MIPS Roadmap Focuses on Bandwidth

"Beast" to Deliver 5 GBytes/s; R12000 Provides Intermediate Step

by Linley Gwennap

Silicon Graphics (SGI) today revealed a new MIPS roadmap that changes the company's processor-development strategy. In the short term, the plan accommodates delays in the 0.25-micron version of the R10000 by replacing that product with a modified device called the R12000, due early next year. For the longer term, the plan outlines a strategy for competing in the post-Merced era, a strategy that emphasizes SGI's role as a high-end computing supplier. The new roadmap includes the first details about the processors code-named Beast (also known as H1) and Alien (H2).

Little Performance Improvement in 1997

When the 0.35-micron R10000 began shipping in 1Q96, SGI expected the design to move to a more advanced IC process within a year, boosting the clock speed significantly. Fab partners NEC and Toshiba both planned to deploy 0.25-micron CMOS in early 1997 (see MPR 9/16/96, p. 11). This plan would have given the high-end MIPS processor a quick speed boost with no design effort.

Had NEC and Toshiba delivered on this aggressive plan, they would have been the first microprocessor makers to reach the 0.25-micron mark. Unfortunately, the Japanese vendors ran into snags with their 0.25-micron processes, which they are still trying to work out. Thus, the companies continue to ship R10000s at a maximum speed of 200 MHz, just as they did more than a year ago.

The delays in the manufacturing process put SGI in a difficult position. If the company continues with its original plan, the 0.25-micron part will not be as competitive when it finally emerges. Making the performance competitive, however, will require design changes that could further delay initial shipments. Doing both a simple shrink and a modified design would stress SGI's limited verification resources.

To maximize performance, SGI decided to go back to the drawing board. The company is now developing a modified version of the original design that should improve its per-cycle performance. The new design will be marketed as the R12000 to distinguish it from the current version. SGI has committed to delivering the chip by 1H98, but if all goes well, it could appear sooner.

In the meantime, the 200-MHz R10000 will fall further off the pace in competitive performance. When the part was first released, it nearly matched the performance of the industry-leading 333-MHz 21164. Since then, Digital has pushed the clock speed of the 21164 to 500 MHz and is sampling 600-MHz parts, opening a 30% performance gap. SGI has not ruled out deploying a faster R10000 as a stopgap measure, but the company is focusing its efforts on the forthcoming R12000.

R12000 Improves CPU Efficiency

In most ways, the R12000 will be identical to the R10000. The two chips share the same basic CPU core, the same on-chip caches, the same pinout and bus structure. The new chip will use a 0.25-micron process, improving clock speed and reducing die size.

The design changes to the CPU core are limited in scope but take advantage of the flexibility of the R10000 design (see MPR 10/24/94, p. 18). For example, the size of the active list has been increased from 32 entries to 48. This change enlarges the "window" within which the processor can reorder instructions, improving its ability to issue instructions in parallel even when several instructions are stalled, waiting for data. This change brings the R12000 closer to the PA-8000, which has a 56-instruction window, and the forthcoming 21264, which can reorder up to 64 instructions.

The R12000 includes a new branch target address cache (BTAC) with 32 entries, adopting a design similar to several PowerPC processors. Because the R10000 has no BTAC, taken branches create a one-cycle "bubble" in the instruction fetch stream. SGI had hoped the deep instruction queues would hide this bubble, preventing it from disrupting the execution stream, but performance tests showed that taken branches with nearly empty queues occurred frequently enough to justify the addition of a small BTAC.

The new chip enlarges the size of the branch history table (BHT) from 512 entries to 2,048, reducing the number of performance-draining mispredicted branches. This puts the R12000 on a par with the 21164, which has the largest BHT among current RISC processors. The R12000 also doubles the size of the way-prediction table, an unusual structure that allows the R10000 to support a two-way set-associative external cache (see MPR 10/24/94, p. 18). The larger table supports L2 caches of up to 4M without aliasing.

SGI did not specify the die size of the R12000, but the core CPU changes shouldn't add much area to the die—perhaps 10%. Since the 0.25-micron process will provide a significant reduction, we estimate the R12000 will be around 200 mm^2, a third smaller than the R10000.

With a simple shrink, the 0.25-micron process was expected to push the R10000 to 275 MHz. The redesign has allowed SGI to rework some critical timing paths, and the company now expects the R12000 to reach 300 MHz or better in the new process.

Assuming the core changes increase performance by 10–20%, we estimate a 300-MHz R12000 will deliver about

16–18 SPECint95 (base) and 25–27 SPECfp95 (base), on a par with the fastest processors expected to ship in 1997. In 1H98, however, Digital's 21264 and HP's PA-8500 are expected to deliver in excess of 25 SPECint95 and 40 SPECfp95, again putting MIPS behind in performance.

SGI Adopts High-End Strategy

As Figure 1 shows, there are two basic strategies used by microprocessor vendors to deploy full product lines; SGI has switched between the two. The first, used by Intel and others, requires designing a new high-end CPU every couple of years while using cost-reduced versions of older CPUs to fill in the lower price points. The second, used by PowerPC and others, creates CPUs for various price points (low, middle, high) and tweaks them over time to improve performance.

The advantage to this latter strategy is that the CPU cores can be designed to meet the needs of specific market segments. The disadvantage is that, since CPU designers are spread across two or three product lines, the high-end line gets a new core every three or four years, versus every two years in the Intel plan. Even with process shrinks and minor tweaks, a single CPU core cannot remain competitive for three years.

The original MIPS strategy was to develop high-end cores, then cost-reduce them. This strategy worked for the R3000 and R4000. With the R10000, the company changed directions, creating high-end, midrange, and low-cost lines based on the R10000, R4600, and R4200, respectively (see MPR 3/6/95, p. 4). For example, the company at that time planned to develop a chip called the D2, a midrange follow-on to the R5000, as well as consumer-oriented processors plus a new high-end chip called the H1.

The latest roadmap goes back to Plan A. SGI has canned the D2 effort and will instead rely on the R7000 and R10000 for its lower-cost workstations. The system vendor will also cease efforts on most parts for consumer products, leaving this market to MIPS licensees such as NEC, Toshiba, and Philips (see MPR 5/12/97, p. 13). One key area of focus remains Nintendo, for which SGI will still develop new processors. This restructuring allows SGI to focus all its CPU designers on new high-performance processors.

Despite this refocusing, the schedule for the H1 has slipped from 1H98 to 1H99. The designers freed up from other projects have been deployed in part on the R12000 and in part to accelerate the next high-end processor, the H2. The company did not provide a specific timeframe for the H2; we expect it to appear in 2001, as Figure 2 shows.

H1 Delivers Memory Bandwidth

SGI believes the H1, when it ships, will have better memory bandwidth than any other microprocessor available at that time. The company expects the chip to sustain more than 5 Gbytes per second of main-memory bandwidth, nearly ten times the sustainable memory bandwidth of the current R10000 and more than twice the bandwidth of the IBM

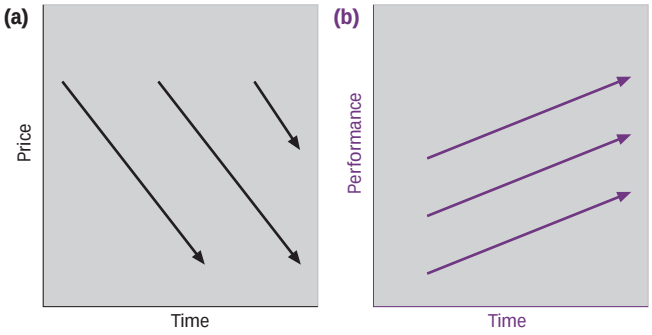

Figure 1. (a) Vendors such as Intel develop high-end processors only and cost-reduce them over time. (b) Other vendors develop a range of specific cores and improve them over time.

P2SC and Digital's forthcoming 21264, the best such processors revealed to date.

Such high bandwidth can be achieved through either very wide or very fast buses. While SGI did not reveal many Beastly details, it did note the device will have in excess of 1,500 pins, plenty to support a 256-bit main-memory bus and a 256-bit cache bus, for example. A 256-bit bus operating at 200 MHz could sustain more than 5 Gbytes/s.

With this bandwidth, Beast will excel on applications with vector data sets too large to fit in the external cache. Since Beast is likely to comfortably support caches of several megabytes, applications must be large indeed to stress the memory system. These types of applications include complex images, data bases, and scientific computation.

Not coincidentally, these categories reflect the main areas of SGI's systems business. The company's workstations are frequently used to create complex 3D images and animation, in both the engineering and entertainment communities. With its acquisition of Cray Research, SGI has become the dominant player in scientific computing; its high-end

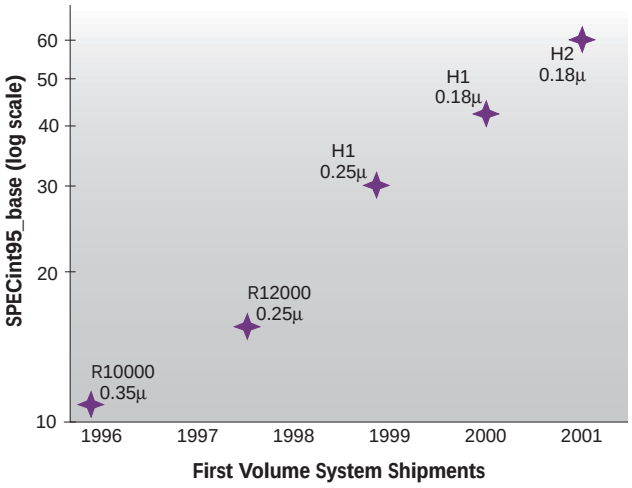

Figure 2. The new MIPS roadmap includes the R12000 (an enhanced version of the R10000) and two high-end processors known as H1 (Beast) and H2 (Alien). (Source: MDR, based on information from SGI)

For More Information

Silicon Graphics has not disclosed full technical details on the R12000, H1, or H2 processors. For more information on Silicon Graphics' MIPS processors, access the Web at *www.mips.com.*

systems are used to solve many large, complicated problems. Today, the Cray systems use proprietary processors, but in the H1 timeframe, SGI plans to convert these systems to the MIPS instruction set.

The H1 will be fully MIPS compatible and will implement the MIPS V extensions for parallel single-precision floating-point operations, as well as the MDMX extensions for parallel multimedia operations (see MPR 11/18/96, p. 24). It is likely to appear first in 0.25-micron CMOS, with a shrink to 0.18-micron CMOS providing a performance boost in 2000. The initial version is likely to measure at least 300 mm^2, which, combined with the enormous package, will make Beast very expensive to manufacture.

Other than the memory bandwidth, SGI did not provide any performance projections for the part. We expect the initial H1 to deliver around 30 SPECint95 and 60 SPECfp95, roughly the same as the 0.35-micron 21264 due a year earlier. These metrics do not capture the real performance of the H1, however, as they do not stress main-memory bandwidth. While the H1 will probably lag 0.25-micron versions of the 21264 and Intel's Merced on SPEC95, it may be able to match these chips on the forthcoming SPEC98 benchmarks and is likely to surpass them on applications that are both memory- and FP-intensive.

The company did not reveal the manufacturer for the H1. The rationale for a semiconductor vendor to market high-end MIPS chips diminishes as the number of MIPS system vendors continues to dwindle. Although five companies produced the R4000, only two build the R10000, and we would not be surprised if only one manufactures the H1. SGI could adopt a foundry model similar to Sun's rather than continue with its open market strategy.

H2 Designed for Parallel Systems

SGI also disclosed a few details about the follow-on to Beast, code-named Alien or H2. This chip will be optimized for systems with large numbers of processors and NUMA (non-uniform memory access) architectures. While the CPU itself will be more powerful than Beast, the emphasis for Alien will be system-level bandwidth.

We believe Alien will include a significant amount of memory on-chip; using a 0.18-micron process, megabytes of cache could fit onto a large die. To facilitate highly parallel systems, the chip could include several high-speed buses, allowing connections in a 2D or 3D array simply by wiring the processors directly to one another. If Alien also included a direct interface to local DRAM, a NUMA system could be built from nodes that each consisted of merely a single Alien CPU plus DRAM.

This type of design would fit well in SGI's line of supercomputers and could also be used in high-end enterprise servers. As SGI's high-end CPU, Alien will also be used in power desktops, but these systems cannot take full advantage of its multiprocessor scalability and thus are less likely to offer performance competitive with that of other contemporary desktop processors.

The H2 takes advantage of the trend of multithreaded software. As multiprocessor and massively parallel systems become more common for scientific applications, these programs are being designed to scale well across large numbers of CPUs. Even mainstream workstation software is being multithreaded to take advantage of MP desktop systems. The H2 will be most effective in these configurations.

Competing With Merced

We believe Merced, with backing from Intel and HP, will set a new performance mark for the microprocessor industry and make it difficult for other CPU makers to compete without developing their own state-of-the-art instruction set. SGI is adamant that all of its processors, including the H2, will continue to use the MIPS instruction set and that there will still be a role for RISC after Merced.

The initial description of Beast lends some credence to this assertion. Both Intel and HP are motivated to provide strong integer and floating-point performance in Merced, but for Intel to have any hope of moving Merced into the PC infrastructure at some point, there are certain limitations. For cost reasons, Merced is unlikely to support a 1,500-pin package, for example. Without such a large package, it will be difficult to match Beast's bandwidth.

Difficult, but not impossible. If Merced were to include four Direct RDRAM channels on the CPU itself, it could sustain the same 5 Gbytes/s of memory bandwidth as Beast while using only 64 pins. Alternatively, the Merced core could easily be repackaged into two products with different bus widths, providing a low-cost version and a high-bandwidth version. Intel has said it will use C4 bonding for some future CPUs, which could allow it to support much higher pin counts without greatly increasing manufacturing cost.

In any case, SGI's new strategy is a clear admission that competing with Merced in mainstream desktops and servers is an iffy proposition at best. SGI may be able to differentiate its low-end systems by their scalability and graphics performance, but Merced-based systems are likely to offer much better price/performance. With the H1 and H2, the company is going upscale, avoiding head-to-head competition with Merced. The questions that remain are how high these Merced systems can go and what the size of the remaining market is. The financially struggling SGI must hope that this remaining market is at least as big as its current market—and more profitable. M