Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

UltraSPARC IV Mirrors Predecessor

Sun Builds Dual-Core Chip in 130nm

By Kevin Krewell {11/10/03-02}

As part of its Throughput Computing initiative, Sun has started down the path to chip-level multithreading (CMT) with its UltraSPARC IV (US IV), revealed at **Microprocessor Forum 2003**. Sun's first US IV consists of two slightly enhanced UltraSPARC III cores that share a systems bus, DRAM memory controller, and off-die L2 cache. Built in the Texas Instruments 130nm process, the dual cores share one large off-die L2 cache, but when Sun migrates the US IV processor to 90nm, the company will add an on-die L2 cache, with the off-die cache becoming a large L3 cache.

Sun's goal was to increase performance with a multicore approach, initially without significant increases in clock speeds. Sun's chip-level multithreading could also be called chip multiprocessing. The UltraSPARC IV processor should ship in Sun systems in 1H04 with clock speeds of 1.05GHz and 1.2GHz. Figure 1 shows the roadmap for the UltraSPARC family through 2006.

The even-numbered member of the UltraSPARC family continues a plan to leverage the infrastructure and system bus created in UltraSPARC III. The goal is to provide a stable platform over an extended period of time. The benefit of this continuity is that US IV can be an on-site upgrade to a number of SunFire systems (V480–V1280, 4800–15K). Solaris can support a mixture of US III and US IV processors.

In simplistic terms, Sun created the UltraSPARC IV by dropping two UltraSPARC III cores onto one die, with one core mirroring the other, as shown in Figure 2. An undesirable effect of mirroring the cores was that it put a hot spot of each core directly next to that of the other (the floating-point unit). Some design changes, such as heat "towers" in the copper interconnect, were required to draw the heat from this hot spot on the die.

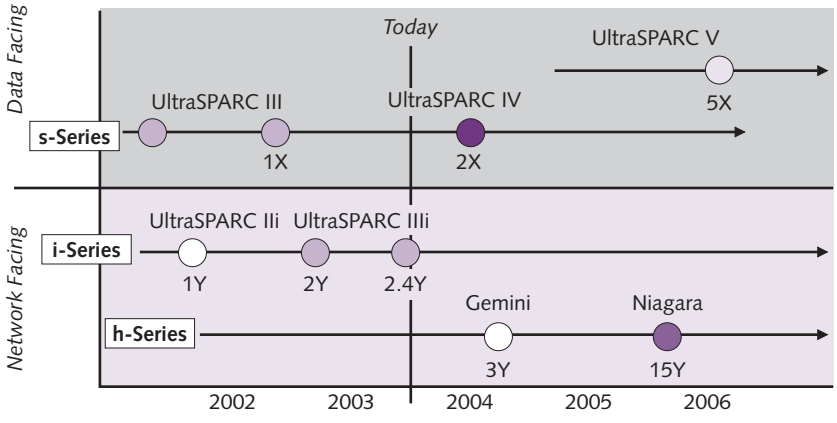

Figure 1. The roadmap for the UltraSPARC family. The UltraSPARC IV will start in 130nm process and then migrate to a 90nm process (with an on-die L2 cache) before the introduction of the UltraSPARC V. UltraSPARC V is planned as a higher-performing single core with dual-threading support. The i-Series will use the US III core in a part called Jalapeno. The i-Series uses a different bus for 1- to 4-way servers and provides higher frequencies in a lower-cost package. The h-Series will be a blade-server design with a simpler core design and multithreading support. The first project is Gemini—a low-power two-thread processor targeted at three times the performance of the IIi. The second h-Series product, Niagara, will have more cores and four threads per core. The h-Series targets web, application serving, and simple databases.

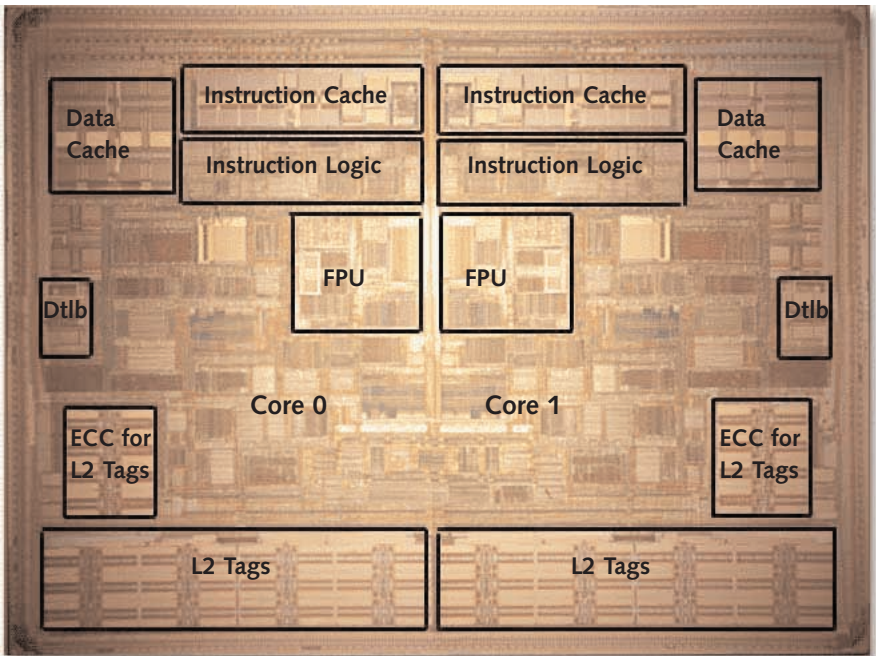

Figure 2. The die photo of the UltraSPARC IV clearly shows the mirror image of the two cores. By mirroring the layout of one core, the shared system bus could be placed in the center, close to each core. Unfortunately, it also had a side effect of placing the floating-point units (FPU) of the two cores right next to each other, creating a challenging hot spot on the die.

Sun also took this opportunity to fine-tune the US III architecture slightly. Previous improvements to the US III design included copper interconnects in the 180nm process, data prefetch logic, improvements to the TLB, and the two-way set-associative L2 cache. (For details of the UltraSPARC III microarchitecture, see *MPR 10/27/97-07*, "UltraSparc-3 Aims at MP Servers" and an update in *MPR 10/16/00-04*, "Sun Shines on UltraSPARC III.") The US IV core improvements focused on instruction fetch, store bandwidth, and improved data prefetching. The floating-point adder was enhanced to handle more NaN (not a number) and underflow cases in hardware, reducing exceptions on certain challenging problems. The on-die 2KB write cache was enhanced with a hashed index to better handle multiple write streams.

The two cores share one external L2 cache, but each core has its own on-chip cache tags. The external L2 SRAM can be up to 16MB, logically divided into two 8MB caches per core, as shown in Figure 3. The cache has 128-byte lines subdivided into two 64-byte subblocks that are two-way set-associative and have a least-recently-used (LRU) replacement policy. The cache bus uses fair dynamic arbitration for bandwidth. The bus runs at some ratio slower than the core clock speeds and has a total bandwidth of 8GB/sec.

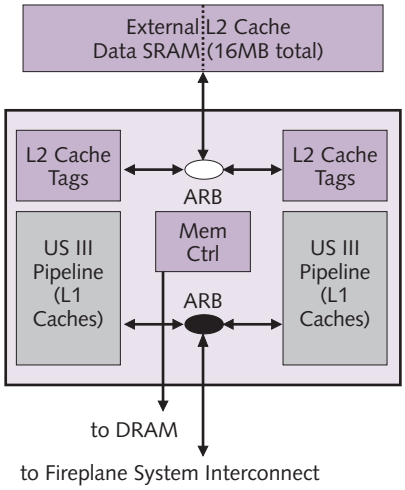

Figure 3. The block diagram of the UltraSPARC IV. The external L2 cache bus and the system and SDRAM buses are arbitrated. This could present a bottleneck to data-movement-intensive applications, but the write cache helps reduce L2 bandwidth requirements.

The US IV is compatible with the US III system bus, the Sun Fireplane. The on-chip memory controller supports up to 16GB of DRAM per core and supports fault-tolerant chip kill DRAM. The two cores arbitrate for the shared system interface and act as a single Fireplane client but support two interrupt IDs (one per core). The Fireplane bus has a hierarchical address bus and is a point-to-point data bus to a data switch that supports two UltraSPARC processors (see Figure 4).

The US IV has a rather modest 66 million transistors in the Texas Instruments 130nm process. The die size came in at 22.1mm × 16.1mm (356mm^2), and the power consumption is said to be 100W at 1.35V. The number of transistors will rise sharply when Sun migrates to 90nm and adds an on-die L2 cache, but Sun can afford to add the on-die cache once the cores are shrunk in the smaller-geometry process. The MDR model indicates that in 130nm, the US IV will cost $310 to manufacture. The 90nm version will probably be roughly similar in die size (300+mm^2), with 2–4MB of on-die L2 cache.

As Sun rolls out these new processors and the CMT initiative, reports are now circulating that Sun and Fujitsu are in talks to co-develop SPARC processors. Fujitsu's latest incarnation of the SPARC 64 was also revealed at Microprocessor Forum 2003 and will be the subject of a future story. If the rumors are true, the Fujitsu SPARC 64 VI core design could find its way into Sun's future s-Series high-end servers. Present reports indicate that the i-Series and h-Series processors would still be developed by Sun.

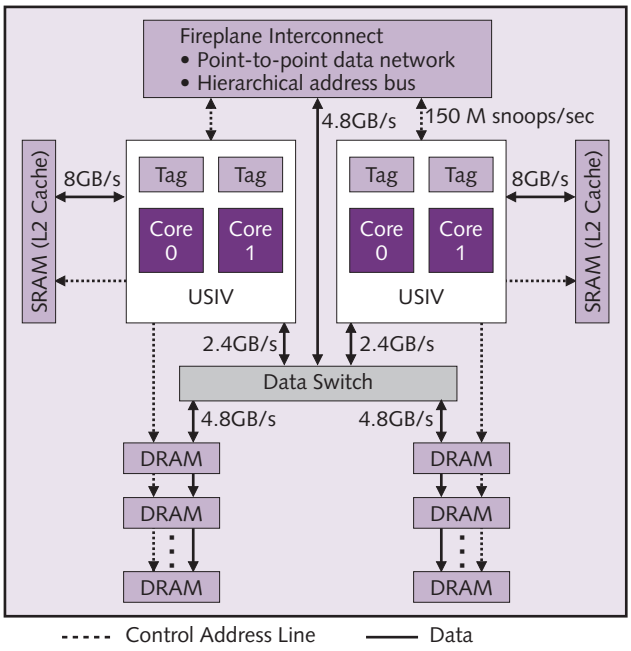

Figure 4. The block diagram of the Fireplane system configuration with UltraSPARC IV processors.

Price & Availability

UltraSPARC IV will ship in 2004 in SunFire systems. Pricing was not available at the announcement.

The UltraSPARC IV represents a significant first step in Sun's long-term plan for Throughput Computing and chip multithreading. Sun's modern support for threading dates to 1992 with its first multithreaded kernel release (Solaris 2.0). Solaris 8 provided the first release of an optional (1:1) threads model in 2000. The company followed with Solaris 9, making the thread model standard. One driving factor was that Java 1.0 had threading support built into the language. The latest release of Red Hat Enterprise Linux 3.0 (one of Sun's Linux partners) also improves threading support by incorporating the Native POSIX Thread Library (NPTL) into a version 2.4 kernel. Continuous improvements are being made in the ability of server software to support extensive threading, and that development is in alignment with Sun's Throughput Computing initiative and the company's plans for the UltraSPARC architecture.