

# TurboSparc Offers Low-End Upgrade

## New Fujitsu Chip Plugs into MicroSparc-2 Systems for Performance Boost

by Linley Gwennap

Eyeing an installed base of underpowered SparcStation 5 workstations approaching half a million units, Fujitsu has developed a new CPU as a field upgrade for those systems. The TurboSparc processor is also appropriate for new low-cost workstations. Following the MicroSparc tradition, TurboSparc is a highly integrated 32-bit processor with on-chip cache, memory, and SBus interfaces, as Figure 1 shows. The performance improvement is modest, however, moving users from the equivalent of a 486 up to a midrange Pentium. TurboSparc is shipping now, both as a standalone device and in an upgrade kit.

Fujitsu estimates the 170-MHz TurboSparc will deliver 3.5 SPECint95 and 3.0 SPECfp95 (base), although current systems need more compiler tuning to achieve these marks. If achieved, these scores would represent roughly twice the integer performance of a 110-MHz MicroSparc-2 (MS-2) and about 50% better floating-point performance. The gain over slower versions of MS-2 will be even greater. This boost makes TurboSparc attractive as a field upgrade. This performance, however, is well below that of a Pentium Pro, PowerPC 604e, or even a high-end Pentium, all of which sell for about the same price as the \$499 TurboSparc.

### New CPU Core Boosts Clock Speed

Although Fujitsu builds and sells MicroSparc-2, that chip design was developed and is owned by Sun. For its upgrade chip, Fujitsu chose to develop its own CPU core, deploying a small team in San Jose (Calif.). For cost and time-to-market reasons, the TurboSparc team chose a simple scalar CPU design based on the 32-bit SPARC v8 architecture. In this regard, TurboSparc is similar to MS-2 but quite different from the superscalar 64-bit UltraSparc.

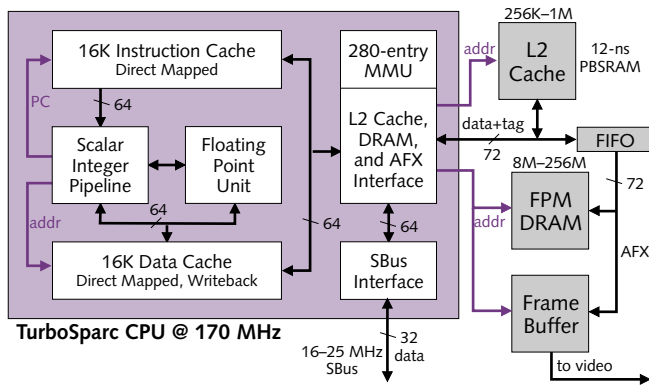


Figure 1. TurboSparc integrates nearly all the memory and system interfaces needed for a complete low-cost workstation.

Because of the similar throughput of the MS-2 and TurboSparc cores, the higher clock speed of the latter part provides a large part of its performance gain. TurboSparc clocks 55% faster than its predecessor. Some of this speed advantage is due to a gate shrink from 0.4-micron CMOS to 0.35-micron.

Most of the speed gain, however, comes from a new pipeline. Whereas MicroSparc-2 uses the classic five-stage RISC pipeline (see 071501.PDF), TurboSparc extends it to six stages for integer instructions and eight for FP instructions. As Figure 2 shows, a new “resolve” stage, after the data-cache access, checks for any faults from the cache access before proceeding to the writeback stage. This new stage provides more time for the cache access to complete, avoiding a critical timing path. The elimination of branch folding, a feature found in MS-2, also eases the timing and helps achieve better clock speeds.

The other new stages handle FP instructions, forming an integrated integer/FP pipeline that simplifies the control logic. Most FP operations, including ALU ops and multiplication, have a four-cycle latency and thus are complete by the FR stage. FP divide and square root both process two bits per cycle and can take from 8 to 50 cycles to complete; the average is 21 cycles for single-precision operands and 35 cycles for double-precision operands. The FP unit also handles integer multiply and divide operations; multiplication has a seven-cycle latency, while division is the same as a single-precision FP division.

One unusual feature of the new CPU is its ability to handle branches without penalty and without prediction. The instruction cache provides two instructions per cycle, while the CPU consumes only one. On a branch, there is enough bandwidth to fetch from both the taken and non-taken paths until the branch condition is resolved, eliminating any branch penalties, as Figure 3 shows. This method avoids the complexities of branch prediction or branch

<b>Fetch</b>	(F) Fetch two instructions from I-cache
<b>Decode</b>	(D) Decode one instruction and read operands
<b>Execute</b>	(E) Execute integer operation
<b>Memory</b>	(M) Read tags and data from D-cache
<b>Resolve</b>	(R) Check tags, abort data if L1 cache miss
<b>Write</b>	(W) Write result to integer register file
<b>FP Resolve</b>	(FR) Complete FP ALU or multiply (4 cycle latency)
<b>FP Write</b>	(FW) Write result to FP register file

Figure 2. TurboSparc’s pipeline adds one stage to MicroSparc-2’s for integer instructions and two more stages for FP instructions.

folding and is similar to the technique used by QED in the RM7000 processor (see 101409.PDF).

As Figure 1 shows, TurboSparc contains 16K each of instruction and data cache. The instruction cache is the same size as in MS-2, but the data-cache size is doubled relative to that chip and is now a write-back cache rather than write-through. These caches are fairly small compared with those of recent microprocessors and are direct mapped, further reducing their hit rate. The MMU is SPARC v8 compliant. It includes a 256-entry TLB for data translations, four more entries for large data pages, a four-entry instruction TLB, and a 16-entry I/O TLB. The total number of TLB entries is four times more than in MS-2, eliminating another performance bottleneck for some applications.

### DRAM Interface Now Supports L2 Cache

The second major performance enhancement from MS-2 is the addition of an external level-two cache. MicroSparc-2 is completely limited to the paltry 24K of cache that is available on the chip, relying on direct access to external DRAM for all other memory references. This shortfall creates a significant performance degradation on SPEC95 as well as on many workstation applications.

TurboSparc supports 256K–1M of direct-mapped L2 cache. This cache has the same 32-byte line size as the on-chip caches and uses a write-through policy. It runs at one-half the core CPU speed, requiring 12-ns pipelined burst SRAMs for the 170-MHz processor. An access that hits in the L2 cache stalls the CPU pipeline for 12 cycles. With a 72-bit interface, it takes four cycles to return a full cache line. A one-third-speed cache is also supported, but this choice will reduce performance.

Using  $\times 36$  SRAMs, the cache tags are stored side by side with the data. Each cycle, the 72-bit interface returns 64 bits of data, 2 parity bits, and 6 tag bits. After the first two accesses, the complete 12-bit tag can be assembled and checked to see if the access has hit in the cache. This design eliminates the need for separate external cache tags (or internal tag storage). The tags cannot be checked until two cycles after the data is received, but the new R stage allows enough time to abort the writeback if the tag check fails.

DRAM accesses are started in parallel with L2-cache accesses and are aborted if the L2 cache hits. This strategy reduces the duration of a pipeline stall on an L2 cache miss by overlapping the DRAM access. The DRAM interface is configurable for page-mode DRAM of various speeds but does not handle more advanced memories such as EDO or SDRAM. These memory types are not supported in the older SparcStation systems, so Fujitsu did not bother to add them, keeping the design as simple as possible to speed its completion. With 60-ns DRAM, the CPU pipeline stalls for 24 cycles on an access to main memory.

The memory controller supports up to eight banks of 32M each, or 256M maximum. This limitation is similar to that of MicroSparc-2.

### Price & Availability

The 170-MHz TurboSparc chip is available now at a list price of \$499 in quantities of 1,000. The 160-MHz TurboSparc upgrade kit costs \$1,500 in quantities of one. To get more information on TurboSparc, contact Fujitsu (San Jose, Calif.) at 800.866.8608 or access the Web at [www.fujitsumicro.com/sparcupgrade/sparcmicro.html](http://www.fujitsumicro.com/sparcupgrade/sparcmicro.html).

### Sun AFX Graphics Supplements SBus Interface

Recently, Sun has added a new graphics interface called AFX. These graphics cards reside on the main memory bus instead of the pedestrian SBus, significantly improving bandwidth. AFX requires adding only a few extra control signals to the existing memory bus, which Fujitsu has done in TurboSparc. This change allows an end user to plug an AFX graphics card into a system that has been upgraded with TurboSparc.


Like MS-2, the new processor supports SBus directly on the chip. The SBus operates at 16–25 MHz, typically one-eighth of the CPU clock speed. Up to six bus masters can be connected to the 32-bit SBus. TurboSparc is fully compatible with the Macio and Slavio chips that supply basic I/O functions in the SparcStation 5 and similar systems.

The integrated memory and bus interfaces make multiprocessor configurations impossible. This fact simplified some aspects of the TurboSparc design. The cache does not support multiprocessor coherency, for example, and the CPU core does not execute certain SPARC v8 instructions for MP synchronization.

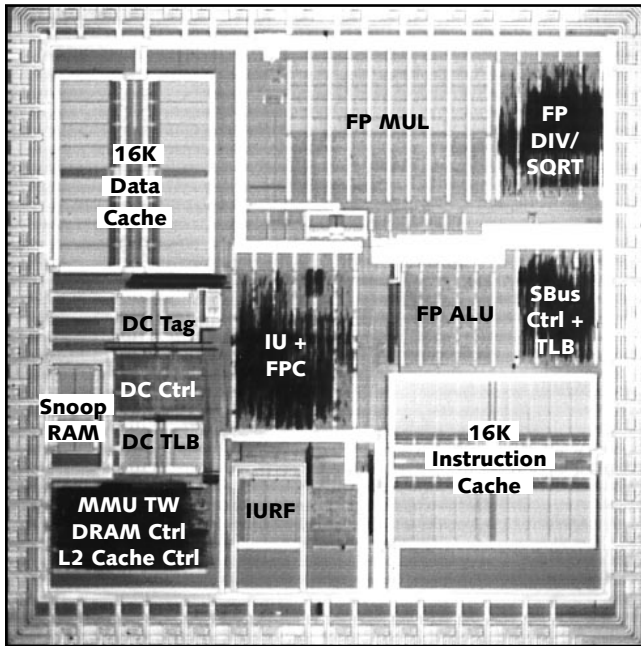
### Manufacturing Cost Shrinks

TurboSparc is built in Fujitsu's 0.35-micron four-layer-metal CS-60ALE, keeping the die size down to 132 mm<sup>2</sup>, relatively svelte for a processor with so much integrated system logic. MicroSparc-2, by comparison, weighs in at 233 mm<sup>2</sup> using the 0.4-micron CS-55 process (see 090905.PDF).

Although the gate shrink is minor, a bigger gain is seen in the metal layers: the CS-55 metal layers are from a 0.5-micron process, whereas CS-60ALE is a complete 0.35-micron process. Thus, the TurboSparc die size is about what we would expect if the MS-2 die was shrunk to the same 0.35-micron process.



**Figure 3.** On a branch, TurboSparc fetches from both the nontaken and taken paths. By cycle 4, the branch is resolved, and either instruction can be executed without penalty. (— indicates stall)



**Figure 4.** Fujitsu's TurboSparc combines a scalar SPARC CPU with a complete set of system interfaces. Sporting 3 million transistors, the die measures 11.5 × 11.5 mm in a 0.35-micron four-layer-metal CMOS process.

The transistor count of TurboSparc is slightly higher: 3.0 million, compared with 2.3 million for MS-2. Most of the increase is due to the extra 8K of cache, with the remainder in the L2 cache and AFX interfaces. The CPU core has about the same number of transistors as in MS-2. Because the physical layout of MS-2 is rather loose, Fujitsu was able to pack more transistors into the same relative die area. Figure 4 shows the TurboSparc die.

TurboSparc is packaged in a 416-contact plastic BGA. The plastic BGA saves cost compared with the old-style ceramic PGA used for MS-2, despite the extra 95 leads required by the new interfaces. Combining the savings from the plastic package and the smaller die, the MDR Cost Model estimates the cost of building TurboSparc at about \$50, a third less than the cost of MS-2. The PowerPC 603e and Pentium chips deliver similar performance at a build cost of \$30–\$40, but these chips cannot match the integrated system logic of TurboSparc.

Like MS-2, the new chip runs at 3.3 V. The maximum power dissipation is 9 W at 170 MHz, matching the maximum power of MS-2 despite the significantly higher clock speed. Fujitsu paid more attention to moderating power in the new design, adding some gated clocking to keep the chip within the same thermal envelope as its predecessor.

### Module Upgrades MicroSparc-2 Systems

Designing a TurboSparc upgrade for MicroSparc-2 systems was no easy task. To gain the necessary performance boost, TurboSparc adds a secondary cache, but this cache is obviously not present on a MicroSparc-2 motherboard. Thus,

Fujitsu has designed a module containing a TurboSparc processor and 256K of L2 cache implemented with two 32K×36 SRAMs. The module is a small PC board with a pin-grid array on the bottom that plugs into an MS-2 socket.

The PC board has an odd “L” shape to fit into the existing SparcStation 5 design, carefully avoiding all obstacles. Fujitsu believes this board will also fit into most other MS-2 workstations. The TurboSparc processor has its own fan mounted on the heat sink to ensure adequate cooling. An on-board voltage regulator delivers the extra current needed by the SRAMs and buffers.

For yield reasons, the modules use a 160-MHz TurboSparc; the company is saving the 170-MHz parts for customers buying standalone chips. Fujitsu is marketing the TurboSparc module directly to end users in an upgrade kit that contains the module, documentation, and an extraction tool to remove the MS-2 chip. The kit, which retails for \$1,500, also contains a new PROM with the appropriate boot code for the new processor.

### From Woeful to Weak

In addition to the upgrade kits, new systems from several small SPARC system vendors are using TurboSparc. Sun, however, has been conspicuously absent among vendors adopting the new chip. This oversight is surprising: with MicroSparc-2 delivering just 1.4 SPECint95 and 1.9 SPECfp95 (base), the performance of Sun's low-end workstation line is quite woeful by current standards.

TurboSparc offers a significant boost but still barely matches the SPEC performance of a good 120-MHz Pentium box on both integer and floating-point code. A high-end Pentium PC will run rings around a TurboSparc workstation on many technical applications while costing less than half as much. Fujitsu argues that SPEC exaggerates the performance difference because Intel's SPEC results rely on far more compiler tuning than Fujitsu's estimates.

In any case, the audience for TurboSparc remains die-hard SPARC advocates who need compatibility with a large installed base of SPARC hardware and software. There are plenty of Sun-only shops around to which TurboSparc is appealing, both as a field upgrade and in new systems. Even some of these diehards, however, are eyeing the low cost of x86-based systems. Users not tied to SPARC will see little attraction in TurboSparc systems.

Sun's longer-term solution for this price point is UltraSparc-2i (see [101301.PDF](#)). This processor is slated to exceed the performance of Intel's P6 chips while including a set of integrated system logic similar to TurboSparc's. US-2i cannot provide a field upgrade for MicroSparc-2 systems, however, and is not due to appear in systems until 4Q97, a year from now. Until then, TurboSparc will have to power low-cost SPARC systems from Sun and others. But the new chip may not have enough power to get these systems safely through the wake created by the Pentium and P6 workstations now entering the market. ▣