

BEST NEW TECHNOLOGY: POWER4

IBM's Chip Multiprocessor Is Analysts' Choice for Technology Award

By Keith Diefendorff {2/7/00-01}

Bringing together the most awesome collection of microprocessor technologies that we have ever seen, POWER4 has been selected by the MDR analyst staff as the **Best New Microprocessor Technology** disclosed in 1999. In recognition of this achievement, the first

annual *Microprocessor Report* Technology Award was given to IBM, beating out five other nominees: Compaq's Alpha 21464, HAL's SPARC64 V, HP and Intel's IA-64 architecture, Sony and Toshiba's Emotion Engine and Graphics Synthesizer, and Sun's MAJC.

Just to make it onto this list of finalists, a company must have created an impressive new technology. Selecting a winner from this group was excruciatingly difficult, as each candidate deserved of the award in its own right. In fact, the race for first place was so tight we also awarded an **Honorable Mention** to Compaq for the Alpha 21464.

The "Unlimited-Class" Award

At our dinner meeting on January 27, MDR analysts presented four **Analysts' Choice Awards** (see sidebar) in the categories of server/workstation processors, PC processors, embedded processors, and 3D accelerators. Nominations for these awards were restricted to processors that were either announced in 1999 or whose announcement was imminent.

Recognizing that not all interesting microprocessor technologies fit neatly within these categories or met the 1999-announcement requirement, however, we decided to also present an unlimited-class award. This award, which we call the *Microprocessor Report* Technology Award, honors the one microprocessor, or microprocessor technology, that the MPR analysts believe has the most potential to improve the performance of the systems in which it will be deployed. There's no time limit on when a company must deploy the

technology, except that the company must disclose enough information for us to make an evaluation. Nor is the award limited to any particular class of processor or type of technology; it could be for a server, PC, or embedded processor, and it could be based on architecture, microarchitecture, implementation, IC-process, or any other innovation that we deem significant.



Innovation Continues Unabated

Despite pundits who claim that all the good ideas in computer architecture have been exhausted, 1999 showed no evidence that the pace of microprocessor innovation is slowing. Simultaneous multithreading, vector units, SIMD, VLIW, trace caches, superspeculation, and chip multiprocessing, although not brand new ideas, have now been refined to the point that they are ready for prime time. Enabled by IC process improvements and huge transistor budgets, these and other new techniques are now being used to build spectacular processors, and we expect the pace of innovation to continue at an increasing rate as far as we can see into the future.

Our first nomination for this year's *Microprocessor Report* Technology Award went to the Compaq Alpha 21464—for its adoption of simultaneous multithreading, or SMT. This clever idea originated with Susan Eggers and Hank Levy at the University of Washington and spread

Microprocessor Report Analysts' Choice Awards

Best Server/Workstation Processor (see [MPR 12/27/99-01](#)):

- *Winner:* Intel Pentium III Xeon (see [MPR 3/29/99-02](#))
- Compaq Alpha 21264 (see [MPR 10/28/96-02](#))
- HP PA-8500 (see [MPR 11/17/97-05](#))
- IBM POWER3/Pulsar (see [MPR 11/17/97-07](#))

Best PC Processor (see [MPR 1/17/00-05](#)):

- *Winner:* AMD Athlon (see [MPR 8/23/99-01](#))
- Intel Coppermine (see [MPR 10/25/99-01](#))
- Motorola G4 (see [MPR 11/16/98-04](#))

Best Embedded Processor (see [MPR 1/17/00-01](#)):

- *Winner:* Sony/Toshiba Emotion Engine (see [MPR 4/19/99-01](#))
- ARC Cores V3 Configurable Core (see [MPR 5/31/99-04](#))
- C-Port C-5 (see [MPR 10/6/99-en](#))
- IBM PowerPC 405GP (see [MPR 7/12/99-03](#))
- Intel/Level One IXP1200 Network Processor (see [MPR 9/13/99-01](#))
- Motorola DSP56690 (see [MPR 12/6/99-04](#))
- Tensilica Xtensa Configurable Core (see [MPR 3/8/99-02](#))

Best 3D-Graphics Accelerator (see [MPR 1/17/00-06](#)):

- *Winner:* NVIDIA GeForce 256 (see [MPR 9/13/99-msb](#))
- 3dfx Voodoo 3 (see [MPR 4/19/99-05](#))
- ATI Rage Fury MAXX (see [MPR 9/14/98-04](#))
- Matrox G400 (see [MPR 4/19/99-05](#))

Best New Technology:

- *Winner:* IBM POWER4 (see [MPR 10/6/99-02](#))
- *Honorable Mention:* Compaq Alpha 21464 (see [MPR 12/6/99-01](#))
- HAL SPARC64 V (see [MPR 11/15/99-01](#))
- HP/Intel IA-64 architecture (see [MPR 10/6/99-01](#))
- Sony/Toshiba Emotion Engine and Graphics Synthesizer (see [MPR 4/19/99-01](#))
- Sun MAJC (see [MPR 10/25/99-04](#))

from there to the Alpha design group through the persistent efforts of Compaq researcher Joel Emer. The concept is head-bangingly simple: exploit thread-level parallelism on basically instruction-level-parallel hardware by interleaving the execution of several independent program threads onto a common pool of execution units. The beauty of this idea lies in the fact that existing out-of-order superscalar processors have most of the hardware necessary to make SMT work, allowing it to be added at little additional cost. Using this trick, the Alpha team believes that, for a given

number of transistors, they can achieve much higher execution-unit utilization, and therefore higher throughput, than can be realized on traditional single-thread superscalar processors.

We concur; SMT is elegant in its simplicity and is likely to be so effective that many high-performance superscalar processors in the future will adopt this technique. One possible downside of SMT is that it's not compatible with VLIW, since it depends on the very superscalar instruction-issue hardware that VLIW machines were invented to avoid. Thus, with VLIW gaining popularity in the form of IA-64, MAJC, and a multitude of media processors and DSPs, the scope of SMT technology could be limited. On the other hand, if SMT works as well as advocates proclaim, SMT could render VLIW unnecessary. Since this conflict may not be resolved strictly on the basis of technical merit, we will have to wait to see how it plays out.

We had one additional concern with SMT: the complete picture of its complexity and its impact on frequency (or pipeline length) are still not well understood. Clearly, for example, a processor needs more registers and cache to support multiple threads; such additional resources could have an impact on performance, the net effect of which is unknown at this time. It was primarily this uncertainty that knocked SMT and the 21464 out of first place for our Technology Award.

Unbelievable Horsepower in a Kid's Game

For packing an astounding level of performance into a sub-\$400 game console, we nominated Sony and Toshiba for their Emotion Engine and its companion graphics chip, the Graphics Synthesizer. These chips, which will be at the heart of Sony's next-generation PlayStation 2 game console, sport an array of processing elements unparalleled by any PC processor in existence today: a 300-MHz 64-bit superscalar MIPS core, 128-bit SIMD extensions to the core, dual floating-point vector units, 10-channel DMA, an MPEG-2 decoder, a two-channel DRDRAM memory interface, and 16 parallel pixel processors that work into an embedded-DRAM frame buffer.

The two vector units are the most fascinating elements in the Emotion Engine. Each unit has five fully pipelined floating-point multiply-add units that can chew through physics models at a rate of 6.2 GFLOPS and pump out 3D polygons at the phenomenal rate of 66 million polygons per second. The Graphics Synthesizer renders these polygons at a rate of 2.4 billion pixels per second, using 48 GB/s of bandwidth to its on-chip frame buffer.

But perhaps the most amazing thing about the EE and GS is that together these parts occupy more than 500 mm² of silicon in a 0.25-micron process (EE = 240 mm², GS = 279 mm²); by comparison, a 0.25-micron Pentium III uses only 128 mm². Even though Sony and Toshiba will likely shrink these parts to 0.18 micron for high-volume production, it is not clear how the duo can produce such

large parts at costs consistent with an inexpensive game console—even if the console is subsidized by game-software revenue. It is clear, however, that the companies are intent on producing both chips inexpensively and in high volume. Two full fabs are under construction for the sole purpose of producing these parts.

VLIW Rears Its Formerly Ugly Head

Reworking an old idea that was once considered to have insurmountable problems for general-purpose computing, two of our nominees—HP/Intel's IA-64 and Sun's MAJC—have resuscitated very-long-instruction-word (VLIW) architecture.

After years of research and work on the problem, HP and Intel are on the verge of hatching their first IA-64-based processor: Itanium (née Merced). For IA-64, the HP/Intel team morphed VLIW into EPIC architecture in an attempt to preserve the good qualities of VLIW while eliminating the bad.

Despite considerable skepticism about this approach in the architecture community, and despite historically unimpressive results in increasing the average number of instructions that processors can execute per cycle (IPC), HP and Intel have steadfastly stood by the notion that there is still much more performance to be gained from instruction-level parallelism (ILP). Furthermore, they are betting the farm that compiler technology can find and expose that parallelism. Since the compiler performs much of the work that out-of-order hardware performs in superscalar processors, it's probably the most critical piece of the IA-64 strategy. We were not able to give the **Technology Award** to IA-64 because we have not yet seen enough evidence to convince us that today's compilers are up to this task.

If HP and Intel are right, however, IA-64 processors will deliver single-thread performance higher than that delivered by out-of-order superscalar processors, and they will do so with hardware that is much simpler than is required by those machines. Indeed, the large number of Intel's OEM customers that have already announced support for IA-64 are counting on HP and Intel to be right. In fact, there is so much industry backing behind IA-64 that its technical merit may be moot.

Like HP and Intel, another of our nominees—Sun's MAJC—is using VLIW. Unlike HP and Intel, however, Sun is not nearly so confident in high ILP. Sun believes that attempts to exploit ILP beyond four-way will yield diminishing returns. This belief is reflected in Sun's four-way MAJC architecture, which is not extensible to arbitrarily long instruction bundles the way IA-64 is.

We nominated MAJC for the **Technology Award** on the basis of what seems to be an extremely sensible approach to parallelism. Rather than placing all its eggs in the ILP basket, Sun uses VLIW as a hardware-efficient way to exploit moderate ILP and then uses chip multiprocessing (CMP) to also exploit parallelism at the thread level. Thread-level parallelism (TLP) is an ideal approach for Sun

because Java is inherently multithreaded. Indeed, MAJC is designed as a Java-centric architecture, and MAJC chips exploit the natural parallelism among Java threads and then go aggressively after additional TLP by speculatively executing Java methods.

MAJC chips employ yet another technique to boost performance: a direct connection to memory. Given the critical relationship between memory latency and performance, and considering that memory continues to get further and further from the CPU in terms of processor cycles, we expect more processors to adopt this technique in the future.

HAL Takes Unconventional Conventional Approach

Proving that there are plenty of new ideas about how to gain more performance, HAL Computer Systems, like HP and Intel, is going after ILP. But HAL believes it can exploit the ILP present in most programs without resorting to the disruptive approach of switching to new VLIW architecture. And unlike Sun and Compaq, HAL believes the software world is not quite ready to take advantage of thread-level parallelism with techniques like SMT and CMP.

Instead, HAL is trying to improve performance using new extensions to conventional out-of-order superscalar processors. In particular, HAL will use a trace cache, which attempts to linearize code for better performance, and superspeculation, whereby the processor attempts to execute past memory dependencies even before they complete. But we didn't nominate HAL's SPARC64 V merely for its use of these techniques. Trace caches and superspeculation are potentially quite complex when piled on top of other superscalar control logic. We were impressed with HAL's implementation, which should allow the SPARC64 V to benefit from the new techniques without sacrificing much clock speed.

HAL's chip will be an eight-issue superscalar processor built to the standard SPARC V9 architecture specification. The chip will use a four-level cache hierarchy with three levels on chip (trace cache, L1, and L2). Unlike other processors with on-chip L2 caches, HAL's L2 is split between instructions and data. SPARC64 V will use 65 million transistors and require 380 mm² of silicon in Fujitsu's 0.17-micron six-layer-copper CS85 process.

A Simply Awesome Collection of Technology

Never before have we seen such an awesome collection of technologies brought together in one chip as IBM has done with POWER4. One POWER4 technology that we are especially enthusiastic about is chip multiprocessing. CMP chips not only exploit high-level parallelism that uniprocessors cannot touch, but they offer significant advantages in construction over similarly performing uniprocessors, such as shorter design time and higher clock rates. And with the high-speed memory sharing that is possible between processors on a single chip, CMPs can exploit finer-grain parallelism and be more efficient than traditional discrete symmetric multiprocessors (SMPs).

While multiprocessing is known to be effective on server applications, it may also become effective for PCs as multimedia grows in importance. As a result, we expect CMP to be used by many processor vendors in the not too distant future.

But POWER4 is not just about CMP. Both of POWER4's two cores are 64-bit, five-issue, superscalar processors that will operate at more than 1 GHz, making each one more powerful than any single CPU in existence today. And unlike most companies that just moan and complain about the problems of memory latency and bandwidth, IBM did something about them. POWER4's two cores share a large on-chip L2 cache with 100 GB/s of combined bandwidth. The chip also provides 45 GB/s of off-chip bandwidth to other POWER4 chips, memory, and I/O. These bandwidths are an order of magnitude higher than found on typical processors today. IBM used wave pipelining to allow POWER4's wide expansion bus to operate at 500 MHz over long distances with good signal integrity.

POWER4 will be deployed in a four-chip multichip module (MCM), providing an eight-way SMP system in a single package. The POWER4 MCM package is not your father's MCM. The package measures about 4.5 inches on a side; it is made of a low-*k* glass-ceramic substrate with copper interconnect layers and a whopping 5,200 I/O pads. The thick metal carrier shown in Figure 1 is required to evenly spread the 700 pounds of force needed to install the package onto its motherboard. The company probably owes the technology used in this impressive package to its long history in mainframes.




Figure 1. Four POWER4 chips packaged in a glass-ceramic MCM provide the complete processor complex for an eight-way SMP system. Each MCM is 4.5 inches on a side, has 5,200 I/O pads, and dissipates about a half kilowatt of power.

Each POWER4 chip implements 174 million transistors and 5,500 C4 solder-ball connections to the MCM. We estimate the power consumption of one chip will be around 125 watts, or about a half a kilowatt per MCM. Chips will be built in IBM's high-reliability CMOS-8S2 process, a process that uses 15% shorter gates than IBM's current copper CMOS-8S, which we have evaluated to be the most advanced IC process in production today. As if that weren't enough, IBM will fabricate the chip on a silicon-on-insulator (SOI) wafer, giving it roughly 25% faster logic than would be expected from the same process on a bulk substrate.

Although POWER4 may sound like a chip for the middle of the decade, it will be available much sooner. Vijay Lund, vice president of technology development and the IBM executive responsible for the POWER4 project, took the opportunity of his acceptance speech at the awards dinner to announce that first silicon had been received, as Figure 2 shows, and that it has already been successfully powered up. Although the company refused official comment, chief architect Jim Kahle, who also attended the dinner, could barely contain his jubilation when asked about the status of his baby. IBM says that POWER4 will ship in systems during the second half of next year (2H01).

On the basis of this impressive collection of technology, it was our sincere pleasure to present this year's *Microprocessor Report Technology Award* to IBM for POWER4. Our congratulations go to the entire team. ♦




Figure 2. Each POWER4 chip implements 174 million transistors, which occupy about 400 mm² in IBM's seven-layer-copper 0.18-micron CMOS-8S2SOI process.

To subscribe to *Microprocessor Report*, phone 408.328.3900 or visit www.MDRonline.com