VOLUME 13, NUMBER 1
JANUARY 25, 1999

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

Server Chips Slip; RISCs Regroup

Are Processors Too Complex or Is It the Curse of Merced?

by Linley Gwennap

Just over a year ago, IA-64 came out of its burrow for the first time and, in the bright lights of Microprocessor Forum, cast a shadow over the rest of the industry. More powerful than the traditional groundhog, IA-64's appearance plunged high-end CPU design efforts into a deep freeze that lasted not just six weeks but six to nine months. Nearly every high-end RISC project suffered a major delay in 1998; ironically, so did the first IA-64 processor, Merced.

There are two theories for this unusually severe set of delays. One is that with all of the RISC CPU designers applying for jobs at Intel and all of the Merced engineers seeking to escape, not enough work got done. The second states that modern microprocessors are getting too complicated, extending the design cycles far beyond those of earlier, simpler chips. We believe there is some validity to both.

Merced's eight-month slip gives the RISC vendors some breathing room in the server market. Compaq's purchase of Alpha and IBM's abandonment of Somerset indicate that both companies will focus their future processor designs almost entirely on high-end multiprocessor servers, aiming to position their processors above Pentium II Xeon and ultimately IA-64.

In the workstation market, Intel's impact is already being felt. Aggressive x86-based products from Compaq, HP, and others have seized about a third of the workstation market. The remaining RISC workstation vendors have seen their margins plummet due to the Intel incursion. Silicon Graphics recently succumbed to the inevitable, rolling out its own x86-based workstations.

To recognize the best and the worst of 1998, we bestow our seventh annual RISCie awards. As usual, we exclude all embedded RISC processors from these categories, but this year we are including Intel's workstation and server processors along with the traditional high-end RISC chips. Note that all performance comparisons are based on SPEC95 (base) unless otherwise indicated.

Compaq Strengthens Alpha Hand

The past year saw Compaq absorb the last remnants of Digital, including its Alpha processor and system designs. After some initial doubt, we now see this move as positive for Alpha, at least for the next few years. Compaq has stabilized a disintegrating design team and committed to using Alpha chips instead of IA-64 in future Tandem systems. We give Compaq the Mighty Mouse award for saving the day.

The Alpha 21264 suffered a second six-month slip during this transition, finally appearing in Compaq systems in November. Adding insult to injury, the performance of these systems was not what Digital had predicted, missing the mark by up to 30%. This shortfall leaves Alpha's performance lead threatened by HP's new PA-8500, as Figure 1 shows. But the 21264 is still quite an impressive part, and with HP's chip still sampling, Alpha

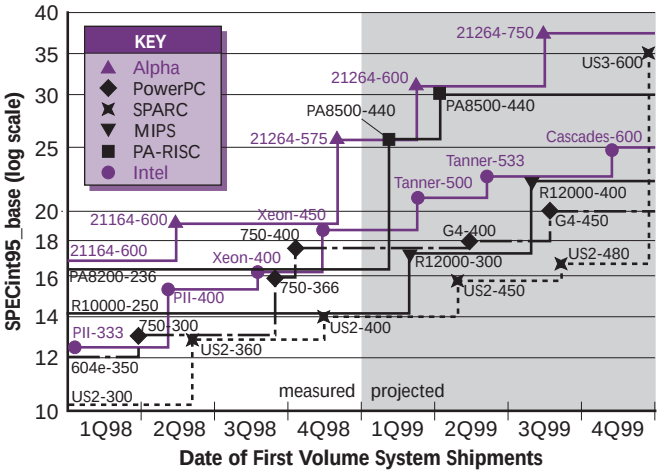

Figure 1. Although HP is poised to gain the integer performance lead, Alpha is likely to end 1999 back on top. Intel's roadmap shows strong integer performance, but the others are well behind. (Where two points are shown for the same processor, the vendor later released better results.) (Source: SPEC, MDR projections)

wins its sixth award as the World's Fastest Microprocessor (shipping).

At 575 MHz, the powerful 21264 delivers an impressive 26 SPECint95 and 41 SPECfp95. It gains this performance by combining complexity (Most Instruction Reordering), clock speed (faster than any non-Alpha CPU), and bandwidth (Most Memory Bandwidth at 2.7 Gbytes/s). Most impressive, the 21264 achieves this speed in an antiquated 0.35-micron process.

Compaq expects the 21264 to reach 750 MHz in a 0.25-micron process and to exceed 1 GHz in a 0.18-micron process. Getting to these advanced processes will take time, however, as Compaq lacks its own fab and relies on Samsung and a potentially recalcitrant Intel. We expect the 0.25-micron version in 2H99 and the 0.18-micron part in 1H00.

The best-case scenario for Alpha would be if Compaq chose to drive it into the PC market. The world's leading PC vendor, however, has shown no interest in this approach, instead seeing Alpha as a line extender for its x86 workstations and servers. As a result, shipments of the 21164PC never ignited, and Mitsubishi has dropped out of the Alpha chip market. Earning the King Sisyphus award, Samsung formed a new subsidiary, Alpha Processor Inc., to promote Alpha but so far continues to find no takers.

HP Fights Back With PA-8500

HP will challenge Compaq with its PA-8500, which will ship early next month at 440 MHz. HP rates the CPU at 30 int/50 fp; initial systems, however, use a slower bus and will not quite reach these scores. By the time HP ships better systems, Compaq plans to ship faster 21264 systems as well. At this point, the performance race is too close to call.

The PA-8500's unique design approach earns several awards, including Most On-Chip Cache (1.5M), Largest Die Size (475 mm^2), and Highest CPU Manufacturing Cost ($330, according to our cost model). The last is somewhat misleading, as other processors must add megabytes of expensive external cache for peak performance, whereas the PA-8500 uses only on-chip cache. Thus, at the system level, the HP chip actually costs less than many of its competitors.

As a result of this cost advantage, its high performance, and its trend-setting use of large on-die cache, we have given the PA-8500 our Editor's Choice award as best RISC processor. The chip gets extra credit for meeting its performance targets (due to a last-minute clock-speed boost) and coming within a month of its shipment goal.

Although HP continues to sell workstations using its low-end 7300LC, the aging chip delivers far less performance (integer or floating-point) than a Xeon-400. HP now offers Xeon-based workstations, but customers must switch from Unix to Windows NT to take advantage of this upgrade, earning HP a citation for Aiding and Abetting Microsoft.

After Merced's delay, HP announced a long-term plan to keep PA-RISC around during the now-extended transition to IA-64. There will be at least two, and possibly four, more PA-RISC processors, although all are likely to be based on variations of the existing PA-8x00 core. The next version, the PA-8600, isn't due until 2000; thus, after challenging for the performance lead in early 1999, HP is likely to fall behind over the course of the year, due to a strategy with the Fewest New Processor Introductions.

IBM Aims to Bounce Back

Last summer, the futility of trying to make PowerPC all things to all people finally hit home, with IBM and Motorola agreeing to separate their design efforts. IBM originated the POWER family and continued to design its own PowerPC-based server processors, even while putting Somerset-designed chips into its workstations and low-end servers. But the company has been losing share in both the workstation and server markets and is de-emphasizing its future workstation efforts. Instead, IBM's future processors are being designed for expensive multiprocessor servers.

Its current chip in this vein, Power3, began shipping last September. Unlike other CPU makers, IBM allowed several extra months to test its latest processor. Thus, the company could claim to be on schedule when it shipped the first Power3 systems a year and a half after the chip taped out. At 200 MHz, Power3 has the Slowest Clock Speed among its key competitors, relying instead on instruction parallelism and bandwidth for high performance. This combination is ineffective, however, on many applications and particularly on SPECint95, where Power3 scores just 12.5. As Figure 2 shows, the processor fares better on SPECfp95 and does particularly well on server applications with large data sets.

In 3Q98, IBM also began shipping a processor known as Northstar, which combines the PowerPC instruction set with extensions for AS/400 compatibility. The company

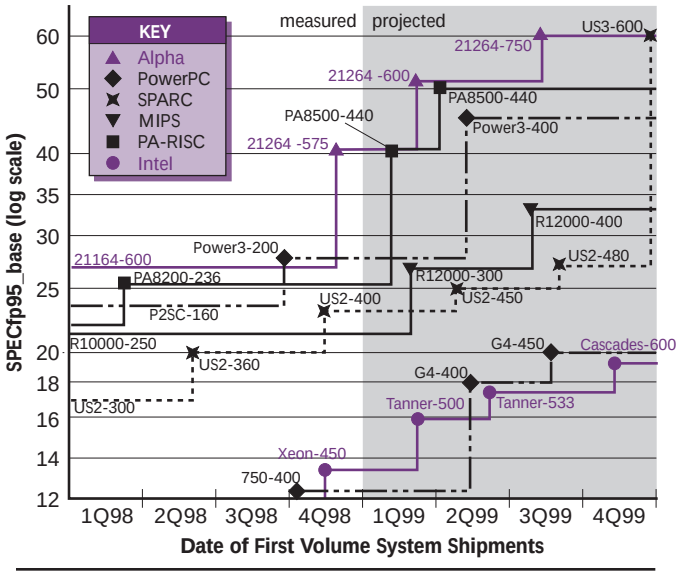

Figure 2. The gap between the floating-point haves and have-nots is large and growing. Sun hopes its UltraSparc-3 will vault to a performance lead in late 1999. (Source: SPEC, MDR projections)

hasn't released SPEC95 scores for Northstar; at 262 MHz it outperforms Power3 on some commercial applications but not on FP-intensive code. Too Big Blue earns a Redundantly Redundant citation for not merging these two projects.

As a vestige of its former strategy, IBM shipped a 400-MHz PowerPC 750 code-named Lonestar, mainly for Apple's use. Lonestar delivers far better integer performance than Power3 or Northstar with a die size of just 47 mm^2, making it the Smallest High-End Processor. As the first high-volume copper processor to reach the market, Lonestar also wins the award for Best Heavy Metal Performance.

With Apple returning from the dead, PowerPC systems tied traditional leader PA-RISC for Most RISC System Revenue, as Figure 3 shows. Most of this revenue, however, is garnered by IBM's RS/6000 and AS/400 lines. The future for these products lies in Power3 and Northstar. Like Captain Kirk, IBM is calling on engineering for more speed. Without changing the laws of physics, the company plans to double the clock speed of both chips by 2H99 and push them to 500 MHz or more in 2000.

For the long term, IBM has a radically new project code-named GigaProcessor. True to its name, GigaProcessor is targeted at 1 GHz or faster in a 0.15-micron process, with first shipments in 2001. The chip is also rumored to contain two independent CPUs, which could make IBM the First to Deploy a Multiprocessor Chip. If the chip meets its aggressive speed goals, it could turn IBM's performance problems around and be a strong competitor against IA-64.

Motorola Inherits Somerset, Macintosh

With IBM taking the high road, Motorola now owns the Somerset design center and all of its future designs, starting with the forthcoming G4. Although IBM will continue to sell G3 processors to Apple for as long as the Macintosh vendor wants, at some point (probably by the end of 2000) Apple will convert its purchases entirely to the G4, and thus to Motorola. Therefore, Motorola's PowerPC efforts are now aimed at supporting Apple's needs and at high-end embedded applications, such as networking and communications. These markets play to Motorola's historical strengths.

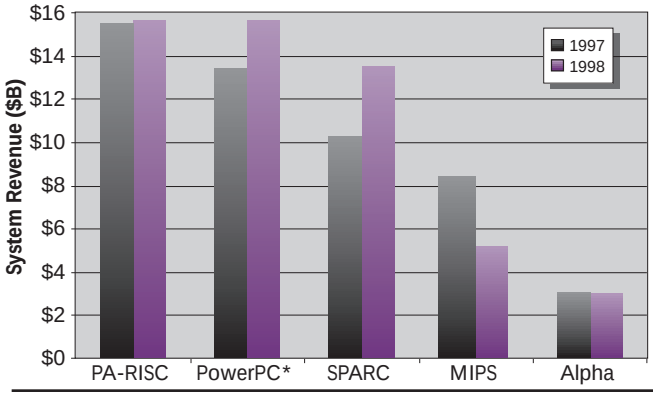

Figure 3. PowerPC system revenue matched PA-RISC's in 1998. SPARC surged while MIPS withered. Total RISC system revenue increased from 1997's $50.6 billion to $52.9 billion. *includes POWER (Source: *Inside the New Computer Industry*)

Despite having its own copper process, the company has chosen not to match IBM's Lonestar, instead waiting to deploy copper G4 processors in mid-1999. This earns Motorola the Best Heavy Metal Lipsync award. As Figure 1 shows, the G4 won't do much to improve PowerPC performance on SPECint95 and similar applications. The G4 core includes the new AltiVec extensions, which will deliver a bigger performance boost on many floating-point and multimedia applications.

Thus, although it will excel on some key applications, the G4 won't give Apple a leg up on Intel in performance. With its grand strategy of displacing Intel crushed, Motorola gets the If You Can't Beat 'Em, Join 'Em award for replacing its old "half the price" plan with more Intel-like prices.

Motorola has confirmed it is now shipping the long-lost PowerPC 620, which earns a Rip Van Winkle award. This processor was announced so long ago, its performance target was 225 SPECint92. It is so old, it was originally designed for 0.5-micron CMOS, although the current version is produced in a 0.35-micron process that boosts its clock speed to an exhilarating 180 MHz. Even with the process shrink, the die measures 250 mm^2, larger than any other Motorola CPU.

With 50 months between architectural disclosure and revenue shipments, the 620 sets a new all-time mark for Longest Time as Vaporware, smashing the previous record of 27 months held by Motorola's own 88110. Motorola attributes the delays to problems with the multiprocessor logic; Groupe Bull, the chip's only customer, had further problems getting its systems out the door.

Sun Designers Strain; Ross Razed

Already trailing its competitors at the start of 1998, Sun could manage only a 33% speed bump for its UltraSparc-2, which is now shipping at 400 MHz. As a result, SPARC racks up its sixth Tail-End Charlie award for laggard integer performance. Perhaps trying to cover this shortfall, the vendor reports SPECpeak results that exceed the corresponding SPECbase values by 24%, earning a citation for Stupid Compiler Tricks. The US-2 fares better in server performance, due to its multiprocessor scalability and memory bandwidth, which exceeds that of any processor except the 21264.

Sun has been gaining share in the server market for the past few years; as Figure 3 shows, SPARC had the Biggest Revenue Gain of 1998. Its workstations, however, have been vulnerable to attacks from Intel-based systems. A year ago, Sun revamped its low-cost workstations with the "Darwin" systems, which use the UltraSparc-2i processor. These systems are selling well despite being far slower than Xeon-based workstations on all but the most FP-intensive tasks.

To improve its performance standing, Sun's hopes lie in the forthcoming UltraSparc-3. Unfortunately, this project was badly affected by the industry deep freeze, slipping more than six months. US-3 is Sun's first processor to support instruction reordering, perhaps explaining the delays. If the

new design meets its aggressive performance goals and late-1999 shipment date, Sun could go from Worst to First, as Figure 2 shows. But US-3's tape-out has not yet occurred, and recent complex processors have taken 12–18 months to debug, making a mid-2000 debut more likely.

Erstwhile SPARC partner Fujitsu has shut down subsidiary Ross Technologies, the HyperSparc vendor, leaving us to ask What Would You Like on Your Tombstone? To compensate, Fujitsu strengthened ties with its other SPARC design firm, Hal Computer. Hal's Sparc64-III appeared a mere four months behind schedule but met its performance goals and is shipping in Fujitsu servers. At 15 int/25 fp, it delivers better performance than UltraSparc-2. Now at 275 MHz, the chip is slated to reach 300 MHz by mid-1999.

MIPS: Dead Man Walking

Silicon Graphics' decision to phase out its high-end MIPS line and spin off its embedded MIPS team has clearly crippled its ability to evolve its processor line. Despite incentives to keep the design team in place, the R12000 has slipped more than six months and is now expected to ship in early 1999. This lack of forward progress earns MIPS the Flatline award for 1998, as depicted in Figures 1 and 2, and undoubtedly contributed to the loss of revenue shown in Figure 3.

Even the 300-MHz part will do little more than fend off the feeble UltraSparc-2 in performance. The MIPS processor still has the Worst Memory Bandwidth, which won't be fixed in the R14000 until early 2000. In the interim, SGI plans to move the R12000 to NEC's next-generation process around midyear, bumping its clock speed to 400 MHz. (NEC will be the only R12000 fab, as R10000 foundry Toshiba has chosen not to pursue the dwindling high-end MIPS market.)

Silicon Graphics just announced Xeon workstations that replace systems using the uncompetitive R5000, but like HP's customers, SGI's must switch to Windows NT to stay at a low price point. The company will move to IA-64 in the long run—not too much longer, it hopes.

Intel Extends High-End Efforts With Xeon

Gaining momentum in the high-end market, Intel crystallized its efforts in August with its new Pentium II Xeon processors. Unlike Pentium Pro, the Xeon brand will be used exclusively for workstation and server processors that compete head-to-head with RISC offerings. While the "server" version with 2M of cache carries a steep list price of $3,692 (see MPR 1/25/99, p. 4), the "workstation" version with a 512K cache lists for just $824.

Intel stumbled in getting the first Xeons out the door, suffering from latent bugs in the CPU and chip sets as well as poor initial yield. The company claims to have all of these problems under control and is shipping Xeons in volume to an impressive array of system makers, including Compaq, Data General, Dell, HP, Micron, IBM, Sequent, and Silicon Graphics, giving Xeon the Biggest Bandwagon award.

Some of the companies just joining the Intel customer list should consider the example of Intergraph, which also happily dumped its in-house RISC architecture in favor of x86. But after getting tangled in a patent lawsuit, Intel made like Lorena Bobbitt and cut off Intergraph's processors; only an emergency judicial procedure made the workstation vendor fully functional again.

The Xeon processor has the advantage of strong integer performance and outstanding software support. Despite Intel's efforts, however, Xeon doesn't address the entire high-

Major RISC Events of 1998

Compaq announced plans to buy **Digital** (2/16/98, p. 4) and to purchase Alpha chips from Intel (5/11/98, p. 4). Compaq's Tandem subsidiary will use Alpha instead of IA-64 (10/5/98, p. 5). The 21264 finally appeared but missed its performance targets (9/14/98, p. 4). The 21364 will include a radical system design (10/26/98, p. 12).

Samsung acquired a license to design Alpha processors (3/9/98, p. 4) and formed a new subsidiary, Alpha Processor Inc. (6/22/98, p. 22). **AMD** negotiated with Compaq for an Alpha license of its own (5/11/98, p. 4).

Motorola pushed the PowerPC 750 to 300 MHz (3/30/98, p. 5) and later 366 MHz (8/24/98, p. 13). **IBM** used its copper process to build a 400-MHz 750 (9/14/98, p. 4). Motorola disclosed its AltiVec multimedia extensions (5/11/98, p. 1), to be deployed in the forthcoming G4 processor core (11/16/98, p. 17).

IBM turned over the Somerset Design Center and the G4 design to Motorola (6/22/98, p. 4). IBM shipped Power3 on schedule at 200 MHz (10/26/98, p. 4).

Sun began shipping its highly integrated UltraSparc-2i processor at 300 MHz (1/26/98, p. 4) and later at 333 MHz (5/11/98, p. 5). UltraSparc-2 jumped to 360 MHz (5/11/98, p. 5). UltraSparc-3 was delayed, but Sun extended its roadmap through 2001 (10/5/98, p. 15).

SPARC design house **Ross Technology** shut its doors (6/22/98, p. 5). **Fujitsu** announced systems using the Sparc64-III chip from **Hal Computer** (12/7/98, p. 4).

Silicon Graphics killed its next-generation MIPS processors in favor of Intel chips (4/20/98, p. 1). The R12000 was delayed (11/16/98, p. 4).

Hewlett-Packard made plans to ship the PA-8500 at 440 MHz (10/26/98, p. 4) and revealed its roadmap through the PA-8900 (11/16/98, p. 4).

Intel rolled out the first Xeon processors and chip sets (7/13/98, p. 1 and p. 11) and after some delays (9/14/98, p. 15) boosted their speed to 450 MHz (10/26/98, p. 4). Merced slipped to mid-2000 (6/22/98, p. 1). Intel disclosed Cascades, Foster, and McKinley (10/26/98, p. 16).

Intergraph gained access to Intel's high-end processors through a court ruling (4/20/98, p. 4) (5/11/98, p. 16). Intel predictably countersued (7/13/98, p. 5).

end market. It has the Worst Floating-Point Performance of any architecture, making it a poor choice for scientific applications. Some vendors sell systems with 32 or more Xeon processors, but these systems require extensive system logic to overcome the processors' limited bus bandwidth.

In the near term, Intel is likely to play to its strengths while ignoring markets where Xeon doesn't fit well. By continually leveraging its fastest PC processor cores into the Xeon line, Intel will rapidly improve clock speeds over the next two years. The progression starts with the 500-MHz Tanner (Pentium III Xeon), based on the Katmai core, in 1Q99 and will continue with Cascades, based on the Coppermine core, in 2H99. We expect Cascades to eventually exceed 700 MHz and its successor, Foster, to exceed 1 GHz in late 2000. But we don't expect much improvement in FP performance or bus bandwidth until Foster appears.

Merced: Hit or Myth?

IA-64 represents Intel's long-term strategy for crushing the RISC vendors. The first IA-64 chip, code-named Merced, was supposed to establish the new architecture's performance superiority when it shipped in 2H99. Now that the project has slipped to mid-2000, Intel has been backpedaling on claims of performance leadership. We won't be surprised if the first Merced systems ship in late 2000, but the chip will ultimately deliver strong performance, outrunning the 1-GHz 21264 on some benchmarks and pacing it on others.

Although it may not be the industry's fastest processor, Merced will still be deployed in a variety of servers and workstations from the aforementioned Xeon system vendors and others. We expect Merced will be modestly successful, due to this industry backing. Intel is counting on its second IA-64 processor, McKinley, to put the final nails in the RISC coffin, but that chip isn't due until late 2001.

The real question is whether any CPU vendor will deliver a strong, on-time product in the interim. Digital and HP have led the industry in performance for the past several years, but many key Alpha architects did not make the transition to Compaq, and many of HP's key designers are now focusing on IA-64. Sun and IBM have aggressive plans to turn around their performance shortcomings, but neither has a convincing track record. SGI's processor line is on life support, and Intel has already blown its chance to make a good first impression with Merced. McKinley looks great on slides, but it's still two years from tape-out.

Despite their best efforts, vendors are taking longer to design and test complex new microprocessors. As Intel has found, simply throwing more engineers at the problem is not the right strategy. Ultimately, microarchitecture itself must change to focus on ease of design along with cost and performance. Until that happens, we may see CPU performance fall behind Moore's Law. M