

# MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 8 NUMBER 13

OCTOBER 3, 1994

## UltraSparc Unleashes SPARC Performance

### Next-Generation Design Could Put Sun Back in Race

by Linley Gwennap

High-end SPARC performance, languishing at sub-Pentium levels, is set to receive a big boost next year when UltraSparc debuts. Sun expects this next-generation RISC chip to triple the performance of a 60-MHz SuperSparc, moving SPARC from the back of the pack to within hailing distance of the lead. The key to this incredible increase is a complete redesign of the processor pipeline to eliminate the constrictions of the SuperSparc design. The result: a projected clock speed of 167 MHz, a huge jump for Sun and a respectable rate compared with other next-generation RISC chips.

Unlike Digital, which has already measured the performance of the 21164, Sun's performance estimates are conjecture, as UltraSparc has not yet seen first silicon. Sun has built test chips to verify the speed of its design and has performed extensive timing simulations, hoping to avoid the embarrassment of its SuperSparc launch. The design avoids SuperSparc's fatal flaws (the double-pumped register file and TLB), but it remains to be seen whether Sun can deliver on its promises and turn a paper tiger into a real man-eater.

The first announced processor to implement the SPARC version 9 architecture (see [070201.PDF](#)), UltraSparc is a full 64-bit design. It can issue as many as four instructions per cycle to nine function units: two integer ALUs, one load/store unit, one branch unit, and five special-purpose units for floating-point and graphics calculations. The chip has moderate on-chip caches for a processor of its generation: 16K for instructions and 16K for data, less than SuperSparc. To make up for these modest caches, UltraSparc connects directly to a synchronous external cache that can return one result per cycle. In addition to SPARC V9, the design implements a unique set of graphics and multimedia instructions.

Sun has not announced price or availability for the new processor, which will be fabricated by Texas Instruments. We expect UltraSparc to begin shipping in volume in 3Q95, six to nine months later than the 21164.

### Flexible Instruction Alignment

Sun, with the largest installed base of any RISC system vendor, has always been concerned about the performance of existing (unrecompiled) binaries on new processors. UltraSparc implements a simple scheme that avoids the instruction-alignment restrictions that prevent the 21164 and other highly superscalar processors from achieving maximum performance without recompilation. The SPARC chip fetches instructions into a 12-entry FIFO buffer; the instruction dispatcher simply issues up to four instructions from the bottom of the buffer.

This scheme works well as long as the buffer is kept reasonably full. For starters, the instruction cache can deliver four instructions (128 bits) per cycle to the buffer, but branches can disrupt this flow. To counter this problem, the cache includes a "next" field that can redirect the fetch stream if the current instruction group contains a predicted-taken branch. For cache lines that do not contain such branches, this field contains the next sequential address. The contents of this field direct the next instruction fetch, eliminating any penalty for correctly predicted taken branches.

As they are loaded into the cache, instructions are partially decoded to determine if they contain a branch and, if so, what the target address is. This information is used to initialize the "next" field. In what is becoming a common superscalar design technique, the instruction cache stores four bits of decode information with each instruction as well as two bits of branch history per cache line. Sun's simulations show an 88% prediction accuracy on SPECint92 using these two history bits.

As Figure 1 shows, instructions are further decoded before being placed in the instruction buffer. Each entry in the buffer is 62 bits wide to contain all the decode information. This extensive information allows the dispatch unit to quickly decide which instructions can be issued and even allows time for a register file access, all in a single clock cycle.

Instructions are always issued in order; if an instruction cannot be issued due to a resource conflict or a register dependency, no subsequent instructions are issued on that cycle. Unlike SuperSparc, the new design does not cascade the ALUs; this change prevents dependent integer instructions from being paired but helps support the high clock rate. One special case is that a store can be dispatched in the same cycle as the instruction that calculates the store data; this case is handled by forwarding the result to the store queue.

There is one flaw that breaks the "no alignment" strategy. The first three instructions can be dispatched to any function unit, but the fourth can be sent to only the branch or floating-point units. Sun says that allowing the fourth slot to contain a general integer instruction would have greatly increased the amount of dependency checking but added little performance. Restricting the fourth slot also reduces the number of ports in the integer register file.

## Long Pipeline Includes FPU

UltraSparc uses a nine-stage pipeline, as Figure 2 shows. The basic integer pipeline is actually six stages, two more than in SuperSparc; the additional stages at the back end support the floating-point and graphics units.

The first two stages perform instruction fetch and decode. As noted above, the decoded instructions are placed in the instruction buffer. If the buffer is not empty (the typical situation), instructions may wait one or more cycles before being dispatched to the function units in the G (grouping) stage. The next two stages are the classic RISC execute and cache-access stages.

Instead of completing with a writeback in the sixth stage, three additional stages are added to wait for long-latency FP and graphics operations. These stages make it easier to resolve FP traps. The completion units hold results until they are written to the register file, reducing the amount of bypassing needed for the long pipeline.

For floating-point and graphics instructions, the fourth (E) stage is used for additional decoding and for accessing the FP register file. Delaying these actions simplifies the predecoder and the instruction buffer. Most floating-point operations have a three-cycle latency and complete in the N2 stage. Any pending traps are resolved in N3 and, assuming no traps, results are written to the register file in the ninth (W) stage.

If a load misses the primary data cache in the C stage, the physical address is used to initiate an external cache access. The load-use penalty for an L2 cache access is seven cycles. Because the external cache uses synchronous SRAMs that run at the same speed as the CPU, these cache accesses can be fully pipelined, resulting in one 128-bit result per cycle.

The front end of the pipeline is quite clean, adding only an extra dispatch stage to the five-stage pipeline used by most scalar RISC processors. This extra dispatch stage is common in highly superscalar processors to allow for grouping and issuing instructions to multiple function units. There are few pipeline hazards, primarily a one-cycle load-use penalty and a four-cycle mispredicted branch penalty. Proper code organization can eliminate the load-use penalty, while the branch-prediction logic strives to avoid the branch penalty. Unlike SuperSparc, there is no penalty for placing an address in a register and then using it in the next cycle.

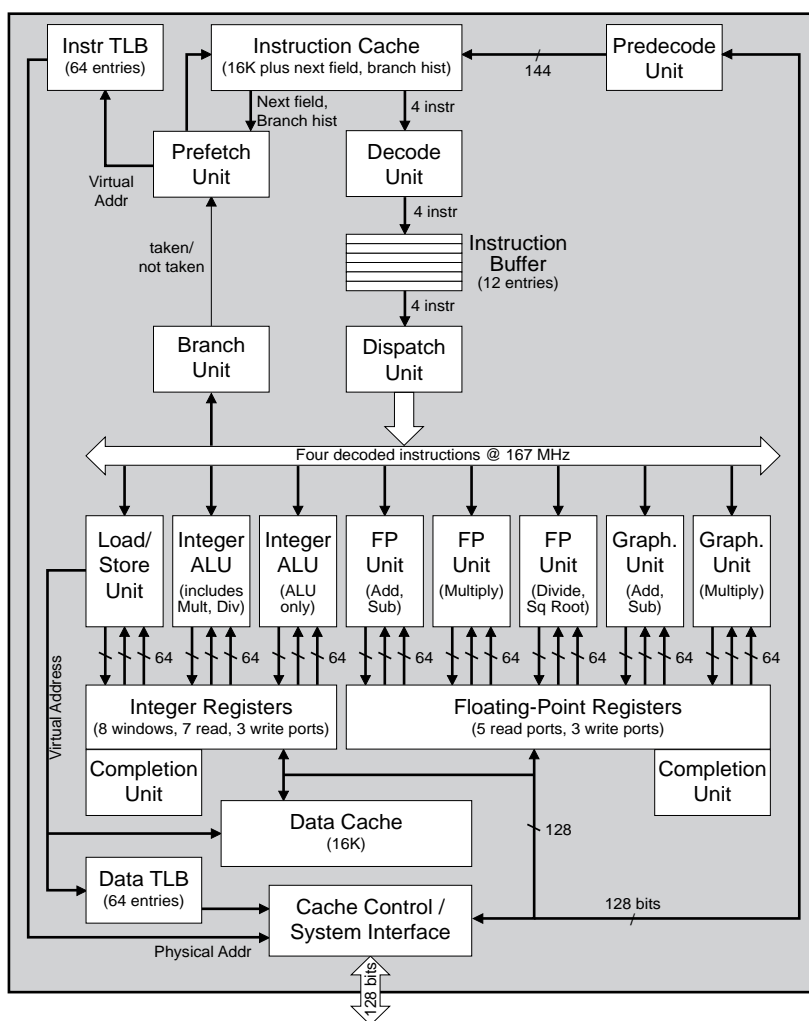


Figure 1. UltraSparc includes five floating-point/graphics units, but only two FP/graphics instructions can be dispatched at a time due to limited register-file ports.

## New Register File Saves Space

One criticism of SPARC's register windows is that, as superscalar processors require greater numbers of register ports, the area penalty of the large SPARC register file would balloon. Indeed, SuperSparc was forced to use a time-multiplexed register file (accessed twice per clock cycle to simulate additional read ports) to save die area at the cost of throttling the clock speed. To reach its frequency goal, Sun's design team could not take this route; instead, it came up with a much better solution.

To support UltraSparc's superscalar execution, the integer register file has 10 ports. Using a traditional design, each cell would have to be as wide as 10 metal traces (to bring in the select lines) and as tall as 10 metal traces (to bring out the data). Even with the tight metal pitches in EPIC-3, this works out to a cell size of roughly 400  $\mu\text{m}^2$ . Using these cells, a block of 144 registers of 64 bits each would be quite large.

The actual storage cell for a single bit, however, is much smaller than 400  $\mu\text{m}^2$ . In fact, Sun discovered enough room for eight bits under the metal grid required for the 10 ports. Because only one register window need be accessed at any given time, these eight bits, one from each window, can share the same set of ports. The only overhead is some selection lines and logic that, since the current window is known early in the cycle, can be set up ahead of time. With this new approach, the eight-window register file takes only 20% more area than a single-window register file.

As a bonus, the register file includes four sets of "global" registers, one each allocated to general use, MMU traps, interrupts, and other traps. These extra register sets, which are not a part of SPARC V9, provide scratch registers for traps and interrupts, allowing them to be handled without saving and restoring registers. The extra register sets are shared with the normal global registers, using the design technique described above.

The integer register file feeds two ALUs and the load/store unit. Both ALUs are 64 bits wide and can execute all arithmetic, logical, and shift operations, but only one has integer multiply and divide units. The integer multiplier calculates two bits per cycle and uses an "early out" algorithm that finishes faster for smaller operands. The divider calculates one bit per cycle.

The load/store unit accesses the 16K direct-mapped data cache. This cache is nonblocking: misses are sent to the external cache, but accesses to the primary cache can continue while the miss is in progress. Once the address is calculated (in the E stage), it is sent to the cache array, cache tags, and data TLB at the same time. The translated address is com-

## Price & Availability

Sun Technology Business (STB) has not announced pricing for the UltraSparc processor or system-logic chips. STB expects that the processor will sample in 1Q95 and reach volume production around mid-1995, with system logic on approximately the same schedule. For more information, contact STB (Sunnyvale, Calif.) at 408.774.8119; fax 408.774.8537.

pared with the physical address in the tags to determine if there is a hit; if so, the 64-bit data is available for use at the end of the C stage.

If a load or store misses the primary data cache, the translated address is queued for access to the external cache. Up to nine loads and eight stores can be queued. Because the primary data cache uses 16-byte sub-blocks, a single 128-bit access to the external cache will service a miss. Like the 21164, UltraSparc performs store compression via the store queue, although it is much more limited than the Digital chip. If two successive entries in the store queue refer to consecutive addresses, the two stores are combined into one entry, and only one 128-bit access is needed to complete the two stores.

The data TLB has 64 entries, each of which can map normal 8K pages or blocks up to 4M in size.

## FPU Includes Multimedia Support

UltraSparc implements a unique floating-point unit that Sun refers to as the floating-point/graphics unit (FGU). With five read ports on the FP register file, the processor can dispatch an FP load or store along with two instructions per cycle among the five function units: FP addition, FP multiplication, FP division/square root, graphics addition, and graphics multiplication. All FGU operations complete in three cycles for either single- or double-precision data. The only exceptions are FP divide and square root, which take 12 cycles for single-precision calculations or 22 cycles for double-precision.

As required by the SPARC V9 architecture, UltraSparc includes an FP register file with 32 double-precision registers, twice as many as SuperSparc and other V8 processors. This register file has three write ports, enough to retire two FGU operations and an FP load on every cycle.

F	D	G	E	C	N1	N2	N3	W
Fetch	Decode	Group	Execute	Cache Access				Write Back
Fetch instructions from cache	Decode instructions and place in buffer	Dispatch up to four instructions RF access	Execute int ALU Calculate mem addr	Start FPU D-cache, D-TLB, hit/miss	Second stage of FP calc	FP units complete	Resolve traps	Commit results to register file

Figure 2. UltraSparc uses a nine-stage pipeline, but the first five stages are similar to other superscalar RISC chips; the latter stages handle floating-point operations.

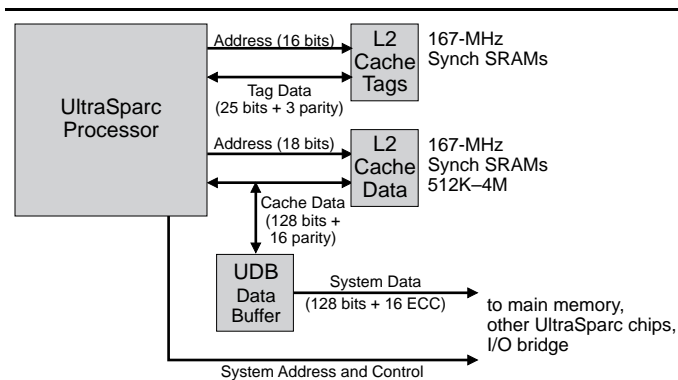


Figure 3. UltraSparc uses fast synchronous SRAM to implement a single-cycle external cache; the same 128-bit data bus is used to connect to the system through the UltraSparc Data Buffer (UDB).

The graphics units implement a set of new instructions that are not part of SPARC V9. From a software standpoint, these instructions are similar in concept to the multimedia extensions in HP's PA-7100LC processor (see *080103.PDF*) but are much more extensive. The new SPARC instructions operate on 8-, 16-, and 32-bit data types in parallel; for example, eight 8-bit results can be calculated by a single instruction.

These types of instructions are ideal for audio and video data that is typically represented in 16 or fewer bits. Sun expects the new chip to decode MPEG-2 video at 30 frames per second and even perform H.320 video encoding at a level adequate for desktop videoconferencing—all without external hardware.

As multimedia support becomes a requirement for high-end desktop systems, the ability to handle these applications without additional hardware will eliminate the need for expensive add-in cards. Starting with a complex high-end processor, the cost of adding multimedia function units is minimal; Sun estimates that the graphics units add about 3% to UltraSparc's die area.

### High-Speed System Interface

UltraSparc uses a single 128-bit data bus to connect to the rest of the system, as Figure 3 shows. This bus connects directly to the external cache SRAMs, which are controlled by the processor chip. These SRAMs must be synchronous parts that operate at the CPU frequency; slower parts are not supported. Sun claims to have multiple sources for 167-MHz synchronous SRAMs and believes that procuring 200-MHz SRAMs will not be a limiting factor to UltraSparc's performance. Given the processor's schedule, volume shipments of these fast SRAMs are not needed until the middle of next year.

The external cache array requires a 144-bit interface to store 128 bits of data plus 16 parity bits. Although ECC would not have required additional storage, it would have forced read-modify-write operations for most

stores; hence the use of byte parity instead. The external cache tags have a separate 28-bit bus and use the same synchronous SRAMs as the data array.

The external cache is a unified direct-mapped cache that can range in size from 512K to 4M. It is not optional. By including a large single-cycle cache, Sun hopes to match up well against HP's PA-RISC processors, which typically include large external primary caches. These caches allow HP processors to deliver good performance on fairly large programs, while SPARC processors often see a significant performance degradation once the small on-chip caches overflow. UltraSparc mimics the HP design but retains significant on-chip cache, allowing it to use a single external cache and reducing pin count, system cost, and system design complexity.

UltraSparc requires an external chip, dubbed the UltraSparc data buffer (UDB), to buffer the cache from the system bus. This chip performs two primary functions. First, it contains queues that buffer the high-speed cache from the slower system bus, which can run at either one-half or one-third of the CPU frequency. The UDB chip also converts data from the byte parity used for the cache to the more reliable ECC used by the system bus. Addresses and control signals for the 128-bit system bus are generated by the UltraSparc chip.

At 83 MHz, the system bus provides a 1.3-Gbyte/s peak transfer rate to main memory. It uses a split-transaction protocol, allowing memory accesses to overlap. UltraSparc supports glueless MP for up to four processors sharing the same system bus, maintaining cache consistency across the bus without any processor interface chips. Sun is developing a memory controller for the UltraSparc system bus as well as a bridge to MBus but has not announced these products.

### Sun, TI Jettison BiCMOS

Texas Instruments will fabricate UltraSparc in a version of its EPIC-3 process (see *080504.PDF*). Unlike the BiCMOS SuperSparc, the new chip uses 0.5-micron gates and four layers of metal in a pure CMOS process. TI says that bipolar transistors do not scale well to half-micron processes and provide little performance gain. The metal pitches are much tighter than in the EPIC-2B process used for SuperSparc, allowing more circuitry to be packed onto the die.

EPIC-3 is also used for TI's MVP DSP, which is currently sampling and is scheduled to enter volume production soon. The DSP chip will let TI debug its process before beginning production of UltraSparc, hopefully avoiding the production problems created when TI used SuperSparc to bring up EPIC-2B. But MVP uses a looser 0.55-micron, three-layer-metal version of EPIC-3, so TI is hoping for a 10% shrink before building UltraSparc chips, an aggressive move that could cause trouble.

This shrink is needed to keep the die size within the reticle limit; even in 0.5-micron CMOS, the UltraSparc die measures 315 mm<sup>2</sup>, larger even than the 21164. This drives the estimated manufacturing cost to \$350, according to the MDR Cost Model (see *081203.PDF*), making it costlier by far than any single-chip microprocessor except for the 21164. The UltraSparc chip, which uses 3.8 million transistors, will burn an Alpha-like 30 W at 167 MHz, according to Sun. The wide system interface and high power require a 521-pin PGA.

### Can UltraSparc Save Sun?

Ever since the SuperSparc clock-speed fiasco, Sun has lagged other workstation vendors in performance. Amazingly, the company has clung fiercely to its leadership share in the overall market, but the company has lost some ground at the high end as buyers turned to faster machines, primarily from HP.

UltraSparc promises to remedy the performance problem. Sun continues to claim that the chip will deliver 275 SPECint92 and 305 SPECfp92 at 167 MHz. This expected integer performance would lag the 21164 by about 20% but should put UltraSparc roughly on par with next-generation RISC processors from the MIPS, PA-RISC, and PowerPC camps.

If Sun can deliver on its promises, it spells bad news for other workstation vendors, taking away their chief weapon against the market leader. Sun still has a larger installed base, more applications available, and lower prices than the competition; with a competitive high-end processor, Sun should be able to defend its current share from HP, IBM, and Silicon Graphics. But if UltraSparc's performance is significantly below target, the SPARC chip could get gobbled up by Intel's P6.

Unfortunately, UltraSparc does not address Sun's biggest problem, which is the collision of high-end PCs and low-end workstations (see *0804ED.PDF*). UltraSparc systems are likely to be quite expensive, due to the high cost of the chip, the large synchronous cache, and the high-speed bus interface. UltraSparc does not address the low end; Sun is apparently counting on MicroSparc-3 to meet this need (see sidebar).

The new chip is appropriate for high-end workstations and servers, where cost is not a significant issue. Its glueless multiprocessor capability will allow Sun to offer relatively inexpensive processor upgrades, furthering its aggressive push into the MP market; the com-

## MicroSparc-3 Revamped

According to Sun's initial SPARC roadmap (see *070404.PDF*), the company planned to upgrade its low-end systems in 2H95 with MicroSparc-3 (MS-3). Originally, this processor was expected to deliver only a slight increase in performance over MicroSparc-2, reaching about 100 SPECint92. As Pentium has become more of a threat to Sun, the company has been forced to revise its plans for MS-3.

The low-cost chip now has a design target of more than 200 SPECint92. This performance will allow MS-3 to outperform, and thus obsolete, SuperSparc and SuperSparc-2, taking over the midrange as well as the low end of Sun's line. To achieve this performance goal, however, Sun has been forced to significantly redesign the MicroSparc core, delaying MS-3 about a year, to 2H96. At this time, Intel's mainstream processors will be 100–150 MHz Pentiums, which will not come close to the new performance goal for MS-3.

There are no details available as to how Sun plans to reach the new goal; a superscalar core with clock speeds approaching 200 MHz will be required. We expect that MS-3 will continue the MicroSparc design style of including memory and bus interfaces on the processor chip. Like its predecessors, MS-3 will include an SBus interface. Sun is investigating providing a PCI interface as well, either as a dual-mode interface or as a different version of the chip. PCI would allow Sun to utilize low-cost PC peripherals and compete better with Pentium-based PCI systems. MS-3 is also likely to be bi-endian, supporting Windows NT as well as Solaris.

pany already ships more MP systems than any other RISC vendor. The new graphics capabilities should help UltraSparc systems compete against popular multimedia boxes from Silicon Graphics and HP.

Sun is taking a chance with its somewhat premature announcement: if UltraSparc fails to fulfill its performance or schedule goals, Sun will again be the subject of well-deserved ridicule for overcommitting itself. The company swears it has done much more extensive timing simulation and other tests than it has in the past and is very confident that its chip will meet or exceed the current performance estimates. If UltraSparc delivers, it will prove that there is nothing wrong with the SPARC architecture that a good implementation can't fix. ♦