

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

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

BEST SERVERS OF 2004

Where Multicore Is the Norm

By Kevin Krewell {1/18/05-01}

If you're in the server processor business and you either did not ship a dual-core processor in 2004 or are promising one this year, you're out of the mainstream. The next few years of server design will involve more cores and higher integration, as companies have more transistors to

use in 90nm process generation in order to put more of the server system onto one die.

The x86 server processor market continued the transition to the AMD64/EM64T/x64 (AMD/Intel/Microsoft names, respectively) instruction set in 2004, started by AMD's Opteron and accelerated by Intel's Nocona. AMD was making progress with AMD64 (also previously called x86-64) instruction extensions, but it has been slowed waiting for Microsoft's production release of x64 Windows operating systems, which didn't happen in 2004 and is now promised for 1H05.

Last year's winner of the Analysts' Choice Award for best server processor of 2003 was AMD's Opteron processor, which continued to gain market traction and design wins in 2004. AMD still hasn't reeled in Dell Computer as a customer, but that doesn't mean Intel hasn't been sweating it out.

AMD appears to be ahead of Intel in the race for the first dual-core server processor shipment, having demonstrated a **dual-core Opteron** (Figure 1) before Intel's Montecito. The dual-core Opteron is a contender for our server processor of the year because it promises to be a very scalable system solution with a nice balance of power, die size, and system performance; it is eligible because it sampled in 2004.

Sun shipped the 130nm UltraSPARC IV and prepared the way for the improved 90nm US IV+ (see Figure 2). While Intel was widely lashed in the press for canceling the Tejas processor, Sun had its own detour in its processor roadmap.



The company canceled two significant CPU projects, Millennium and Gemini. To fill the gap the cancellation created between the launch of UltraSPARC IV+ and the "Rock" processor scheduled for 2008, Sun announced it would collaborate with SPARC partner and business competitor Fujitsu on systems under the product name Advanced Product Line (APL), which would bring to Sun systems based on the Fujitsu SPARC64 VI. Fujitsu

is well along in the design and has a working test chip of its SPARC 64 VI dual-core processor.




Figure 1. AMD's dual-core Opteron processor was sampled in 2004. Both cores on die are almost identical, except that only one core's north bridge is active and is shared by both cores.

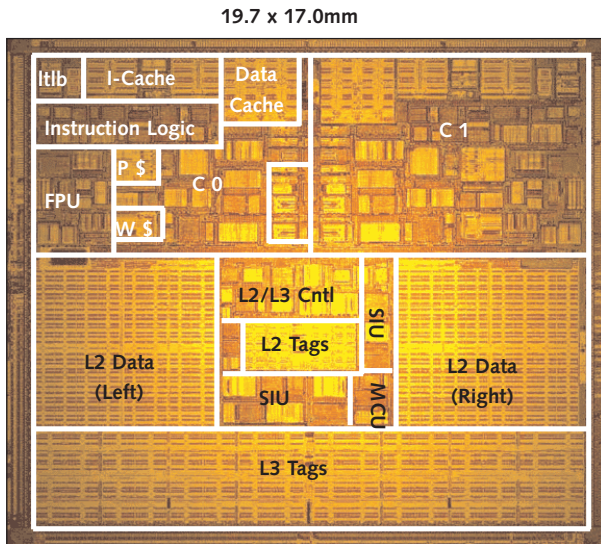


Figure 2. Die photo of the UltraSPARC IV+. The dual-core processor is built in Texas Instruments' 90nm bulk process, which has a 37nm gate length and nine layers of metal. The processor is built with 295 million transistors. Still, the die is 5% smaller than that of the US IV. The initial chips will be available at 1.8GHz and 1.65GHz, with a 2+GHz part coming later. The power is only 90W at 1.8GHz (1.1V), helped by using aggressive, dynamic power-management features. Sun had first silicon in 1H04.

One of the most aggressive and daring designs for multicore, multithreaded processing will be Sun Microsystems' **Niagara** processor, with eight quad-threaded cores on one die (Figure 3). The design uses a stripped-down UltraSPARC II core with minimal floating point. Design emphasis was placed on memory bandwidth and a reasonable amount of on-chip cache memory.

The Niagara design taped out in 2004, and first silicon is reportedly working very well. Sun has stated publicly that

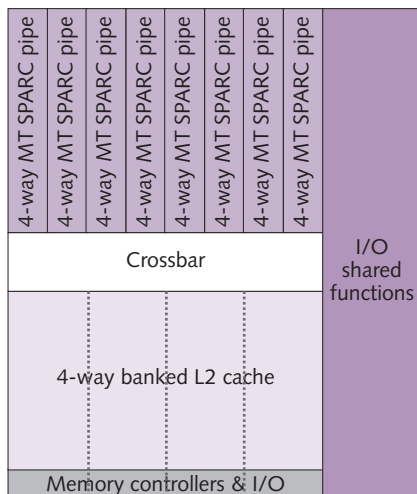


Figure 3. The Niagara processor has eight UltraSPARC II-like cores with a shared and banked L2 cache.

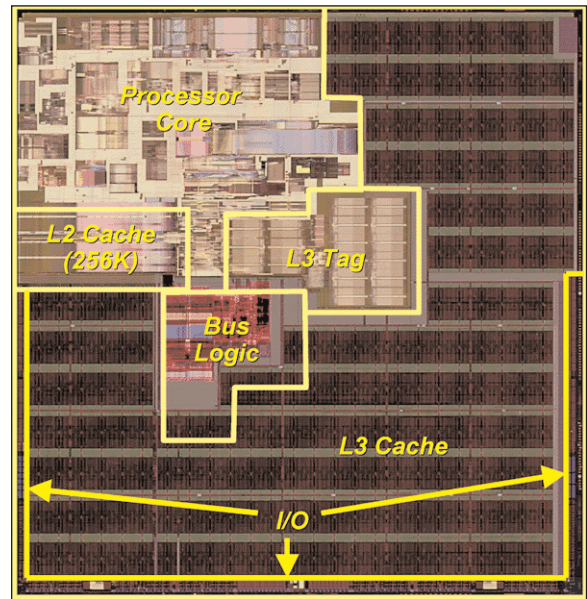


Figure 4. The Madison/9M version of the Itanium 2 is shown in the die photo. The integrated 9MB L3 cache can clearly be seen wrapping around on the right side and bottom of the die.

it will ship systems with Niagara in early 2006, but don't be surprised if Niagara-based systems start showing up considerably earlier. The challenge for Sun will be software that can efficiently manage 32 threads on one die. If Sun succeeds, expect to see other vendors scrambling to catch up in 2007.

Niagara is another potential server processor of the year because it takes multithreading and multicore design beyond all the other processors in this category. It's a bold design that attempts to get the right core and system on chip for particular server workloads. It could turn out to be a Xeon killer if Sun's Niagara with the Solaris operating system can deliver the promise of Throughput Computing.

In 2004, IBM shipped the **Power5** in systems and the **PowerPC970FX** in Apple Xserver rack servers. The Power5 is a dual-core, dual-threaded processor that is near the top in many system benchmarks. The Power Architecture was also well represented in the latest Top500 supercomputer list (www.top500.org) with the BlueGene/L beta system taking the number one spot.

Intel's Xeon processor made modest progress, as the 90nm XeonMP hasn't shipped yet, and the volume Xeon processors were hampered by the same power and frequency limits as the Pentium 4. The Itanium line saw the production release of the **Madison** core with 9MB of L3 cache (Figure 4).

Intel's Itanium processors continue to improve year after year. In 2004, Intel introduced another version of the 130nm Madison die with 9MB of L3 cache. The L3 cache is also protected with error correction code (ECC) logic. This Itanium 2 processor has 592 million transistors and a 432mm² die size, making it the largest server die in full production today. But

MPR 2004 Servers News in a Nutshell

The SPARC architecture made a comeback in 2004, at least in roadmap news if not SPEC scores. Early in the year, Sun released the code-name and little details of its next major project after Niagara, which will combine throughput computing with higher single-thread performance (*MPR 3/8/04-02*, "Sun Rolls Forward With Rock"). Soon after that announcement, Sun announced the cancellation of the Millennium and Gemini projects. To fill the gap between the launch of UltraSPARC IV+ and Rock, Sun and Fujitsu would collaborate on the Advanced Product Line (APL) (*MPR 8/02/04-02*, "SPARC's New Roadmap"). At FPF04, Sun and Fujitsu updated the roadmap for SPARC (*MPR 10/25/04-02*, "SPARC Turns 90nm"). Sun continues to dole out details on Niagara (*MPR 9/13/04-02*, "Sun's Niagara Pours on the Cores").

At IDF Spring '04, Intel abruptly announced it would also adopt the same AMD64 instruction set (*MPR 3/15/04-01*, "Intel Addresses the 64-Bit Question"). Intel apparently had to reverse-engineer the instruction-set extension from AMD public documents (*MPR 3/29/04-01*, "AMD and Intel Harmonize on 64" and *MPR 7/19/04-01*, "A Tale of Two Instructions"), setting the stage for a leader/follower role reversal with AMD. Intel was now forced to copy AMD instruction-set extensions, whereas AMD usually copies Intel's MMX and SSEn extensions (where n=1,2,3,...).

Intel's processor roadmap detoured in 2004 with cancellation of Tejas and its Xeon sibling Jayhawk. Instead of pushing higher clock frequencies with ever-increasing

power requirements, Intel chose to accelerate its plans for dual-core processing (*MPR 5/31/04-02*, "Intel's PC Roadmap Sees Double").

Soon after Intel's announcement, AMD announced it intended to ship a dual-core Opteron in mid-2005 and that it had completed tape-out (*MPR 7/06/04-01*, "AMD vs. Intel in Dual-Core Duel"). The company demonstrated its first working silicon just prior to Intel's Fall IDF. At FPF04, AMD also revealed some additional details of the dual-core Opteron (*MPR 10/11/04-02*, "Double Your Opterons; Double Your Fun").

Itanium architecture continued to make steady progress with production shipments of Madison/9M, first previewed earlier in the year at ISSCC (*MPR 1/05/04-02*, "ISSCC Promises Progress"). The next generation of Itanium is the dual-core, dual-threaded Montecito processor with a whopping 24MB of total L3 cache.

The Power architecture did move closer to volume servers with the Apple Xserve system (*MPR 1/20/04-02*, "Apple Debuts 90nm G5 in Xserve") and began production shipments of Power5 systems. The Xserve, with a 2.3GHz 970FX processor, was the basis of Virginia Tech's clustered supercomputer (System X) that is ranked 7th on the latest Top500 list of supercomputers.

IBM is looking to a new consortium, Power.org, to spread the Power Architecture and, in the process, develop more volume server designs (*MPR 12/27/04-02*, "Bringing Power to the People").

this is only a warm-up for the next version of the Itanium, code-named Montecito, which began sampling in 2004.

By far the most ambitious server-processor design in terms of transistor count is Intel's **Montecito**, using 1.72 billion transistors (Figure 5) in 90nm to create a dual-core processor with more than 26MB of on-die cache. Intel's manufacturing prowess continued to amaze with its ability to build gigantic chips with these enormous caches.

The chip promises to produce improvements in performance, due to faster clock speeds, larger caches, and coarse multithreading. As with most of this first generation of dual-core processors, the two cores share the same front-side bus as previous single-core processors; therefore, some of the other performance gains of the second core will be lost because of bandwidth conflict for the front-side bus. Even with all this logic integration, Intel managed to lower the power envelope as well. Still, we project a single-core SPEC2000_int performance of 1,700 at 1.7GHz. In a press briefing, Intel revealed three multicore Itanium processor code-names after Montecito: Montvale, Tukwila, and Dimona.

Montecito will likely be the highest-performing mainframe-class processor in 2005 and the highest-scoring

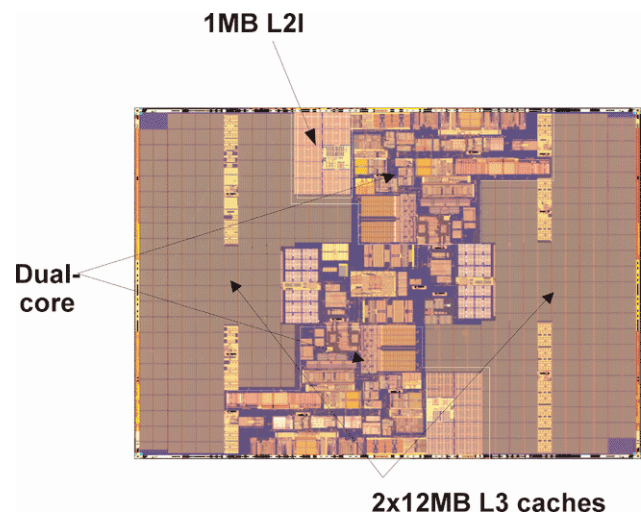


Figure 5. With 1.72 billion transistors, the Montecito die has three times the number of transistors as the Madison/9M it will replace and is fabricated in Intel's 90nm process. Montecito has two multithreaded cores. Each core is supported by a 1MB L2 cache and a 12MB L3 cache. To provide reliable operation with such large on-die caches, Montecito has soft error detection and correction.

processor on SPEC2000_fp. Despite all the criticisms and the relatively modest sales to date, Intel continues to show unwavering support for the architecture. The chip is such a massive die (although Intel has not officially revealed the die size, it's likely to be 500mm² or larger), it may remind some of the Death Star from the *Star Wars* movie. Montecito won't vaporize the competition, however, because, despite the large caches, it is still hobbled by an antiquated off-chip memory controller and system-bus design. Nevertheless, by late 2005, it will still have become the standard to judge against in "big-iron" systems.

And the Winner Is...

This year the decision came down to AMD's dual-core Opteron, IBM's Power5, Intel's Montecito, and Sun's Niagara processors. All are multicore processors, and all but Opteron support multithreading. Power5 is the most proven at this time, as it shipped in systems in 2004, whereas Opteron, Montecito, and Niagara are still in early sample stages. Picking the winner in this group is much like trying to pick the car of the year when the contestants are a big 600HP luxury car, a fuel-efficient hybrid, a high-powered exotic, and an SUV. Each offers different trade-offs of power, complexity, and cost.

IBM's Power5 is an impressive processor, but it's expensive, and it doesn't top the Itanium in SPEC performance.

It shipped in 2004, but we've been familiar with the design for a number of years.

AMD's Opteron should offer a compelling price/performance solution for many work loads and has the potential to scale beyond the eight-way with help from chips like Horus from Newisys. But we still don't know the clock frequency at which it will ship and how performance scales compared with the single-core Opterons.

The Niagara processor is a challenge to the industry to design systems for high-throughput workloads, not just for number crunching. Details are still leaking out about the chip, but it's obviously a bold design, and the word we hear is that silicon evaluation is looking very good. Our only concern is how the 8-core/32-thread processor actually performs under typical workloads. It's just too early to tell.

When you look for an all-out go-for-broke design that includes as many transistors as possible, however, the Montecito processor is the big kahuna. There is no question that, in terms of transistor count and die size, this is the most ambitious chip of the group. Intel crammed it with two cores with even more cache memory than the single-core Madison/9M, using a record-breaking 1.72 *billion* transistors, and it will have a lower power envelope than its predecessor. That is why the **Intel Montecito** processor is the winner of our *MPR* Analysts' Choice Award for **Best Server Processor of 2004**. ♦

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