IBM TRIMS POWER4, ADDS ALTIVEC

64-Bit PowerPC 970 Targets Entry-Level Servers and Desktops

By Tom R. Halfhill {10/28/02-02}

Rarely does a downsized product raise expectations for high performance. But by trimming down the awesome 64-bit Power4 server processor and adding AltiVec media extensions, IBM has created an impressive and affordable PowerPC chip for smaller

servers, graphics workstations, and desktop computers. IBM disclosed the first technical details of the new PowerPC 970 at the recent **Microprocessor Forum 2002** in a presentation by Peter Sandon, the senior PowerPC architect at IBM Microelectronics. Sandon confirmed that the widely rumored chip will be a lower-cost 64-bit incarnation of the Power4, a radical design for high-end servers (see *MPR 11/20/00-03*, "IBM's Power4 Unveiling Continues"). Nobody at IBM would confirm rumors that a leading customer for the PowerPC 970 is Apple—and Apple is even more tight-lipped. Nevertheless, the 970 is such an obvious improvement over today's Motorola G4-family PowerPC chips that it's hard to imagine Apple using anything else in its top-of-the-line desktop Macs and servers.

Indeed, the 970 seems tailor-made for professional publishing and media-processing applications. It has 64-bit datapaths and memory addressing, yet it's natively compatible with 32-bit PowerPC software. It has a much deeper pipeline than G4-family processors, so it's in a better position to compete with a rival architecture that's marketed largely on the basis of clock speed. It has tremendous bus bandwidth, pushing effective bus rates to 900MHz and peak data bandwidth to 6.4GB/s. And it's the first IBM PowerPC chip with AltiVec extensions—the single-instruction, multiple-data (SIMD) operations that accelerate data-intensive processing tasks.

Downsized, Not Neutered

IBM's challenge was to design an affordable version of the Power4 without losing the basic traits that make it so impressive to begin with. Clearly, something had to go. The 170-million-transistor Power4 integrates two 1.3GHz processor cores on a single die with 1.5MB of shared L2 cache, an L3 cache controller, and a chip-to-chip fabric controller that lets IBM pack four Power4 chips (totaling eight processors and 680 million transistors) into a single multichip module (MCM). Four MCMs can be linked together without glue logic. The resulting 32-way multiprocessor subsystem has more than 20,000 I/O pads and dissipates about 2kW. It also requires 700 pounds of insertion force: to plug it into a socket without using special tools, you'd have to get two beefy NFL linemen to stomp on it.

Obviously, the Power4 is overkill for a desktop computer or local server, especially one that doesn't have to run Windows. So IBM removed one of the two processor cores, the L3 cache controller, and the complex chip-to-chip fabric controller. For vector processing, IBM added the AltiVec extensions, which IBM developed jointly with Motorola and which are currently available only in Motorola's G4 and G4+ PowerPC chips.

The physical characteristics tell the story. The Power4 has 170 million transistors and a 415mm^2 die, and it comes packaged in an 85 × 85mm MCM with 5,184 pins. The PowerPC 970 has 52 million transistors and a 118mm^2 die;

it comes packaged in a 25 × 25mm ceramic ball-grid array with 576 pins. These are enormous reductions. Although IBM hasn't announced pricing, the crash diet should make the 970 affordable enough for entry-level servers and high-end desktop systems. The 970's die is 40% larger than AMD's Thoroughbred-core Athlon XP ($84mm^2$), but it's 10% smaller than Intel's NorthwoodX-core Pentium 4 ($131mm^2$)—remarkable for a 64-bit processor that bridges the performance desktop and server markets.

What's left after all that downsizing? A powerful but more-conventional single-core microprocessor without the costly MCM packaging and neighborhood-brownout power requirements. The 970 is still a 64-bit, out-of-order-execution, five-way superscalar machine with even deeper pipelines, dynamic branch prediction, exceptional bus bandwidth, and enough coherency logic to keep it an interesting choice for SMP systems. Almost all these features set it apart from existing G4 and G4+ chips.

The deeper pipeline is especially welcome because it could allow the 970 to reduce the growing clock-frequency gap between today's PowerPC chips and the speedy x86 competition. G4 processors and their G4+ derivatives are stunted by short pipelines of only five or seven stages, compared with at least twenty stages in the superpipelined Pentium 4. That some G4+ chips manage to run at 1.25GHz is a testament to careful circuit design and good fabrication process. Imagine what might be possible with the 970, whose pipelines are 16 stages deep for integer operations and up to 25 stages deep for SIMD operations. Table 1 compares some vital statistics for these processors.

For now, IBM is being conservative. The 970 will debut at clock rates from 1.2GHz to 1.8GHz when fabricated in a 0.13μm CMOS process with silicon-on-insulator (SOI) transistors and eight layers of copper interconnects. IBM estimates power consumption at 19W (typical) at 1.1V for a 1.2GHz part, and 42W (typical) at 1.3V for a 1.8GHz part. Sampling will begin in 2Q03, with volume production scheduled for 2H03.

Another Shot at 64 Bits

Contrary to some reports, the 970 isn't the first 64-bit Power-PC. In fact, a 64-bit PowerPC was planned from the start, when IBM, Motorola, and Apple began creating the PowerPC architecture at the Somerset design center in Austin in the early 1990s (*MPR 7/24/91*, "Apple/IBM Deal Catapults RS/6000 to Prominence"). At that time, the PowerPC alliance promised to deliver three 32-bit processors—the 601, 603, and 604—and one 64-bit implementation, the 620.

All four chips eventually reached the market, but only the 32-bit processors succeeded. The ill-fated 620 first appeared on the PowerPC roadmap in 1991 and was first described at Microprocessor Forum in 1994 (see *MPR 10/24/94-02*, "620 Fills Out PowerPC Product Line"), but it didn't ship until 1998. By then, it had grown so complex it was uneconomical. At $250mm^2$, it was Motorola's largest slab of silicon. The 620 never made it into a Mac and soon vanished.

Indeed, the 620 was so late that a few other 64-bit Power-PC chips beat it to market. In 1995, IBM shipped the A10 and A30 processors, a pair of AS/400 chips modified to support the 64-bit PowerPC architecture (see *MPR 7/31/95-04*, "IBM Creates PowerPC Processors for AS/400"). IBM followed those introductions with the 64-bit PowerPC-compatible RS64 in 1997 and Power3 in 1998 (see *MPR 11/17/97-07*, "IBM's Power3 to Replace P2SC").

	IBM PowerPC 970	IBM Power4	Motorola G4+	Motorola G4	IBM/Motorola G3
Virtual address range	64 bits	64 bits	32 bits	32 bits	32 bits
Real address range	42 bits	42 bits	36 bits	32 bits	32 bits
Scalar datapath width	64 bits	64 bits	32 bits	32 bits	32 bits
CPU cores per die	1	2	1	1	1
Superscalar execution	4 + 1 branch	4 + 1 branch	3 + 1 branch	2 + 1 branch	2 + 1 branch
Pipeline depth (int)	16 stages	12 stages	7 stages	4 stages	4 stages
AltiVec extensions	Yes	No	Yes	Yes	No
FPUs	2 + AltiVec	4	1 + AltiVec	1 + AltiVec	1
L1 cache I/D (ways)	64K/32K (DM)	2 x 64K/2 x 32K	32K/32K (8)	32K/32K (8)	32K/32K (8)
L2 cache (internal)	512K	1.5MB	256K	None	None
Core frequency (max)	1.8GHz	1.3GHz	1.25GHz	550MHz	700MHz
FSB frequency	450MHz	433MHz	133MHz	133MHz	100MHz
FSB effective bit rate	900MHz	433MHz	133MHz	133MHz	100MHz
FSB width	2 x 32 bits	2 x 128 bits	64 bits	64 bits	64 bits
FSB data bandwidth	2 x 3.2GB/s	2 x 6.9GB/s	1.0GB/s	1.0GB/s	800MB/s
Transistors	52 million	170 million	33 million	10.5 million	6.5 million
IC process	0.13μm Cu + SOI	0.18μm Cu + SOI	0.18μm Cu + SOI	0.18μm Cu	0.22μm Al
Die size	$118mm^2$	$415mm^2$	$106mm^2$	$83mm^2$	$47mm^2$
Voltage (core)	1.3V (1.8GHz)	1.5V (1.3GHz)	1.6V (1.25GHz)	1.8V (550MHz)	2V (600MHz)
Power (typical)	42W @ 1.8GHz	~125W @ 1.3GHz	21.3W @ 1.0GHz	5.3W @ 500MHz	4W @ 400MHz
Production	2H03	Now	Now	Now	Now

Table 1. The PowerPC 970 is clearly in a different class than existing G4+, G4, and G3 PowerPC chips. Its deeper pipelines and much faster front-side bus (FSB) fix the most serious shortcomings of today's PowerPCs.

The PowerPC architecture has always supported 64-bit memory addressing. IBM made sure the architecture could switch between 32- and 64-bit modes to support AIX, IBM's version of Unix. Building on this foundation, the 970 supports a flat address space of 64 bits for effective addresses and a 42-bit real address range. A fully associative segment-lookaside buffer (SLB) can cache up to 64 entries from the off-chip segment table.

When running 32-bit programs, the processor clears the upper 32 bits of all effective addresses and uses the first 16 entries of the SLB as on-chip segment registers, reducing the number of off-chip table accesses. There's also a 1,024-entry TLB (256 × 4 ways) and a pair of effective-to-real address tables (ERAT) that allow the processor to quickly look up memory addresses without translation. Each ERAT holds 128 entries.

All integer and floating-point datapaths and registers in the 970 are 64 bits wide. It has the normal RISC complement of 32 general-purpose registers for integer operations and 32 floating-point registers. The AltiVec unit, which we describe below, has 128-bit-wide datapaths and 32 registers of the same width. In addition to these programmer-visible architectural registers, each of the 970's three register files also has 48 rename registers for pending operands.

As a more-efficient 64-bit implementation than the old 620, the 970 has a much better chance of bringing 64-bit processing to PowerPC desktop computers and small servers. Although it's a coincidence that the 970 will appear in systems within a few months of AMD's x86-64 Hammer processors, it will be interesting to see which 64-bit architecture wins more success in the marketplace. AMD has the greater momentum of the x86, but it faces the obstacle of major server customers already committed to Intel products. IBM has one top-tier server customer virtually guaranteed (IBM) and another potential customer with few other options (Apple), but it faces the prospect of a smaller overall market.

Five-Way Superscalar Execution

One important characteristic the 970 inherits unabridged from the Power4 is wide superscalar execution. As Figure 1 shows, the 970 has ten function units: two ALUs, two FPUs, two SIMD units, two load/store units, a branch unit, and a condition-register unit.

The 970 can fetch up to eight instructions per clock cycle, dispatch up to five instructions per cycle to the function units, issue up to eight instructions per cycle from the function units, and retire results for up to five instructions per cycle. It's the results that count, so we regard the 970 as a five-way superscalar processor backed by extra resources in the execution stages.

In comparison, G4+ processors have two ALUs for simple integer instructions, one ALU for complex integer instructions, one FPU, two SIMD units, one load/store unit, and a branch unit. Judged solely by their complements of function units, a G4+ processor doesn't seem greatly outmatched by the 970. But a G4+ can complete only four instructions per cycle (if one is a branch), compared with the 970's five instructions per cycle (four plus a branch). Earlier G4 processors can complete only three instructions per cycle (two plus a branch). The 970 approaches the practical limit of superscalar RISC execution with typical program code. It's doubtful that wider pipelining would significantly improve the instructions-per-cycle efficiency.

The additional execution resources and their queues allow the 970 to juggle an unusually large number of instructions in its pipelines simultaneously: more than 200 instructions can be "in flight" at any moment. IBM says those 200-plus instructions may include as many as 20 "dispatch groups" (five instructions per group) while 100-plus other instructions are in the various fetch, decode, and store stages. The large number of possible in-flight instructions gives the 970 much more scheduling flexibility than other PowerPC processors, which have a relatively limited ability to execute instructions out of order.

The depths of the 970's pipelines vary greatly, depending on the instruction type. All the pipelines begin with nine fetch and decode stages. This large number of initial stages is extraordinary for a classic RISC architecture. It's more reminiscent of today's x86 processors, which must overcome the
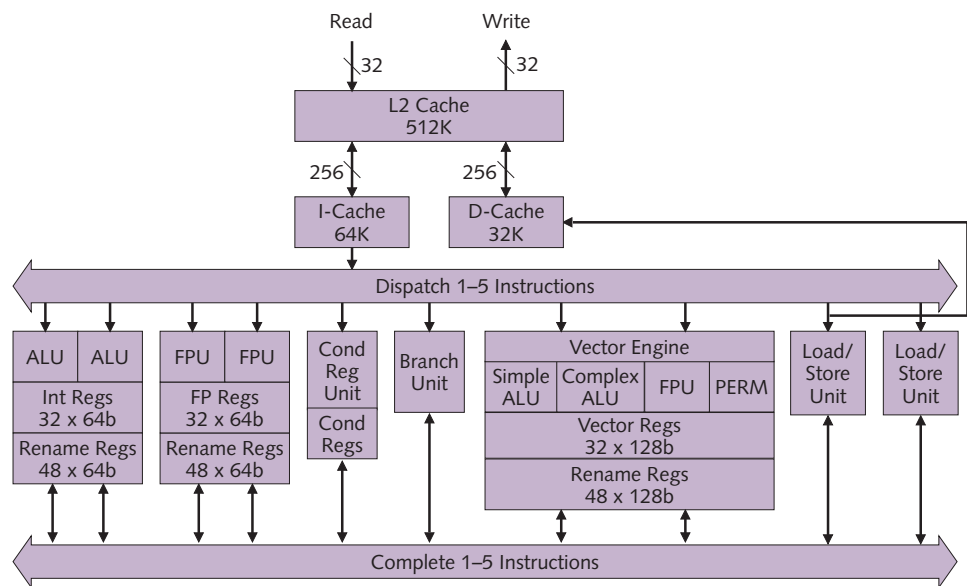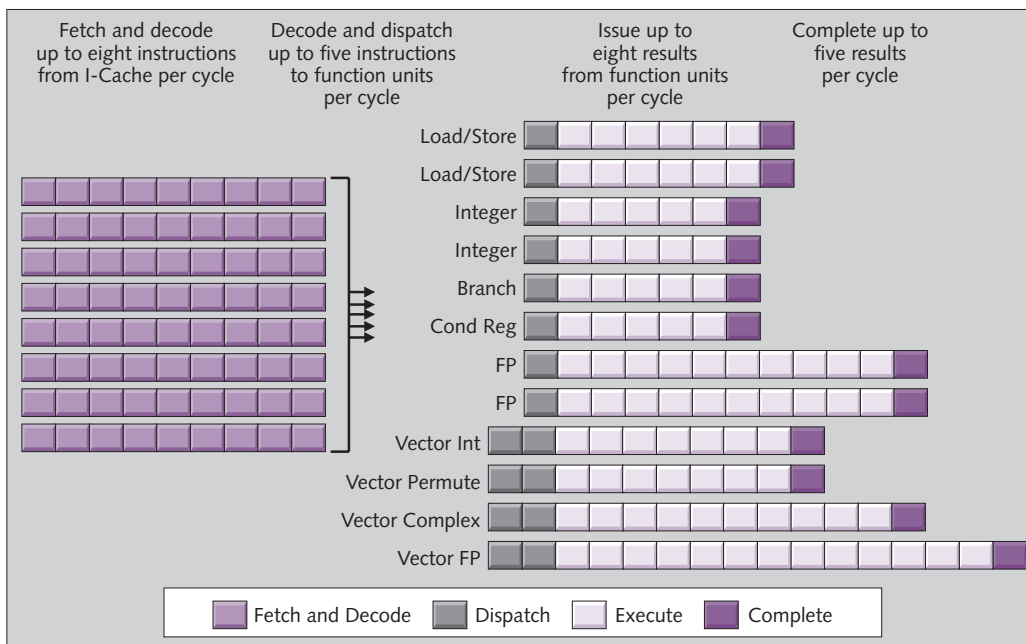

Figure 1. The PowerPC 970 has the same basic processor core as the Power4. Note the prodigious number of function units.

CISC legacy of variable-length instructions by painstakingly decoding the longer instructions into smaller RISC-like operations for easier digestion by the inner core. The Pentium 4 and Athlon XP each devote several stages to that purpose.

IBM says the 970 needs nine fetch and decode stages for similar reasons. The processor checks for constraints (such as data dependencies) that would impair instruction reordering; organizes the instructions into dispatch groups; and—most surprisingly—performs the same kind of instruction fission that modern x86 processors perform. IBM calls it "instruction cracking": breaking some of the more-complex RISC instructions into smaller, easier-to-schedule micro-ops. A typical example is the PowerPC's multiply-add instruction, which the 970 cracks into separate multiply and add operations. (Perhaps we should call this technique FRISC: further reduced instruction-set computing.)

As the pipeline diagram in Figure 2 shows, the 970 has 16 stages for integer-type instructions, 17 stages for load/store instructions, 21 stages for floating-point instructions, and up to 25 stages for SIMD vector instructions.

Pipeline depths vary for SIMD operations, because the AltiVec vector-processing engine contains two function units divided into subunits. The engine requires this level of specialization to handle the 162 different AltiVec instructions (see *MPR 5/11/98-01*, "AltiVec Vectorizes PowerPC"). With its 128-bit-wide datapaths and registers, the vector engine can execute multiple operations on various combinations of operands: four 32-bit integers, eight 16-bit integers, sixteen 8-bit integers, or four 32-bit single-precision floating-point values.

One function unit in the vector engine handles permutation-type instructions, and the other unit executes integer and floating-point instructions. The permute unit has a 16-entry buffer that queues those types of instructions for issue to the execution stages of the pipeline. The other function unit has a 20-entry queue, which is larger because it issues instructions to three subunits: one for simple integer operations, a second for complex integer operations, and a third for floating-point operations.

When using its two regular FPUs for scalar operations, the 970 can execute 7.2 GFLOPS at 1.8GHz (two 32- or 64-bit multiply-adds per cycle). When using its vector FPU for SIMD operations, the 970 can execute 14.4 GFLOPS at 1.8GHz (four 32-bit multiply-adds per cycle). The vector FPU doesn't support double-precision 64-bit operations, but that's not a serious handicap. The vast majority of floating-point-media data types (such as the operands for 3D graphics) are single-precision 32-bit values.

AltiVec is at least as powerful as Intel's MMX, SSE, and SSE2 extensions, as well as AMD's 3DNow and 3DNow Professional extensions. All are based on 128-bit SIMD operations. Although older Macs with G3 and earlier CPUs lack AltiVec (which Apple calls the Velocity Engine), we consider AltiVec compatibility an absolute requirement for any new PowerPC chip intended for Macs. Some strategic Mac software, such as Adobe PhotoShop, relies on AltiVec to accelerate critical functions. Of course, a vector-processing engine would be equally welcome in a Unix workstation, a 3D-graphics render-farm server, and some embedded systems. But IBM's decision to equip the 970 with AltiVec is a strong indication that Apple is a defining customer for the chip.

I/O Bandwidth Raises the Bar

With core clock frequencies rising much faster than bus frequencies in recent years, CPUs face an unhappy future of data starvation. The 970 tackles that problem head-on with the fastest front-side bus (FSB) of any processor likely to appear in 2003.

The 970's FSB consists of two 32-bit unidirectional buses, one for loads and the other for stores. Technically, they're source-synchronous point-to-point interfaces instead of true multidrop

Figure 2. The PowerPC 970 has deep pipelines in all function units, beginning with an unusually long nine-stage section that fetches and decodes the instructions.

buses. In that respect, they're similar to the Alpha EV6-derived FSB on Athlon XP processors. But the 970 drives its FSB at a phenomenally high effective clock speed: up to 900MHz. This compares with 533MHz for the fastest Pentium 4 and 333MHz for the fastest Athlon XP.

To achieve that speed, the 970 uses a 4-to-1 core-to-bus clock ratio, so the 1.8GHz part has a 450MHz base bus frequency. That compares with up to 167MHz for the Athlon XP and up to 133MHz for the Pentium 4. The 970 transfers two data phases per clock period, so the effective bit rate is twice the base frequency—or 900MHz for the 1.8GHz part. A 1.4GHz PowerPC 970 would have a 350MHz base clock and an effective 700MHz bit rate—still very impressive.

At 900MHz, the FSB provides 7.2GB/s of raw aggregate bandwidth, but it needs some of that capacity to carry the multiplexed address and control signals on the same wires as the data. Therefore, the peak data bandwidth is 6.4GB/s. That's about 52% more aggregate bandwidth than the fastest Pentium 4 and about 137% more aggregate bandwidth than the fastest Athlon XP. However, the x86 processors have bidirectional 64-bit FSBs instead of unidirectional 32-bit FSBs. For individual read or write transactions, the 970 has 3.2GB/s of peak data bandwidth, compared with 4.2GB/s for the Pentium 4 and 2.7GB/s for the Athlon XP. In some instances, the Pentium 4 would have an advantage over the 970. We believe that in a typical server application that rapidly moves lots of data on and off the chip, the 970's higher aggregate bandwidth would be more telling.

Of course, the Pentium 4 and Athlon XP are shipping now, while the 970 is at least eight months away from production. Bus speeds could very well change in the meantime. In any case, the 970 is a huge improvement over the hard-pressed G4+, which has a single-data-rate 167MHz FSB that provides a paltry 1.3GB/s of peak bandwidth.

The 970 pipelines the bus transactions, can handle 21 outstanding transactions at a time, and can read and write data out of order—additional improvements over the G4+. Although the 970's bus protocol multiplexes the data, address, and control signals, it uses sideband signals for cache snooping and acknowledgments.

IBM describes the 970's FSB as an "elastic interface," because it can tolerate multiple-cycle wire delays between the CPU and the core-logic companion chip. Special hardware corrects any clock skew in the data lines. This should make it easier to design motherboards, especially for SMP systems that can't place every CPU in close proximity to the companion chip. There are a few other provisions for SMP, such as mechanisms for processor synchronization, memory sharing, and memory coherency, but IBM hasn't described those features in detail.

IBM's Peter Sandon describes the PowerPC 970 at the recent Microprocessor Forum.

The on-chip caches of the 970 are comparable to those in the latest x86 desktop and server processors. The L1 instruction cache is 64K, is direct-mapped, and has a 32-entry prefetch buffer. As mentioned earlier, the 970 can fetch eight instructions per cycle from this cache. The direct mapping is unusual: primary instruction caches in desktop/server processors are usually multiway set-associative. However, IBM says its profiling revealed that a multiway L1 cache would add only a negligible amount of performance (about 1%) with this processor.

The L1 data cache is 32K and two-way set-associative. The L2 cache is 512K and eight-way set-associative, and it is linked to the CPU core by an interface that runs at the core frequency. The caches can prefetch from as many as eight datastreams at a time.

Performance Should Be Competitive

For the 1.8GHz chip, IBM estimates the performance at 937 SPECint2000 and 1,051 SPECfp2000. IBM deems those estimates "conservative." With the less-useful but often-quoted Dhrystone 2.1 benchmark, the 970's estimated performance is 5,220 mips, or 2.9 mips/MHz. IBM also offers an unusual metric: the 970 can encrypt or decrypt 18 million 64-bit keys per second with the RC5 algorithm, a chaining block cipher developed by RSA Data Security.

The SPEC scores are the most relevant for comparisons with other desktop and workstation/server CPUs. IBM's SPECint2000 estimate is higher than the score of any other workstation processor we track, edging out the faster-clocked 2.53GHz Intel Xeon by 5% (see *MPR 8/26/02*, "Chart Watch"). The 970 is even 19% faster than the Power4 that sired it, although the 1.8GHz 970 has the advantage of a 38% faster clock speed with a similar microarchitecture.

Floating-point performance is strong but not a trendsetter. IBM's SPECfp2000 estimate would put the 1.8GHz 970 in third place on our most recent chart of workstation processors, trailing Intel's 1.0GHz Itanium 2 by 29% (1,356 vs. 1,051) and the 1.3GHz Power4 by 4% (1,098 vs. 1,051). Still, the 970 would outperform a wide range of other workstation processors currently available, including the 1.0GHz Alpha 21264C; 1.8GHz Athlon XP; 750MHz PA-RISC 8700; 600MHz MIPS R14000; and 1.05GHz UltraSparcIII.

Again, however, remember the 970 probably won't appear in systems for almost a year—and a lot can happen in a year. Intel's Pentium 4 will be pushing 4.0GHz by then, allowing the x86 to maintain the same 2-to-1 clock-frequency advantage over the PowerPC that it enjoys today. Other competing processors will be faster by then, too. Of course, clock frequency isn't an absolute indicator of performance, especially when comparing two disparate RISC and CISC architectures,

but benchmarks prove that a heavenly clock frequency is effective penance for the original sin of CISC.

More to the point is what the 970 can offer users wedded to the PowerPC. The G4+ is falling further behind the x86 with every new Intel and AMD announcement, and Motorola's G5—announced in 2001—has yet to appear in systems. Mac users have a new operating system and would appreciate a new processor to make it shine. OS X is also Apple's first robust server operating system, so an SMP-ready processor with strong server parentage is an attractive option. Recently, Apple has been making inroads into the server market (see *MPR 6/3/02-03*, "Apple Serves Up XServe").

Certainly, IBM could find a nice home for the 970 in its own Unix and Linux workstations and entry-level servers. But it's a good bet the 970 will also end up in a Mac—unless Apple's thinking is even more different than advertised. ◇