

M I C R O P R O C E S S O R

www.MPRonline.com

◆ THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE ◆

SUN WEAVES MULTITHREADED FUTURE

Afara Acquisition Brings New Life to SPARC

By Kevin Krewell {4/14/03-01}

At the Sun Worldwide Analyst Conference, the company revealed plans to take its UltraSPARC processor family in a direction heavily dependent on multicore and multithreaded microarchitectures. Sun hopes its new direction will lead to a blade processor up to 15 times

faster than Sun's present blade server processor and a SPARC enterprise system processor 30 times faster than today's big iron server processors; a traditional single-threaded processor might expect to triple its performance during the same time frame. These claims are quite fantastic, and, if true, we could see these SPARC processors leapfrog over Intel's Itanium and IBM's Power families of processors. Sun's focus on applications that are highly threaded could be at the expense of single-threaded application performance.

Sun's control of both the software (Solaris) and hardware (UltraSPARC) facilitates its ability to embrace the radical multithreading concept, unlike the so-called Wintel alliance (Microsoft Windows and Intel processors), where two separate companies are involved. Radical multithreading will be efficient only if the operating system can efficiently support many concurrent threads. That's not to say everything is perfectly flexible at Sun. Solaris is considered by the software department to be the crown jewel and is fiercely protected to maintain stability and reliability. Sometimes, it's even harder to change the

software than the hardware. Still, Sun has publicly committed to this heavily multithreaded platform.

New SPARCs in Sun's Roadmap

Sun will be moving toward more-radical multithreading in future processors, as shown in Figure 1, starting gradually with a dual-core UltraSPARC IV in late 2003 and a dual-threaded blade processor (Gemini) in 2004. In 2005, though, Sun takes multithreading to the extreme in the Niagara processor, which will have eight processor cores per die and be able to support four threads per core.

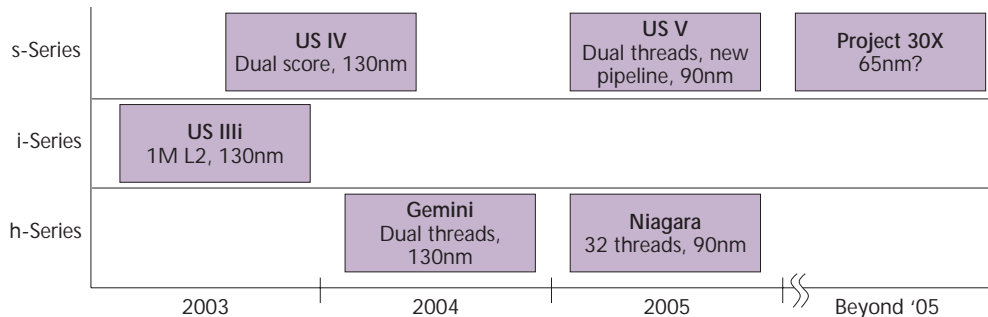


Figure 1. Sun's aggressive processor roadmap shows an ambitious plan to introduce five new processors during the next three years. This processor blitz starts this year with the UltraSPARC IIIi and IV. Beyond US V is an unnamed processor Sun claims will perform 30 times faster than today's enterprise processors; we'll call it Project 30X. Sun divides its systems into three main categories: s-Series for big iron enterprise computers; i-Series for one- to four-way systems; and the h-Series, which is designed for dense servers and blade servers.

Sun's first multicore SPARC will be the UltraSPARC IV, which integrates two US III cores on one die. It is pin compatible with the US III socket to simplify system upgrades, but it still offers memory subsystem improvements and improved reliability, accessibility, and serviceability (RAS) features.

The next new processor core will be the UltraSPARC V in 2005. The US V will support two threads, visual instruction set (VIS) 3.0 multimedia extensions, additional RAS features, and what Sun terms "aggressive performance enhancements." We believe US V will have a deeper pipeline (than US III/IV) and may finally add out-of-order execution to UltraSPARC (which has been available in Fujitsu's SPARC64 V but never in a Sun SPARC processor).

During 2003, Sun will bring the advanced features and RAS of the UltraSPARC III into the i-Series. The i-Series processor integrates a 1MB L2 cache on-die and is optimized for one- to four-way systems. Integrating the cache lowers system cost at the expense of cache expandability.

With the h-Series of products, Sun is also taking its 64-bit UltraSPARC solutions into blade servers and other dense and power-sensitive applications. The new form factor is a challenge best handled with high integration and advanced semiconductor process. In 2004, Sun will introduce a dual-threaded processor code-named Gemini, made in TI's 130nm semiconductor process. Sun is promising twice the performance of the 650MHz UltraSPARC Ili used in the Sun Fire B100s Blade Server.

Just one year later, in 2005, Sun plans to introduce a processor it promises will offer 15 times the throughput of the 650MHz UltraSPARC Ili—an audacious goal. To hit that target, Sun is using a design from recent acquisition Afara. The processor, code-named Niagara, integrates eight UltraSPARC II cores per die, and each core can support four threads. In addition, Niagara integrates four Gigabit Ethernet links, the memory controller, and a security coprocessor. Niagara is designed for blade servers and high-density servers, but it does not include support for larger multiprocessor configurations. Larger processor arrays would probably be configured using clustering software and the Ethernet links. As a startup, Afara




Figure 2. In this simplified execution model, memory access latency can dominate processing time, because each time a cache miss occurs, hundreds of clock cycles can be lost while a single-threaded processor's waiting for memory. Increasing clock frequencies has only a minor effect by shortening the processing time, but it does not help shorten memory access times, as the top bar shows.

developed this processor to accelerate Web content delivery; we would expect it to excel in TCP/IP processing as well.

Many details of Niagara are still missing, but with eight cores on die and up to 32 active threads possible, the processor will probably require either large caches or huge amounts of memory bandwidth. Niagara will likely use DDR2 main memory that will be in high volume in 2005.

Multithreading is an important design choice when dealing with the widening gap between processor core speeds and memory interfaces. As Figure 2 shows, the time spent waiting for memory access during a cache miss can swamp the performance improvements realized by higher clock speeds. Adding the capability of the processor to perform real work during memory access times can significantly benefit performance. Even single-threaded application performance can be improved using the threaded capabilities to prefetch data and speculatively execute instructions. The choice of four threads for each Niagara core, as shown in Figure 3, was probably a trade-off of expected memory latencies and design complexity.

Sun Relies on Past Stars

Les Kohn is leading the Afara team developing the radical multithreaded, multicore Niagara processor. In the 1980s, Kohn led the development of Intel's i860, one of the earliest dual-issue long-instruction-word (LIW) microprocessors (see *MPR 8/25/97-11*, "Arc of an Architecture: The i860") and worked on C-Cube's DVx MPEG-2 encoder chip (see *MPR 12/8/97-01*, "DVx Sets New Standard for Digital Video"). His experience at Sun includes lead roles in the UltraSPARC I and II. Having worked on seminal Sun processors yet bringing a radical approach to data processing makes Kohn an insider who thinks out of the box.

Another leader of this architectural movement within Sun with experience in multicore and ground-breaking processors is Dr. Marc Trembley, who was chief architect of the dual-core MAJC processor (see *MPR 10/25/99-04*, "Sun Makes MAJC With Mirrors") and the VIS extensions. In addition to multithreading and chip multiprocessing, Trembley is responsible for the next generation of SPARC processors (past US V; what we call Project 30X in Figure 1) and for

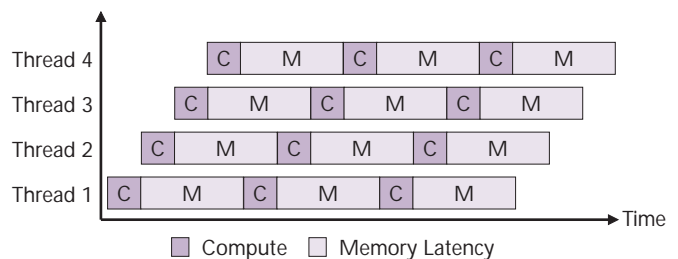


Figure 3. In this simplified model, using four threads allows the processor to overlap the execution times of each thread with memory access times.

implementing a new collection of buzzwords: hardware scout, space-time computing, transactional memory, and execute ahead. The goal of Project 30X is to produce a SPARC processor that has 30 times the performance of today's mainstream 1.2GHz UltraSPARC III processor in the s-Series. This is an order of magnitude faster than the goals of conventional processor-architecture design.

From what we can gather, *hardware scout* uses multithreaded capability to scan ahead in the instruction stream to look for opportunities to prefetch data and speculatively traverse an instruction path. In a sense, it's a more-intelligent and farther-reaching form of prefetch, in that it prefetches the instruction stream and applicable data without actually executing the program. The *execute ahead* concept could be a form of speculative execution using multithreaded capabilities of future UltraSPARC processors. Although *space-time computing* sounds like science fiction, it may have more to do with thread locality and execution timing in multiprocessing

systems than with Albert Einstein and wormholes, whereas the rather cryptic *transactional memory* sounds very businesslike. Details of these innovations are still undisclosed by Sun, as it is awaiting patents.

Sun's new roadmap represents a huge challenge for the company, but it has thrown caution to the winds and is making significant processor promises for the next three years. Our initial reaction is mixed enthusiasm and skepticism. Enthusiasm, because we're excited to see Sun taking aggressive steps to reinvent and redefine server processor performance and challenge the industry with this daring move. Yet we harbor some skepticism, because Sun's execution on its UltraSPARC III was, at best, lackluster and uninspiring. With a renewed sense of excitement and challenge, however, and new leadership (including Dr. David Yen as executive vice president of the processor group), Sun is once again taking up the role of innovator and is showing it will not allow Intel's Itanium juggernaut to roll over it. ♦



SUBSCRIPTION INFORMATION

To subscribe to *Microprocessor Report*, contact our customer service department in Scottsdale, Arizona by phone, 480.609.4551; fax, 480.609.4523; email, emckeighan@instat.com; or Web, www.MDRonline.com.

One year	U.S. & Canada*		Elsewhere		Two years	U.S. & Canada*		Elsewhere	
Web access only		\$795		\$795	Web access only		\$1,395		\$1,395
Hardcopy only		\$895		\$995	Hardcopy only		\$1,595		\$1,795
Both Hardcopy and Web access		\$995		\$1,095	Both Hardcopy and Web access		\$1,795		\$1,995

*Sales tax applies in the following states: AL, AZ, CO, DC, GA, HI, ID, IN, IA, KS, KY, LA, MD, MO, NV, NM, RI, SC, SD, TN, UT, VT, WA, and WV. GST or HST tax applies in Canada.