.3, Feb. 5, 9:30 a.m.

Analysis of future tera-scale applications indicates that running multi-threaded workloads on many cores will significantly increase the demand for memory bandwidth. On-chip SRAM memory is very fast but costly in terms of the chip area it requires. DRAM, used for "main memory," is much denser, though much slower. It also cannot be integrated on the microprocessor due to differences in the manufacturing process. It can be closely integrated via 3-D stacking, but even this does not approach the speed of on-chip memory.

This paper demonstrates a new category of integrated DRAM memory that can be made on a standard microprocessor process. This provides a new choice for chip designers to provide more fast memory on-chip and increase the performance of future application. The memory must be "refreshed" regularly as with other dynamic memories, but this would in turn provide:

- Twice the memory density vs. on-chip SRAM.
- Much faster speed than DRAM: 128GB/s running at 2GHz.

◆ ◆ ◆ ◆ ◆

PROCESS TECHNOLOGY

"A 45nm Logic Technology with High-k+Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193nm Dry Patterning, and 100% Pb-free[1] Packaging"
Session 6, Highlights of IEDM 2007, Feb. 5 (paper originally presented at IEDM 2007)

This paper will describe Intel's breakthrough 45nm processing technology, the world's first with high-k metal gate transistors. The new gate stack is combined with enhanced third-generation strained silicon to produce n-type metal oxide semiconductor (NMOS) and p-type metal oxide semiconductor (PMOS) transistors with the highest drive currents reported to date. Logic gate delay will be reported to improve by more than 20 percent, compared to logic gate delay in 65nm. The technology has produced multiple functional microprocessors and is already in high-volume manufacturing. Intel recently launched its first 45nm processors based on its high-k metal gate transistor technology.

- The paper will highlight another first for the technology: the use of trench (rectangular) contacts to replace square contacts, providing improved performance and local routing capability for improved layout density.
- In addition, the paper will discuss key design rules to achieve density scaling. The technology features the smallest transistor pitch reported at the 45nm generation, providing better transistor packing density as well as a small static random access memory (SRAM) cell size of 0.346μm2. High transistor performance at a small transistor pitch shows that no fundamental conflict exists between performance and density as some others have suggested.
- The process features nine copper interconnect layers with extensive use of low-k interlayer dielectrics for improved power and performance integrated with lead-free[1] packaging.
- For the first time, the process integrates a very thick copper power redistribution interconnect layer using a polymer inter layer dielectric (ILD).

-- 30 --

[1] 45nm product is manufactured on a lead-free process. Lead-free per EU RoHS directive July 2006. Some E.U. RoHS exemptions may apply to other components used in the product package.