

M I C R O P R O C E S S O R

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

ON THE ROAD TO MANY-CORE μ Ps

Cores Proliferate at Intel's Spring '05 Developer Forum

By Kevin Krewell {3/21/05-02}

In 2004, Intel publicly committed to a processor roadmap that embraced dual- and multi-core processing from servers to desktops to mobile processors. At the Fall 2004 Intel Developer Forum (IDF), however, Intel had only limited demos to show. Fast-forward to March's

Spring 2005 IDF. At the Spring 2005 IDF, Intel had many demos on the stage and a demo pavilion with multiple systems running on multiple versions of Intel dual-core processors. What a difference six months make! Intel also revealed some surprising implementation details of its multicore manufacturing strategy.

Intel branding for the dual-core processors will drop the numeral reference. The dual-core desktop processor, based on the Net-Burst (Pentium 4) architecture, will be called the Pentium D processor. For the Pentium D, Intel will not enable Hyper-Threading, so this version will perform almost identically to a Xeon dual-processor SMP system today. Intel will also offer a version of the dual-core Net-Burst processor with Hyper-Threading enabled, which will be branded as a Pentium Processor Extreme Edition. The Pentium EE will support four threads for the operating system.

So far, Microsoft has decided to support the one-socket/one-operating-system/one-license pricing model for WindowsXP. That should mean that Microsoft does not care if the dual-core processor consists of one or more actual die in the package. When we get to virtualization technologies, with multiple operating system instances running, however, we will likely see Microsoft charging per instance.

What Makes a Dual-Core Processor?

The biggest surprise from the show was that a number of future dual-core processors are not on a monolithic die. The first dual-core processor Intel will ship, code-named

Smithfield, will consist of two Net-Burst cores on one die and will ship in desktop systems as the Pentium D processor. The follow-on 65nm dual-core part, code-named Presler, will have two Net-Burst cores on two separate die, but both die will be mounted in one package. This is the first time Intel has ventured into volume microprocessor multichip modules (MCM) since the original PentiumPro. (Intel did ship Pentium II and Pentium III processors in modules; those modules actually consisted of a processor core with cache SRAM chips on a daughtercard, enclosed in a plastic shell.) Although there is no user-perceivable difference between the monolithic and separate die versions of Net-Burst processors (which we'll refrain from calling the Pentium 4 architecture, as the "4" seems to be on its way out), there are manufacturing issues. There is also the question of what constitutes a "true" dual-core processor. Does die-packaging integration count—or does it even matter?

Intel has certainly blurred the line with Presler and the Xeon server version, code-named Dempsey. If MCM packaging of cores counts as dual- and multicore processors, then the IBM Power 4 processor should have been considered an eight-core processor, because IBM shipped an MCM that incorporated four dual-core die. (See *MPR 10/06/99-02*, "Power4 Focuses on Memory Bandwidth.") Like Presler, the Power 4 doesn't require glue logic to attach multiple cores inside the MCM package.

The clear way to describe Smithfield, Presler, and Dempsey would be as dual-CPU processors. All these desktop




Figure 1. The Smithfield (left) and Presler processors generate two bus-interface loads on the Pentium 4 bus, unlike a single-core processor. The extra bus loading of these dual-core designs will limit the bus speed to 800MHz. The Xeon version of Presler is called Dempsey, and a single-core/die version is called Cedar Mill.

and Xeon processors have two CPUs in one package, and they enumerate to the operating systems as CPU 0 and CPU 1. The enumeration difference between a dual-processor SMP system and a dual-CPU Pentium D is trivial.

There are some disadvantages to this independent-core approach to multicore processors. Smithfield, Dempsey, and Presler use the Pentium 4 front-side bus to connect the cores, just as if it were a dual-processor SMP (see Figure 1). The problem is that Smithfield, Dempsey, and Presler processor cores generate two bus-interface loads on the processor bus. A slightly more integrated Paxville processor has the cores share a single bus interface and have only one bus load (see Figure 2) on the bus. The impact of the extra bus loading of Smithfield is that the high-end Pentium Extreme Edition processor Intel plans to ship this year will have to slow its bus to 800MHz—from the 1,066MHz of the Prescott version of the Pentium 4 Extreme Edition. The slower front-side bus will be a potential bottleneck for the Extreme Edition processor, because it will be a high-clock-speed processor (likely 3.2GHz) with up to four threads. Any program that is bandwidth sensitive may not scale well on Smithfield.

The Platforms for the Dual Core

The desktop dual-core processors will be supported by new desktop chip sets called the Intel 945G and 955X Express chip sets. The dual-core processors will not boot in the existing 915 and 925 Express chip sets. The 915/925 chip sets were




Figure 3. A dual-processor system based on the dual-core Dempsey processor will connect to the Blackford chip set with two independent buses. The loading of the bus is similar to the desktop part and will likely run at 800MHz.




Figure 2. The bus interface for the dual-core Xeon processor, code-named Paxville, has an integrated bus interface, limiting the bus loading to one load. Intel had to integrate the bus interface in order to support two sockets on one bus and to be able to run the bus at speeds exceeding 533MHz (the frequency limit for a bus with four processor bus loads).

designed for single-core processors and do not support the enumeration of multiple processors, which is the domain of the server chip sets today. The 945G/955X chip sets add that capability to the desktop platform.

For the Xeon processor line, Intel will introduce a system architecture called dual-independent buses (DIB). The dual-bus design (see Figure 3) is similar to the bus structure of AMD's original server processor, the Athlon MP. The Athlon MP was limited to dual-processor systems because the chip set could support only enough pins for two point-to-point processor connections.

Beyond two-processor systems, Intel mixes the DIB with its traditional shared-bus architecture (see Figure 4). The Paxville processor bus interface could run at 667MHz or 800MHz, because of the processor's integrated bus interface, significantly faster than the 533MHz of today's quad-processor XeonMP systems.

The disadvantage of putting the chip set in the middle of the processors is that coherency traffic must pass through the chip set, adding latency and a potential bottleneck. This was the scaling problem of the AMD Athlon MP processor. The advantage for the Intel Xeon lineup is that Intel can significantly improve bus bandwidth by reducing bus loading. The design trade-off will have different effects, depending on the application needs. For the moment, we believe the bandwidth improvement has the performance edge over the burden of the higher-coherency-traffic latency.

Yonah and Montecito Take Dual Core to Heart

The desktop Pentium and Xeon processor designs are more modular and provide quicker time-to-market solutions, but the Itanium and Pentium M versions of dual-core processing offer greater promise for future core integration. The Montecito and Yonah die integrate the two cores and use that close core proximity to build a more efficient processor design. The die layout of Montecito shows there is no clear, straight demarcation line between the two cores. The desktop Smithfield part looks like one could snap it in half, and each half is a complete and independent processor. Montecito uses Foxtan technology to control the power

of both cores. (See *MPR 2/28/05-02*, "ISSCC '05 Brings Out the Big Die.")

Although Intel did not reveal details of Yonah, there were a number of strong hints. It is likely the design will include a 2MB shared L2 cache and a unified bus interface. In a briefing, Mooly Eden, vice president of Intel's Mobility Group and GM of Intel's Mobile Platforms Group, said Yonah will not increase the power envelope from Dothan while adding the second core on the dual-core processor.

When asked about support for the EM64T 64-bit extension, Eden wouldn't make a commitment; instead, he talked about the power costs of adding 64-bit processing support. The message seems clear: Don't expect 64-bit support in Yonah. The processor will not include an integrated memory controller, but we expect Intel to deliver a mobile processor beyond Yonah that supports EM64T and has an integrated memory controller.

Intel's Israeli engineers are making a number of evolutionary improvements to the Banias core. The presentation slides showed that the two cores can go into independent power states but will need to maintain the same core frequency and voltage. The cores have to coordinate when to go to a lower voltage and frequency. We expect the cores also need to maintain synchronization if both cores share a synchronous relationship with the L2 cache.

The cores themselves have been improved with the addition of SSE3, fused SSE/SSE2 micro-ops, improved floating-point performance, and an improved instruction decoder. The original Banias decoder could decode only one complex and two simple instructions per cycle. The Yonah instruction decoder can decode a greater number of simultaneous complex instructions, including multiple SSE instructions. Eden promised improved floating point performance in Yonah which will translate to better gaming performance.

The Montecito Itanium processor is on track to ship in late 2005, and Intel demonstrated a number of systems running. Intel also revealed the long-term roadmap for Itanium. After Montecito comes the Montvale processor, also in 90nm. Sometime in 2007, Intel plans to ship the Tukwila MP server processor; Tukwila contains four or more cores and is part of the "Richford" platform. Intel might use the MCM technology to put two or more dual-core processors into one package; Tukwila is probably a 65nm part. After

Price & Availability

The following table summarizes the key dual-core x64 processors and derivatives discussed at the Spring '05 IDF.

Processor Code Name	Smithfield	Presler	Cedar Mill	Dempsey	Paxville
Process	90nm	65nm	65	65nm	65nm
Die Size (mm ²)	~202*	110*	110*	110*	—
Core/die	2	1	1	1	2
Die/package	1	2	1	2	1
MPR die cost est	\$48	\$16	\$16	\$16	—
MPR die cost est/package	\$48	\$32	\$16	\$32	—
Cache/die	2x1M	2M	2M	2M	TBA
Cache/package	2M	4M	2M	4M	TBA
Market	Desktop	Desktop	Desktop	Server	MP Server
Availability	2Q05	1Q06	1Q06	2Q06	1H07*

*In-Stat estimate, TBA = To be announced.

Tukwila will come "Poulson," but no other information about it was revealed.

The Richford platform is the first unified platform that both Itanium and Xeon will share. The Xeon processor for the unified platform is code-named "Whitefield." The new platform will provide for processors with an integrated memory controller, and Intel will also integrate the new system interconnect in future IDFs. We have learned that the new bus is called CSI, an abbreviation we believe stands for "coherent scalable interface." This new system architecture may have been influenced by earlier work on the Alpha EV7 by numerous Alpha engineers who joined Intel after the Alpha EV8 was canceled, as well as by those engineers' EV8 work. (See *MPR 03/24/03-02*, "EV7 Stresses Memory Bandwidth.")

Many-Core Processors in the Future

Standard operating systems will have no problem with two to four cores, but Intel's roadmap will eventually go beyond a single-digit number of cores. When the number of cores exceeds four, programming complexity increases, but, so far, these are symmetric cores. As we get to tens of cores on a die, greater core specialization may offer performance and power benefits. In the Q&A following his keynote, Justin Rattner,





Figure 4. The quad-processor system with the Paxville processor. Intel was forced to integrate the two cores' bus interface in order to limit the loading on the processor bus.

Intel senior fellow and director of the corporate technology group, would not rule out the possibility that cores could be developed that were optimized for improved media, networking, or security performance. These specialized cores could provide more performance on certain tasks for equal or less power. More-complex operating-system schedulers will be required to handle nonsymmetric cores or a virtual machine layer that can recognize certain code sequences and

then direct those sequences to specific cores optimized for those code profiles.

Future workloads like data mining, photorealistic 3D, and others should be amenable to multiprocessing system design. While general-purpose software has to catch up, there is ample university research and programming expertise in specialized areas such as high-performance computing that can be used to lead us to this many-core nirvana. 

To subscribe to Microprocessor Report, phone 480.483.4441 or visit www.MDRonline.com