

Sun Spins Low-Cost UltraSparc *UltraSparc-2i to Boost Low-End Workstations in Late 1997*

by Linley Gwennap

Aiming to revitalize the feeble performance of its low-end workstations, Sun is developing a cost-reduced version of its high-end UltraSparc processor. Although the chip itself, dubbed UltraSparc-2i, is no less expensive to build than UltraSparc, it reduces system cost by integrating main-memory and system-bus interfaces and supporting lower-cost SRAM and DRAM chips. US-2i is the first Sun processor to integrate a PCI bus interface, presaging a move from SBus to PCI in Sun's own workstations.

The company expects the modifications to have relatively little impact on SPEC performance, particularly on the integer side: US-2i is estimated to deliver 10.6 SPECint95 and 14 SPECfp95 (base) at its target clock speed of 300 MHz. These scores would be roughly seven times better than the anemic ratings of the 110-MHz MicroSparc-2, currently anchoring Sun's low-end workstation line. This help is far away, however, as US-2i has not yet taped out, and Sun projects system shipments no sooner than 4Q97.

As the high-end PC and low-end workstation markets merge, this performance boost is critical for Sun to be competitive. Based on SPEC95, a MicroSparc-2 offers the integer performance of a 486DX4 and the floating-point power of a midrange Pentium. MS-2 systems are selling today because of the huge base of software for Solaris and SPARC. With Windows NT becoming a viable workstation platform, Sun will have to compete head to head with x86-based workstations, and US-2i is needed to match P6-class performance.

Integration Reduces System Cost

Sun has a history of building low-cost systems around highly integrated microprocessors, stretching back to the original MicroSparc-1 (see 061402.PDF). US-2i adopts the same strategy, including a memory interface and bus interface, as Figure 1 shows. One key difference between US-2i and the MicroSparc chips is the former's external L2 cache. The MicroSparc chips had on-chip primary cache (24K in the

case of MS-2) but no provision for external cache. This limitation kept pin count down but caused a significant performance drop between SPEC92 and the larger SPEC95 benchmarks, which, like most application software, come nowhere near fitting into 24K of cache.

US-2i remedies this problem by retaining the direct L2 cache interface of UltraSparc, although it trims the width of this bus to 64 data bits instead of 128. This design still allows up to 2M of external cache, although most systems are likely to use 512K for cost reasons. This amount of cache memory is much better suited to both SPEC95 and real application code than are the tiny MicroSparc caches.

The second big difference relative to MicroSparc is the PCI interface, a strategic change we predicted two years ago (see 081301.PDF). The MicroSparc processors connect directly to Sun's SBus. Although SBus is relatively popular for a workstation bus, PCI cards are more widely available and far less expensive than SBus cards, due to the higher volumes in the PC market. PCI also offers significantly better bandwidth than SBus. One drawback is the need for PCI card vendors to port their drivers to SPARC.

The third major change is the CPU core. US-2i includes essentially the complete UltraSparc-2 CPU, a four-way

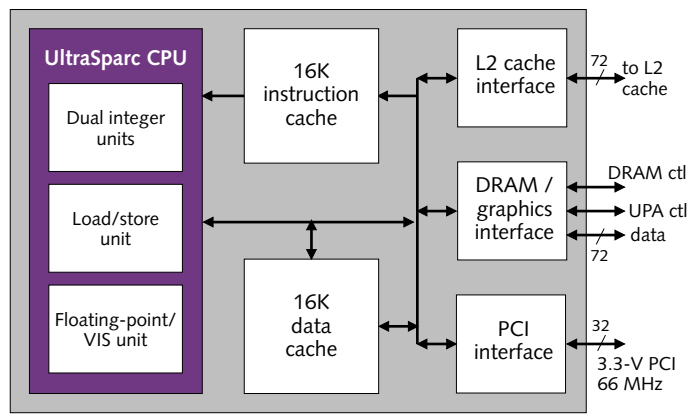


Figure 1. UltraSparc-2i combines an UltraSparc CPU core with DRAM and PCI controllers to lower system cost.

superscalar in-order processor with 16K of instruction cache and 16K of data cache (see 081301.PDF). This core is far superior to the scalar MicroSparc CPU, which topped out at 110 MHz in a 0.4-micron process. With a deeper pipeline, US-2 reaches 167 MHz in 0.42-micron CMOS, can execute four times as many instructions per cycle, and contains the VIS multimedia extensions.

US-2i will be built in 0.35-micron five-layer-metal CMOS by a foundry Sun would not identify. Sun believes US-2i will reach speeds of 300 MHz in this process, comparable to the speed of US-2 in Texas Instruments' 0.29-micron five-layer-metal CMOS process. The die size of US-2i should be around 144 mm², slightly smaller than the 149-mm² US-2, but the die size is not final, since the part has not yet taped out. The added system logic takes up about 14% of the US-2i die, indicating that the new chip would be slightly larger than US-2 if built in the same process.

Separate Cache Bus Eliminates UDBs

The 72-bit cache interface includes 64 bits of data and 8 bits of parity, allowing a minimum 256K cache to be built from just three parts: two 32K×36 data SRAMs and one 32K×18 chip for tags. Separating the cache bus from the main memory bus eliminates the need for the two external UltraSparc data buffers (UDBs) required in all US-1 and US-2 systems. At a list price of \$65 each, the UDBs add cost to the system and increase the footprint of the processor subsystem. US-2i sidesteps these problems.

The new chip supports two different cache speeds. For high performance, the cache latency is four cycles, with a new access starting every other cycle. This mode requires 167-MHz late-write synchronous SRAMs with a 300-MHz CPU and is similar to the UltraSparc-2 cache timing.

To take advantage of lower-price cache chips, US-2i also supports a mode with a cache latency of six cycles; the repeat rate remains at two cycles. In this mode, the SRAMs run at exactly half the CPU clock speed, for example, 150-MHz cache chips support a 300-MHz CPU. This slower configuration reduces performance by about 8%, according to Sun.

Sun is hoping that the price of these fast synchronous

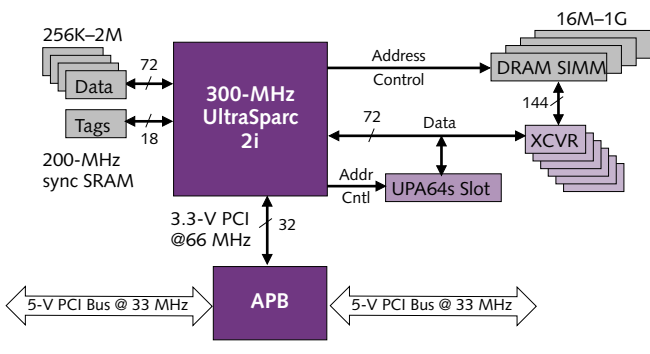


Figure 2. A typical UltraSparc-2i system includes the CPU, 512K of external L2 cache, a 144-bit EDO memory subsystem, and a PCI bridge that supports two 5-V PCI buses at 33 MHz.

SRAMs will drop as Intel's forthcoming Klamath (P6) processor becomes prevalent. Klamath will use an external half-speed cache. Since we expect to see Klamath (and successor) parts at speeds of up to 300 MHz in 1997, this implies vendors will be producing synchronous SRAMs at up to 150 MHz in high volumes. These parts may have a special interface for Klamath, but Sun hopes the parts for US-2i will ride a similar curve to low prices.

At eight bytes per access, it takes four accesses (10 cycles in high-performance mode) to fill a line in the instruction cache; data-cache lines are divided into two subblocks that each fill in two accesses. This is the biggest performance difference compared with US-2; with its 128-bit interface, that chip can fill its on-chip cache with half as many accesses. An instruction-cache line, for example, takes just two accesses (6 cycles) to load in US-2.

Commodity DRAM Lowers Cost

Staying with low-cost commodity memory chips, the integrated processor connects directly to either asynchronous or EDO DRAM, as Figure 2 shows. With a 72-bit bus (including ECC), the chip can sustain 350 Mbytes/s to DRAM. Sun considered supporting synchronous DRAM, but didn't because US-2i's EDO memory subsystem delivers the same performance as a 75-MHz SDRAM subsystem.

It achieves this performance by using a double-wide bank of 60-ns EDO DRAM, resulting in a 144-bit memory subsystem. Each bank operates at roughly 37.5 MHz, extending the EDO CAS cycle to a relatively leisurely 26.5 ns. The transceivers shown in Figure 2 funnel the data onto the processor's 72-bit data bus at 75 MHz.

This bus speed is conveniently one-fourth of the CPU clock speed at 300 MHz, although the timing of the memory bus can be adjusted to support different CPU and DRAM timings. Using faster DRAM would improve performance beyond Sun's stated estimates, which assume the 60-ns EDO subsystem described above.

Integrating the memory controller reduces system cost by avoiding the need for an expensive custom memory-control ASIC. It can also improve performance by decreasing memory latency. On a miss in the L2 cache, the US-2i takes 50 cycles before receiving the first (critical) word from memory. In contrast, a US-2 system, which has main memory on Sun's proprietary UPA bus, carries a latency of 64 cycles. The US-2 system has a bandwidth advantage, however, since the UPA bus is 128 bits wide.

Compatible with UPA Graphics

Although eschewing it for memory, the US-2i designers wanted to retain some compatibility with UPA, for which Sun has designed several graphics accelerators. The designers wanted to avoid the high pin count of yet another bus interface, however. As a compromise, US-2i includes about 35 pins devoted to the UPA control and address signals, enough to support a slave-only subset of UPA that Sun calls UPA64s.

As Figure 2 shows, a single 64-bit UPA graphics card can connect to the main-memory data bus; the same transceivers that multiplex the DRAM data onto the 64-bit bus isolate the graphics card from the DRAM. Data transfers to the UPA card occur at one-third of the CPU clock speed, up to 100 MHz. This configuration is likely to be supported in Sun's low-end workstations.

A very low cost system, however, could use a standard PCI graphics card to reduce cost. In this configuration, the UPA slot could be eliminated, but the cost and board area of the external transceivers remain, since they are needed for the double-wide DRAM subsystem. A typical design requires six 24-to-12-bit transceivers (e.g., TI's 74ALVC16268).

The initial US-2i design appears focused mainly on Sun's workstation needs, where UPA is a necessity. To push US-2i into a broader range of low-cost markets, including high-end embedded systems, Sun will add SDRAM support in a future version and may cut out the extra UPA pins to save package cost. With SDRAM and no UPA slot, the external transceivers could be eliminated as well with no performance penalty. We expect the price premium for SDRAM to be nearly zero by 1998, increasing the attractiveness of this configuration.

Several PCI Options

The new processor supports a 32-bit PCI bus that can operate at either 33 or 66 MHz. It is fully compliant with PCI version 2.1 and supports up to four bus masters, but it supports only 3.3-V I/O. Sun plans to provide a bridge chip called APB that supports two 33-MHz 5-V PCI buses while connecting to US-2i at 66 MHz, as Figure 2 shows. In a maximum configuration, with four bridge chips, the processor could support eight PCI buses and more than 32 devices.

Some PCI chips, designed to be mounted on the motherboard instead of add-in cards, support 3.3-V I/O. In a simple US-2i system, with only a PCI-based super I/O chip and an Ethernet chip, for example, the PCI bridge can be left out. Without the PCI bridge, however, the chip can't connect to standard 5-V devices.

The chip has two PLL circuits, including a separate PLL for the PCI interface. This design allows the PCI bus to operate in its own clock domain, running at either 33 or 66 MHz regardless of the CPU clock speed. Sun claims that its synchronizer circuits impart little extra latency, and this design avoids the need to slow the PCI bus for certain CPU clock speeds. Most PC system-logic chip sets today allow the PCI bus to run at 33 MHz regardless of the CPU speed.

The PCI interface includes a DMA engine that can sustain the full bandwidth of a PCI device to the main memory. A 16-entry I/O TLB translates 32-bit virtual PCI addresses to 34-bit physical addresses. This large physical address space is needed for compatibility with Solaris. On a miss, this TLB can access page tables in main memory without interrupting the CPU.

Long Gestation Period for US-2i

When Sun began shipping its first MicroSparc-2 systems in 2Q94, it planned to follow up with an enhanced device called MicroSparc-3 by 2H95. The original MS-3 was just a process shrink with minor modifications, boosting performance by about 50%. But within a few months, the company, realizing the performance weaknesses of MS-2 needed to be addressed, changed plans. The revised MS-3 was to have a revamped CPU core delivering three times the performance of the original MS-2, delaying shipments until 2H96.

By the end of 1995, the company changed tacks again, this time saying that MS-3 would be based on the UltraSparc CPU core and would be available in 1H97. The latest effort, now dubbed UltraSparc-2i, is slated to provide seven times the performance of MS-2 with a 4Q97 delivery date, two years later than the original MS-3 plan.

Sun is not the first company to redirect and delay its processor-design projects. But these delays have left Sun's MS-2 workstations without an upgrade for nearly 30 months, an eternity in the microprocessor market.

High Power, Pin Count

Inheriting the high performance of UltraSparc-2 comes with the drawback of that processor's power consumption, which maxes out at about 30 W (*see 091505.PDF*). With the additional on-chip logic, US-2i is estimated to burn up to 36 W (maximum). Although this power level has not yet been measured, UltraSparc-2 itself has been characterized, so the likelihood of the US-2i rating falling significantly is small. The high power comes despite a core operating voltage of just 2.5 V; the chip also requires a 3.3-V input to drive its I/O signals.

To position US-2i for high-end embedded systems, Sun needs to reduce the power considerably. It plans to offer a 150-MHz version, dropping the power to about 18 W. In fact, this version could probably operate at 2.0 V instead of 2.5 V; if so, this embedded version might use 12 W, in the range for a network router or high-end printer.

Alert readers have already realized that supporting three external buses—the 72-bit cache bus, 72-bit memory bus, and 32-bit PCI bus—plus extra control signals for UPA64s adds up to a large pin count. Furthermore, the chip requires extra power and ground connections to pull in the maximum 14 A. Sun has not yet determined the final packaging configuration but expects 550–600 connections.

To support this large number of pins, US-2i, like many recent CPUs, will use a BGA package. Sun has relied on a 521-pin plastic BGA for its previous UltraSparc chips, so it has plenty of experience with this packaging. Unlike these earlier chips, US-2i will be mounted in the package using flip-chip attachment.

Price & Availability

Sun projects UltraSparc-2i to sample in 3Q97, with volume production in 4Q97 at speeds up to 300 MHz. The company has not yet announced pricing. For more information, contact Sun Microelectronics (Sunnyvale, Calif.) at 415.336.1558 or www.sun.com/sparc.

To minimize system cost, the US-2i and cache subsystem can be placed on a daughtercard. This card would contain the fastest signals, requiring an eight-layer PC board, but it could be fairly small. In this design, the motherboard would have only the 66-MHz PCI interface, 75-MHz DRAM interface, and optional 100-MHz UPA64s interface. These signals, with the possible exception of the UPA bus, could be implemented on a low-cost four-layer board.

Workstations, PC Collide

The impending collision of the workstation and PC markets has been predicted for years (*see* [0804ED.PDF](#)), giving Sun and other workstation vendors plenty of time to react. The emergence of the high-performance Pentium Pro processor and the maturity of Windows NT, embodied in the recent 4.0 release, are stoking the furnaces of these oncoming trains. With PC vendors such as Compaq jumping into the workstation market and workstation vendors such as HP offering Pentium Pro-based workstations, a spectacular impact is drawing very near.

Sun is hoping not to be one of the casualties. As a backup plan, the company is madly diversifying into Java products and high-end servers, but workstations remain its core business. Sun's high-priced UltraSparc workstations have boosted the performance competitiveness of its product line, particularly on applications that take advantage of either floating-point math or VIS. The company's sub-\$10K family, however, still relies on the aging and VIS-less MicroSparc-2, which is outperformed by a high-end Pentium on integer, floating-point, and multimedia code.

US-2i will greatly strengthen Sun's low-cost workstation line. A 300-MHz P6, which should be available at the same time as the 300-MHz US-2i, will offer slightly better integer performance but far less floating-point performance. By the time US-2i appears, Intel will have added MMX to the P6, so the advantages of VIS on multimedia code will be reduced. Sun's ability to acquire US-2i at foundry prices, rather than with Intel's high profit margins, may enable Sun to match the integer price/performance of P6-based workstations while offering superior FP. A key issue is whether Sun can match the low overhead of P6 workstations from PC makers.

Positioning US-2i against US-2 in Sun's own product line may be more of a challenge. The US-2i designers believe their chip will come within 5% of US-2 on SPECint95 and

within 15% on SPECfp95 at the same clock speed. The main performance difference is the lower cache and main memory bandwidth. The performance difference may be bigger on large applications that put a greater strain on the cache and memory systems, but perhaps not enough to justify the higher system costs.

US-2 will still be required for multiprocessor systems or for systems that require compatibility with SBus or the full UPA128 specification, but these requirements will be outside of Sun's mainstream by 1998. Fortunately, UltraSparc-3 is scheduled to appear shortly after US-2i, providing a performance boost for the high end. Thus, by mid-1998, Sun's workstation line should consist mainly of US-2i systems at the low end and US-3 systems at higher price points.

Looking to Broader Markets

While US-2i will shore up Sun's own systems, it doesn't offer a compelling incentive for most other system vendors. The few remaining SPARC-compatible system vendors will have to pay significantly higher prices than Sun for US-2i, since Sun obtains its chips at foundry prices but charges the market rate. These vendors may stay with SPARC to fit their products under Sun's system-price umbrella, but if Sun has to cut its system prices to compete with x86-based workstations, there will be little value in staying in the low-volume SPARC-compatible market.

Vendors that are not building SPARC systems today will find no reason to switch to US-2i. With Sun's NT port lying dormant, the only operating system available for SPARC is Solaris; an unaligned vendor is much more likely to be interested in the larger market for Windows NT systems, most likely using P6 processors.

Sun is also interested in the high end of the embedded market. Designers of high-end color printers and network routers are always looking for more performance, particularly from an integrated processor that sports standard memory and bus interfaces. The power dissipation of the 300-MHz US-2i, however, is far higher than most embedded applications can handle. The 150-MHz version reduces the power level but also cuts the core CPU performance in half, making it no better than a chip like today's R5000.

The best opportunities for getting the design into embedded systems rests with a future 0.25-micron shrink. In fact, given the late debut of US-2i, it is surprising the company has not targeted a 0.25-micron process from the start, and a quick shrink is likely. Combined with a move to a 1.8-V core, this shrink could bring power down to about 15 W while also reducing manufacturing cost and maintaining the core performance. This shrink version (perhaps called UltraSparc-3i?) will keep SPARC competitive with faster P6 processors. Thus, we expect Sun to have a competitive workstation lineup in 1998 and 1999, but the vendor will have to rely more on loyalty than performance to keep its customers until then. 