MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 6 NUMBER 7 MAY 27, 1992

TI and Cypress/Ross Battle for SPARC Leadership

SuperSPARC Has Inside Track at Sun; hyperSPARC Hoping for Upset

By Brian Case and Michael Slater

At long last, both Texas Instruments and Cypress/Ross have formally announced their superscalar SPARC implementations. TI's SuperSPARC (formerly know as Viking) was designed primarily by Sun, and the first chips were fabricated last September; TI and Sun claim the chip is now production-ready, albeit at a lower clock rate than originally targeted. This processor is at the heart of Sun's new SPARCstation 10 (see p. 11), and Sun's 600MP multiprocessor server line is being upgraded to use SuperSPARC modules.

SuperSPARC's challenger from Cypress subsidiary Ross Technology, code-named Pinnacle and now officially "hyperSPARC," was designed without backing from Sun and takes a different approach than SuperSPARC. It is not as aggressive in its issue capabilities or level of integration (it requires a minimum four-chip set), but it strives for higher performance by using a faster clock rate. The hyperSPARC chip set began fabrication only this week, so there won't be any silicon to test until late June. Samples have been promised for the third quarter, but it will be a challenge to get such a complex chip set debugged on this schedule; volume production is unlikely before year-end.

Descriptions of both processors have been public for some time, and we will not repeat here details that have appeared in previous articles. (See µPR 12/4/91 for an overview of SuperSPARC, and 3/25/92, p. 15 for Pinnacle/hyperSPARC.) Some new information has come to light as part of the formal announcements, however. Perhaps of greatest interest are the clock rates: SuperSPARC initially will be available at lower clock rates than expected—33 and 40 MHz—while Cypress/Ross is shooting for a starting clock rate of 66.7 MHz. Doing a balancing act between performance and chip yield, Sun will use a 36-MHz SuperSPARC in its low-end SPARCstation 10.

Of course, both companies expect higher clock rates in the future. TI is promising 45- and 50-MHz parts by the end of this year, with a 75-MHz version in 1993. Cypress/Ross is promising 80- and 100-MHz versions in the future.

The performance of systems based on the two initially available implementations should be in the same ballpark, and a true comparison must await measured performance figures and shipment dates for hyperSPARC. While hyperSPARC has the potential to deliver higher performance than SuperSPARC, this will depend on when Cypress/Ross gets the chip set into production, what clock rates it actually achieves, and whether SuperSPARC makes it to 50 MHz by the time hyperSPARC is in production.

Performance

While these processors offer an integration advan-

Continued on page 6

In This Issue

SuperSPARC and hyperSPARC

Continued from front page

tage over previous-generation SPARCs, increased performance is the key objective. Table 1 summarizes the performance data provided by TI, Sun, and Cypress. For reference, SPARCstation 2 performance is also shown. At 36 MHz, SuperSPARC doubles the SPARCstation 2 SPECint92 performance and has 2.3 times its SPECfp92 performance. At 40 MHz, SuperSPARC achieves 2.5 times the SS-2's integer performance and 3 times its floating-point performance.

Note that the SPARCstation 10 ratings are for the latest, not-yet-released compilers, so this actually overstates the performance improvement that can be expected by users with existing binaries—or even today's production compilers. In fact, the compilers used for the SPEC benchmarks are from Apogee Software—the compilers from Sun's own SunPro division don't perform as well. Apogee customized its compilers to optimize for the SuperSPARC pipeline.

If Cypress/Ross' simulations are correct and if they are able to achieve their target 66.7 MHz clock speed, hyperSPARC's performance should exceed that of a 36-MHz, no-external-cache SuperSPARC and roughly match that of a 40-MHz SuperSPARC with a 1-Mbyte external cache. If Cypress/Ross is able to increase its clock rate to 80 MHz before TI gets to 50 MHz, it may be able to maintain a performance lead. It seems likely, however, that hyperSPARC will occupy a middle ground in performance, and perhaps in price as well, between a SuperSPARC without external cache and a SuperSPARC with external cache. (See p. 12 for performance comparisons with other architectures.)

Table 2 compares computation latencies for SuperSPARC and hyperSPARC. SuperSPARC has significantly lower latencies in clock cycles. When comparing a 40-MHz SuperSPARC to a 66.7-MHz hyperSPARC, however, SuperSPARC is actually slower on floating-point add and multiply and comparable on divide and square root because of its lower clock rate. Both processors have fully pipelined floating-point units (i.e., an operation can be started every cycle) for single- or double-precision add or multiply. Integer multiply and divide are relatively slow on hyperSPARC, which uses an iterative implementation, while SuperSPARC implements integer multiply and divide using the floating-point unit's computational resources.

Function	SuperSPARC		hyperSPARC	
	Clock Cycles	ns at 40 MHz	Clock Cycles	ns at 66.7 MHz
Integer Multiply	4	100	17	255
Integer Divide	15	375	36	540
FP Add (SP/DP)	3	75	3	45
FP Multiply (SP/DP)	3	75	3	45
FP Divide (SP)	6	150	10	150
FP Divide (DP)	7	175	14	210
FP Sq. Root (SP)	8	200	13	195
FP Sq. Root (DP)	10	250	19	285

Table 2. Latencies for integer and floating-point arithmetic functions, assuming normalized operands and results.

SuperSPARC

The single-chip SuperSPARC microprocessor has been given the TI part number TMS390Z50; the companion external, second-level cache controller, with on-chip tags and off-chip data RAM, is called the 390Z55. Power consumption for the 390Z50 is 5.5 W typical (8 W maximum) at 33 MHz, which is somewhat less than the R4000 with its smaller transistor count but higher clock rate.

TI has shown a variety of packaging options for SuperSPARC. Single-processor MBus modules, using either a lone 390Z50 or a processor/cache controller/SRAM chip set, initially will use conventional PGA packages for the VLSI devices and SOJs (small-outline, surface-mount plastic packages) for the RAMs and interface logic. Figure 1 shows a single-processor module with external cache.

The first dual-processor modules from TI will be built using TAB (tape automated bonding) for the processor and cache controller. Sometime in 1993, TI plans to begin producing dual-processor MBus modules using one multichip, silicon-substrate module for each processor/cache subsystem. Prototypes of the silicon substrate module have already been fabricated. Currently, sili-

Processor	SuperSPARC						hyperSPARC			7C601
Source	TI Estimates				SPARCstation 10		Cypress Estimates			SPARC-station 2
Clock Rate (MHz)	33	40	45	50	36	40	66.7	80	100	40
External Cache	—	—	1M	1M	—	1M	256K	256K	256K	64K
SPECmark89	40-50	50-60	60-70	80+	—	—	70	85	101	25.0
SPECint92	—	—	—	—	44.2	52.6	55	65	77	21.8
SPECfp92	—	—	—	—	52.9	64.7	64	75	89	22.8

Table 1. SuperSPARC and hyperSPARC performance claims. The only measured numbers are those for the SPARCstation 10 and SPARCstation 2.

con-substrate modules are too expensive for all but the very high end of the market, but TI believes that it can significantly lower their cost. Modules have potential advantages in making higher clock rates possible, but the system is currently limited by speed paths within the processor chip, not by I/O speeds.

The bus of the 390Z50 operates in one of two modes: MBus mode, for operation as a single-chip subsystem, or VBus mode, for use with the external cache controller. In MBus mode, the on-chip data cache operates in write-back mode with a write-allocate policy. In VBus mode, the on-chip data cache uses a write-through protocol so the external cache can maintain inclusion (data in the internal cache is also always in the external cache). Inclusion allows the 390Z55 cache controller to implement the system-wide cache coherency protocol without explicit help from the microprocessor.

Figure 1. A SuperSPARC MBus module with external cache.

The 390Z50 implements the new instructions required by version 8 of the SPARC architecture, plus one additional debug instruction. The integer multiply and divide instructions are executed by the floating-point hardware and take 4 and 15 cycles, respectively. Actually, integer divide is executed in hardware only if the 64-bit dividend (register pair Y and register source 1) has no numerically significant bits beyond bit 51 (i.e., bits 52–63 must be either all ones or all zeros). If significant bits occur above bit 51, a trap is taken and the divide is emulated in software.

SuperSPARC has one instruction not in the SPARC version 8 architecture specification: SIGM (signal emulator). This instruction either executes as a no-op or causes a trap, depending on the state of a bit that can be set only by the JTAG interface. Thus, system software can place SIGM instructions at interesting points, but debugging action will not take place unless JTAG boundary scan is being used and enables the trap.

Some other powerful software debugging facilities are provided in the form of traps that are triggered on a code address or a data address, or by underflows in cycle and instruction counters. A single breakpoint register is available for either a code or data address, and the address can be either virtual (32-bit) or physical (36-bit).

For counting instructions and cycles, a single 32-bit register programs two 16-bit down counters. Larger capacity registers can be simulated by keeping track of underflows in software. By dividing the instruction count by the cycle count, the dynamic native MIPS value can be measured.

SuperSPARC Cache Controller

While the 390Z55 cache controller has on-chip cache tags for 1 MB of cache, it requires external 128K × 8 or × 9 SRAMs for the cache data store. These 1-Mbit SRAMs are currently available from Sony, Paradigm, and AT&T. The cache controller is expected to be available from TI only as part of a complete SuperSPARC MBus module, although TI says exceptions may be made for special customers.

The 390Z55 serves two purposes: it provides a second-level cache controller and a non-MBus interface. XBus is the generic name given to the non-MBus interface, and it is designed to allow a variety of bus protocols to be used depending on the needs of the system. XBus is not, in itself, the system bus; it is designed to connect to an external bus interface chip that implements the system bus protocols. Sun is rumored to be developing a high-end multiprocessor server, code-named Dragon, that uses the XBus to provide an interface to the DynaBus, a high-end multiprocessor bus designed at Xerox PARC. Both the XBus and the DynaBus use GTL (Gunning Transceiver Logic) levels, which have a smaller signal swing than conventional CMOS signals and are designed to support higher-speed operation.

The XBus consists of 73 bused wires plus 3 point-to-point wires per bus interface chip; up to four bus interfaces can connect to a single 390Z55. This implies that a single 390Z50/390Z55 pair can give rise to four separate system buses. Clearly, Sun and TI have designed SuperSPARC to be used in very large multiprocessor systems.

When used in XBus mode, the pins of the MBus connector are redefined as XBus signals. The XBus signals cannot be directly bused between processors, how-

Price & Availability

Both TI and Cypress have given very limited pricing information. The only figure TI will provide is "under $400" for the 33-MHz version, in 1993, in quantities of 10,000. Given the difficulties TI has had with increasing the clock rate, the higher speeds are likely to be much more expensive; TI will probably have a surplus of 33-MHz parts as a side effect of building enough chips to meet Sun's demand for 36 MHz, 40 MHz, and faster parts. TI has not provided any 1992 pricing, pricing for modules with external cache, or pricing for higher clock rates. Production of 45- and 50-MHz modules is promised for the fourth quarter.

TI claims to be in production on 33- and 40-MHz versions, but it is currently making them available only to Sun and a few other selected customers, which has caused considerable dissatisfaction among smaller makers of SPARC-based systems. General availability of the 33- and 40-MHz versions is planned for the fourth quarter, with one catch: TI expects all customers to purchase a $50,000 "system design starter kit" (promised for the third quarter), which includes a simulator, Verilog model, a JTAG-based testing tool, data sheets, a user's guide, technical support by e-mail, and several chip samples. This approach is designed to minimize TI's support burden, but it raises the price of using the chip to unprecedented levels. With Sun likely to consume the vast majority of the volume, TI apparently feels no need to be especially accommodating to a broad range of customers. If TI had put aside just a few dozen chips and made them available to all participants in the SPARC market, it could have gone a long way toward realizing the stated goals of Sun's "Silicon Partners" program.

While TI is giving a low-ball price, Cypress is quoting an unrealistically high $3500 in quantities of 100, for a 66.7-MHz uniprocessor module with 128 Kbytes of cache. Cypress acknowledges that volume pricing will have to be much lower for the device to have any chance of success, but it would not provide any other pricing information. Samples of the 66.7-MHz version are promised for the third quarter, with production in the fourth quarter. Faster versions will follow at an unspecified later date.

Given the only released pricing information of under $400 in 10K quantities for the 33-MHz SuperSPARC, compared to $3500 for 100-piece lots for the 66.7-MHz hyperSPARC, it is impossible to determine how the volume pricing will compare. Furthermore, hyperSPARC should be compared to the SuperSPARC module with second-level cache, for which TI has not released any pricing. Perhaps both vendors will see fit to release 1000-piece pricing for the full range of configurations once they gain some experience with the yield of the chips.

Texas Instruments, P.O. Box 809066, Dallas, TX 75380; 800/336-5236 ext. 3990, or 214/995-6611 ext. 3990.

Cypress Semiconductor, 3901 North First St., San Jose, CA 95134; 408/943-2600.

ever, so dual-processor modules will not support the XBus interface. The primary application envisioned for dual-processor modules will be in four-processor desktop systems, not in high-end servers.

In MBus mode, the 390Z55 can be used with either no external cache or a cache of 1 MB. In XBus mode, cache sizes of zero, 512 KB, 1 MB, or 2 MB can be used.

When the 390Z55 is in XBus mode, a separate 12-pin "boot bus" is provided. This bus is intended primarily for accessing bootstrap EPROMs, but with external controller logic, it can also support other devices such as diagnostic I/O hardware.

Pinnacle Becomes hyperSPARC

In an apparent attempt to beat Sun/TI at the microprocessor naming game, Cypress/Ross has chosen "hyperSPARC" as the official name for its superscalar SPARC microprocessor, previously known as Pinnacle. The hyperSPARC chip set consists of a 7C620 processor, a 7C625 CMTU (cache controller/MMU), and either two or four 7C627 CDU (cache data unit) SRAMs. The processor uses about 1 million transistors, the CMTU has about 700,000 transistors, and the half-megabit CDUs use 2.2 million transistors each. Two of the synchronous CDU SRAMs yield 128 KB of combined cache, while four CDUs give 256 KB.

Just as with SuperSPARC, hyperSPARC implements the current SPARC standards: version-8 instruction set architecture, reference MMU, and level-2 MBus (MOESI cache coherency). HyperSPARC does not implement an XBus interface.

The hyperSPARC intra-module bus connecting the processor, CMTU, and CDUs uses 3.3-V logic levels to help accommodate high clock rates. The 7C625 is based on the previous-generation 7C605 cache-controller/MMU but adds the logic necessary to synchronize the 40-MHz MBus to the intra-module bus. The 7C620 has an on-chip, 8-KB instruction cache but no data cache; thus, the external cache serves as a second-level instruction-cache and a first-level data cache.

As with SuperSPARC, hyperSPARC processors will be available only in the form of MBus modules, which Cypress/Ross calls "SPARCore" modules. TAB packaging is planned, with one large heat sink covering all four (or six) chips. Four varieties of modules will be offered: single or dual processor, with 128K or 256K cache per processor. Samples of these modules are promised for the third quarter, with production planned before the end of the year. All modules are promised in 66.7-MHz speed grades initially, with 80- and 100-MHz versions in the future.

Not content with having hype only in the microprocessor name, Cypress/Ross has invented a new acronym for superscalar: MILE (multiple instruction launch and execute). In addition, they claim that hyperSPARC is

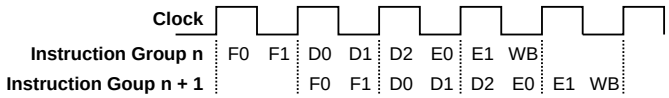

Figure 2. SuperSPARC pipeline operation.

superpipelined because the floating-point unit spreads the execution hardware over more than one pipeline stage. While it is true that this organization satisfies the definition of superpipelined, no such claim is made for other microprocessors that have a similar floating-point organization because the technique is so universal. In any case, hyperSPARC is not superpipelined in the same sense that the R4000 is superpipelined (i.e., hyperSPARC's integer unit is not superpipelined), and the relatively high clock rates of hyperSPARC (at least compared to SuperSPARC) are not the result of superpipelining.

SuperSPARC vs. hyperSPARC Pipelines and Execution Resources

SuperSPARC and hyperSPARC are both superscalar processors, but they differ significantly in the degree of simultaneous instruction issue and execution that they achieve. HyperSPARC implements execution resources essentially identical to those found in first-generation SPARC processors but uses them in a much more parallel manner. (The first DEC Alpha and RS/6000 implementations also use this approach.) SuperSPARC combines the hyperSPARC technique of making better use of traditional resources with a significantly more sophisticated integer unit offering dual ALUs.

HyperSPARC contains five functional units: load/store, branch/call, integer ALU, FP add, and FP multiply. In a single cycle, the issue logic can dispatch two instructions as long as they are executed in different functional units. For FP instructions, the issue logic can actually do better: two FP instructions can be sent to the FP instruction queue in the same cycle. Once in the queue, however, the instructions enter execution units—add/subtract or multiply—only one per cycle.

HyperSPARC realizes a benefit from its dual-issue capability only when the instruction stream contains pairs of instructions that are executed in separate units (the FP exception just mentioned notwithstanding). For example, hyperSPARC can issue an integer ALU instruction along with a branch. HyperSPARC's pipeline is a traditional—if somewhat longer—RISC pipeline: fetch, decode, execute, cache lookup/read, writeback, and cache update (write) are each allocated one cycle.

As shown in Figure 2, SuperSPARC implements a much more intricate pipeline where decode takes one and a half cycles, execute straddles a clock-cycle boundary, and writeback is allocated only half a clock cycle. Adding the address operands for a load or store, a task traditionally done in the execute stage, is performed in D2. An integer ALU operation using the results of a load cannot be executed in the same three-instruction group as the load, but a floating-point operation can use the results of a load instruction executed in the same instruction group.

A unique advantage of SuperSPARC is its ability to simultaneously issue and execute two integer ALU instructions; further, the two instructions may be dependent. This is permitted by the allocation of three ALUs to the two half-cycle E0 and E1 substages. E0 contains two ALUs that feed the inputs to the third ALU in E1. Thus, an add instruction can be executed in one of the ALUs in E0 (depending on register-file port assignments), and a subsequent add instruction that uses the result of the first can be executed in E1. Because E0 and E1 take only

Tsunami Previewed

Since it may be years before SuperSPARC is cheap enough for really high-volume, low-end systems, Sun has been developing another chip, code-named Tsunami, to meet the need for a highly integrated but low-cost SPARC microprocessor. Texas Instruments has already fabricated the Sun-designed chip, which is expected to debut (both as a chip and in systems) later this year. Tsunami is the outgrowth of a Sun/TI research project, code named Genie, that originally had even more aggressive functional integration goals, including an on-chip frame buffer controller.

The TI Tsunami contains an integer unit, an FPU, an MMU, a 4-KB instruction cache, a 2-KB data cache, an SBus interface, and a DRAM controller, in a 288-pin package. Both caches are direct-mapped. Since the focus is on low cost, it does not support an external cache or multiple processors. The chip is large (15 mm, or 590 mils, on a side) but contains only about 1 million transistors. The relatively large die size is a result of a more mainstream, 0.8-micron, two-level-metal CMOS process and the use of automatic design tools. The result is a chip with a modest development cost that can be shrunk if sufficient market demand materializes.

The TI Tsunami will not have spectacular performance; its goal is to cut system cost, not to boost performance. The processor is not superscalar, the floating-point hardware has reduced performance to save die area, and the caches are tiny. The clock rate target is rumored to be 40 to 50 MHz; at 50 MHz, the chip could provide near-SPARCstation-2-level performance at a much lower cost.

Fujitsu is rumored to be working on a higher-performance follow-on to Tsunami, but no information has been released on this project.

a half cycle, SuperSPARC appears to execute them simultaneously, in the same full cycle.

Both SuperSPARC and hyperSPARC treat floating-point loads and stores as if they were integer memory-reference instructions. Thus, both processors can issue and execute an FP computation and an FP memory reference simultaneously. HyperSPARC does have a slight advantage in being able to send two FP computation instructions to the FP queue simultaneously. While it cannot then send both to execution units in a single cycle, it may still realize a performance benefit simply by being able to get two FP operate instructions "out of the way" of subsequent integer instructions.

Conclusions

Despite the hype in hyperSPARC, the design tradeoffs it represents are just as valid as those made by SuperSPARC's designers, at least within the context of MBus modules. Cypress/Ross claims that its multi-chip, small die-size implementation has significant manufacturability advantages over the large-die implementation of SuperSPARC. For MBus modules, where a fixed, relatively generous amount of board space is available on the module, the ability to implement the processor/cache subsystem in one chip instead of four provides no advantage in physical size. The multi-chip implementation allows hyperSPARC to have a much larger first-level data cache.

In defense of the single-die SuperSPARC, it will have an advantage in lower-end systems as soon as manufacturability ceases to be an issue. Even though the MBus module has plenty of space, a single-chip module is potentially less expensive than multiple chips. In addition, the on-chip SuperSPARC first-level data cache can be accessed faster than the off-chip hyperSPARC data cache. Faster access compensates somewhat for SuperSPARC's smaller first-level data cache. SuperSPARC also has a wider range of configurations; with the external cache controller and 1M byte of cache SRAMs, it offers more than four times as much total cache, and it also has the XBus option.

The biggest validation of hyperSPARC's approach is that its performance may exceed SuperSPARC's, despite its less-aggressive superscalar implementation. Even though hyperSPARC is only a two-issue design and cannot simultaneously issue two integer ALU instructions, the small, simpler hyperSPARC chips have a clock-speed advantage that compensates for the three-issue capability and dual ALUs of SuperSPARC.

TI appears to be having significant problems getting SuperSPARC up to 50 MHz, as evidenced by the plan for baby-steps from 33, to 40, to 45, and finally to 50 MHz. In such a large chip, the problem is likely to be related to layout and could therefore take considerable effort to solve. Ironically, the current speed-limiting path has nothing to do with the superscalar capabilities—it is access to the data cache. This suggests that circuit design and layout problems, rather than microarchitectural issues, are responsible for SuperSPARC's limited clock rate. Cypress/Ross may also find its target clock rate difficult to achieve; TI and Sun were probably more optimistic about their clock rates before they had silicon to test.

In theory, hyperSPARC should have had a time-to-market advantage over SuperSPARC because of its use of smaller chips and less-exotic technology. In reality, however, it is at least six months behind, primarily due to Cypress' more limited resources. (It was also started later than SuperSPARC.) As a result, Cypress/Ross has missed its best window of opportunity—providing a product before SuperSPARC became available.

The real problem for Cypress/Ross is the future of hyperSPARC. Because Sun has so much invested in SuperSPARC, both financially and emotionally, it is clearly the first choice for Sun in computer system designs. SuperSPARC was designed primarily by Sun (with some help from TI), while hyperSPARC was developed independently. Putting SuperSPARC in as many systems as possible has software benefits as well: fewer ports of system software are necessary and optimizing compilers can be written to focus on generating good code for just one or two processor designs.

On the other hand, if Sun does not use hyperSPARC, Cypress may find it difficult to justify continuing its Ross division, and such a failure would be bad for SPARC overall. This factor alone may motivate Sun to find a place for hyperSPARC in its product line.

Cypress' recent layoffs, while not directly related to the Ross division, are symptomatic of the difficult position it is in. Cypress' SPARC sales are shrinking, and they probably will continue to shrink—at least for the near term. In the Sun product line, Cypress SPARC processors are used only in the SPARCstation 2, which will suffer declining sales as soon as the SPARCstation 10 begins shipping in volume, and in the 600MP servers, which are being phased over to SuperSPARC.

These shifts in Sun's product lines will make Cypress/Ross more dependent than ever on the Sun-compatible market, which remains small. If they can get hyperSPARC into production quickly, they may be able to take advantage of SuperSPARC's restricted availability.

Unless a significant consumer of SPARC processors other than Sun materializes, Cypress may have to hope for catastrophic problems with SuperSPARC. Given this situation, reports that Cypress may hedge its bet and begin producing DEC's Alpha microprocessor begin to make sense. ♦