

UltraSPARC[®] IV Processor Architecture Overview

Technical Whitepaper

**February 2004
Version 1.0**



<http://www.sun.com>

Copyright © 2004 Sun Microsystems, Inc. All Rights reserved.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY EXPRESS REPRESENTATIONS OR WARRANTIES. IN ADDITION, SUN MICROSYSTEMS, INC. DISCLAIMS ALL IMPLIED REPRESENTATIONS AND WARRANTIES, INCLUDING ANY WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

This document contains proprietary information of Sun Microsystems, Inc. or under license from third parties. No part of this document may be reproduced in any form or by any means or transferred to any third party without the prior written consent of Sun Microsystems, Inc.

Sun, Sun Microsystems, the Sun Logo, Java, Jini, SunFire, Netra, Solaris, Jiro, Sun Enterprise, Ultra, Write Once, Run Anywhere, SunNet Manager, and The Network is the Computer are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. UNIX is a registered trademark in the United States and other countries, exclusively licensed through x/Open Company Ltd.

The information contained in this document is not designed or intended for use in on-line control of aircraft, air traffic, aircraft navigation or aircraft communications; or in the design, construction, operation or maintenance of any nuclear facility. Sun disclaims any express or implied warranty of fitness for such uses.

The UltraSPARC IV Processor Architecture

1.1 UltraSPARC IV Processor Brief

The UltraSPARC[®] IV processor is among the first Chip Multithreading (CMT) processors that follows Sun's Throughput Computing strategy and continues the tradition of binary code compatibility by complying with the 64-bit SPARC[®] International Version 9 Instruction Set Architecture (ISA).

The UltraSPARC IV processor is a dual-thread processor which supports up to 16 MB of external level-2 (L2) cache. With the exception of one pin¹, the UltraSPARC IV processor features the same footprint as a single UltraSPARC III processor. This design point minimized motherboard modification and time to market.

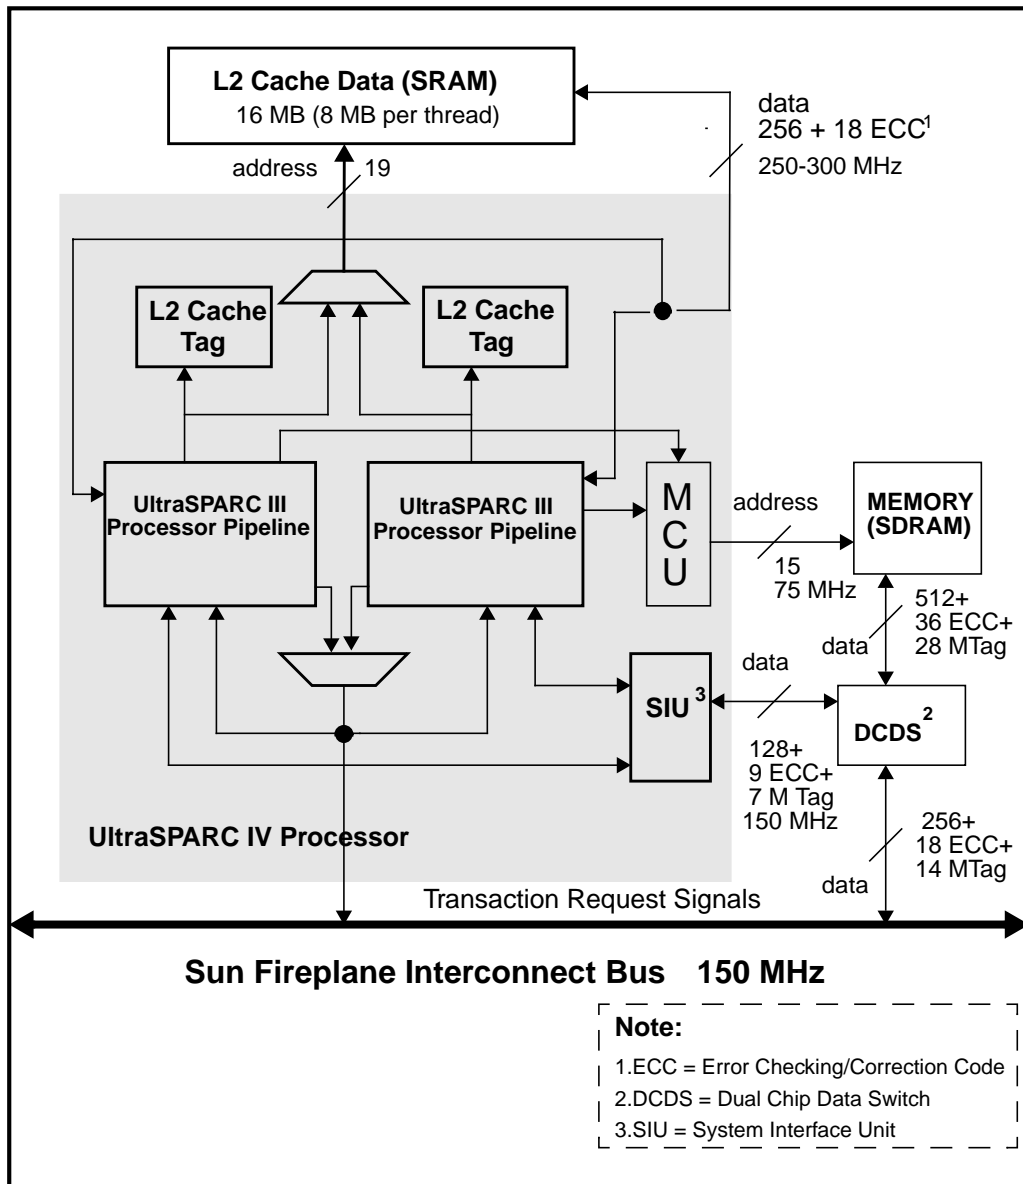
The primary design goal for the UltraSPARC IV processor is to improve throughput performance in commercial applications such as databases, web servers, and High Performance Technical Computing (HPTC). The following key techniques are used to improve the UltraSPARC IV processor's performance:

- *Implementing dual-thread processing capabilities*
This CMT technology nearly doubles current compute densities and reduces overall heat dissipation, resulting in significant end-user cost-benefit savings.
- *An improved Level 2 (L2) cache configuration*
Each thread in the UltraSPARC IV processor can access 8 MB of 2-way, set-associative L2 cache. The L2 cache line sizes changed from 512 bytes to 128 bytes to reduce data contention associated with sub-blocked caches. This change balances cache efficiencies over a wide range of data sets, enhancing performance throughput over an even broader range of general applications. In addition, a Least-Recently-Used (LRU) L2 cache eviction policy is used for more optimal use of caching resources.
- *An enhanced Floating Point Unit (FPU)*
The enhanced Floating Point Unit produces higher performance on codes such as large radix Fast Fourier Transform (FFT). The FPU has additional hardware assist for IEEE 754-1985 exception processing.
- *An enhanced write cache*
The write cache, used to relieve the cache bus from inefficient use of write bandwidth, has been enhanced with a hashed-index algorithm to reduce write cache latency.

1. The single pin difference is due to an extra address line required to switch from one of either 8 MB L2 cache address blocks.

Threads running on the UltraSPARC IV processor share the following: address and data bus to access the L2 cache data, Memory Control Unit (MCU), and the Sun Fireplane™ Interconnect port. The bus to the L2 cache and the physical SRAM modules containing the L2 cache is shared. Although the L2 cache is logically separate for each thread running on the UltraSPARC IV processor, the L2 cache is physically contained in one SRAM module. Figure 1-1 illustrates the UltraSPARC IV processor block diagram.

Figure 1-1 Basic UltraSPARC IV Processor



1.2 Summaries of New Features in the UltraSPARC IV Processor

Table 1-1 and Table 1-2 summarize the UltraSPARC IV processor enhancements with respect to the UltraSPARC III processor. Table 1-1 lists enhancements in single-thread execution including clock rate increments and new cache organization; Table 1-2 lists changes due to the CMT technology implemented in the UltraSPARC IV processor.

1.3 RAS Architecture Improvements

The UltraSPARC IV processor inherits all of the reliability, availability, and serviceability features (see the UltraSPARC III Processor summary end note) implemented in the latest UltraSPARC III processor with the addition of L2 cache address bus error protection.

1.3.1 Cache Bus Error Protection

The data bus between the UltraSPARC III processor and the external L2 cache is ECC-protected by splitting the bus into two 16 byte sections. Each 16 byte section has its own address bus and control signal bus. Since the UltraSPARC IV processor inherits architectural characteristics of the UltraSPARC III processor, these upper and lower 16 byte buses are maintained. On the UltraSPARC IV processor, the lower 16 byte address bus can access the ECC of the upper 16 bytes of data. Likewise, the upper 16 byte address bus can access the ECC of the lower 16 bytes of data. By splitting the ECC methodology in this manner, the entire address bus used to access the L2 cache are implicitly protected.

1.4 Advanced Compiler Analysis Techniques

The UltraSPARC IV processor's architectural improvements provide a performance enhancement over its predecessor. By using Sun's latest compilers, even higher performance levels can be achieved by the UltraSPARC IV processor.

The latest Sun compilers use advanced analytical techniques demanded by classic multiprocessor systems and CMT processors. These new techniques include accurate dependence analyses to assist automatic code execution parallelization -- thereby optimizing the utilization of available processor resources and increasing throughput performance. Sun's new compiler analyses can be applied to a variety of code scenarios and include techniques such as index-association dependency analysis and transformations of complex loop nests. With the new analyses, performance on many industry-standard benchmarks is also significantly higher.

As an example, consider memory disambiguation analysis which is a common need in loop optimization. The primary problem here is to determine whether more than one array reference can access the same memory location. Traditional compiler analysis techniques only handle array subscripts that are linear functions of enclosing loop indices. Such tests are often successful on simple loops, but are unable to extract meaningful alias information from more complex loops. Sun's index-

association-based dependence analysis can handle more complex subscripts, allowing loop nests to be parallelized that would otherwise be serially executed. Additionally, by using cache locality knowledge during parallelization, throughput performance is further enhanced on CMT designs. Using such techniques, the UltraSPARC IV processor has shown preliminary performance increases of 1.60x and 1.14x over its predecessor for the SPEC CPU2000 benchmark tests **swim** and **lucas**. Similar performance enhancements are expected for larger multiprocessor systems.

Sun's compilers also contain extensive inter-procedural and inter-module analysis capabilities. With these techniques, variables and pointers can be tracked well beyond small local regions, cross function and file boundaries. Alias information can then be gathered with much better knowledge of the whole program. Versioning and cloning are employed to take full advantage of the information gathered. If the profile feedback information gathered from training runs.

Every release of the compiler is carefully tuned to maximally leverage new releases of hardware. As an example, consider the increasing gap between processor and memory speeds. Cache misses and memory latency are rapidly becoming a critical bottleneck for many applications. Sun's compilers employ techniques to detect potential cache misses and utilize prefetch to fetch the data before it is needed. Moreover extensive efforts have been made to match the prefetching to the processor and system characteristics.¹

Sun also provides a suite of highly tuned libraries. In some common cases, the compilers can automatically detect particular code patterns (idioms) and call a corresponding tuned library function. Users can also directly use the provided library functions saving development and tuning costs.

Sun's compilers employ various algorithms for optimal performance. Various traditional and Sun-developed techniques have helped extract the most out of processor resources and increase throughput performance.

1. *Processor Aware Anticipatory Prefetching in Loops*, Partha Tirumalai, Yonghong Song, Spiros Kalogeropoulos, Vikram Rao, and Raja Mahadevan, roceedings of the 10th International Symposium on High Performance Computer Architecture, Madrid, Spain, Feb 14-18, 2004.

TABLE 1-1 UltraSPARC IV Processor Enhancements

Feature	Benefit
Each thread running on the UltraSPARC IV processor can access 8 MB of L2 cache with 128-byte line size (2 sub-blocks per line) or 4 MB with 64-byte line size (no sub-block)	Higher thread performance by making caching more efficient over a larger range of applications
L2 cache employs least-recently-used (LRU) eviction strategy	LRU is a cache eviction policy resulting in better cache hit rates leading to faster execution and better system throughput
L2 cache control a wide range of interface clocking	Allows for a larger range of L2 cache SRAM clock speeds as semiconductor technology progresses and processor clock speeds increase
Supports higher system clock divisors	Increases overall system throughput performance by allowing for high clocks multiples of 150 MHz
L2 cache Address Bus error protection	Higher reliability
Hash-indexing for write cache	Decreases conflict misses during multiple write streams resulting in higher write store bandwidth and overall performance
Additional hardware assist for IEEE 754-1985 floating point exception processing	Decreases system overhead by having processor hardware logic perform exception processing rather than relying on the operating system software
Software prefetch semantics used with hardware prefetch cache	Higher floating point performance

TABLE 1-2 Enhancements Due to CMT Technology

CMT Enhancements
Resources such as MCU registers, pins, and Sun Fireplane Interconnect registers are shared so that either thread in the processor is able to access these registers. For example, allowing either thread to modify memory controller timing values.
New registers have been added to support the Sun Standard CMT model, allowing compatibility with current and future operating system standard interaction with all Sun CMT processors.
Processor registers initialized with values associated with CMT. For example, establish which of the two threads in the processor will address a cache error.
New per-thread ASI_CESR_ID register is added. This is a thread ID associated with the Sun Fireplane Interconnect block I/O. A remote device can identify the source thread of a block data move.

1.5 Conclusion

The dual-thread UltraSPARC IV processor is Sun Microsystems' first-generation CMT processor targeting mid-range to high-end servers. Furthermore, end users' software investment is protected by maintaining binary compatibility. With a modular upgrade path for current UltraSPARC III processor deployments, investments in current data centers are both secured and enhanced.

Follow-on processors in the UltraSPARC IV processor family will further enhance the performance already attained by the first generation of UltraSPARC IV processors. Single-thread performance, exploiting Texas Instruments' 90 nm semiconductor process technology, increases in clock frequency, bandwidths, and the addition of a large level-3 cache are just few of the new features for future UltraSPARC IV processors.

UltraSPARC III Processor Summary

- 64-bit SPARC Version 9 Instruction Set Architecture
As with the UltraSPARC I and UltraSPARC II processors, UltraSPARC III processors are compatible with the SPARC Version 9 Instruction Set Architecture. These 64-bit processors are binary compatible with earlier SPARC V7 and SPARC V8 ISAs and VIS 2.0 compliant.
- Non-aligned instruction fetch
No penalty is introduced to perform a quad-word boundary aligned instruction fetch.
- Single clock throughput from on-chip caches
- Over 1000-way scalability
- 14 Stage, non-stalling pipeline
When the wait or hold condition is removed, the instructions that would have stalled the pipeline are simply re-executed.
- Six peak instruction issue rate
Four instructions are dispatched from the 16-entry instruction buffer, six instructions can be issued into six parallel execution units including 2 integer, 1 branch, 1 load/store, 2 floating point units (consisting of 1 floating point multiply/divide, and 1 floating point add/subtract).
- 95% branch prediction accuracy
- 64 KB L1 data cache, single cycle (wave pipelined) throughput
- 32 KB L1 instruction cache, single cycle (wave pipelined) throughput
- 8 MB L2 external cache, single cycle (wave pipelined) throughput
- On-chip memory controller (addressing 16 GB)
- On-chip L2 cache tags
- 2 KB write cache, single cycle (wave pipelined) throughput
The write cache enhances general-purpose application performance by having weighing on read-intensive data cache bandwidth with little cost on write bandwidth.
- 2 KB prefetch cache
This cache is accessed in parallel with the data cache for floating point loads. Floating point load misses, hardware prefetches and software prefetches bring data into this cache.
- Speculative execution of instructions after branches
- Speculative memory loads
- ECC or parity on all major SRAMs both internal and external
- Diagnostic bus
- 130 nm Texas Instruments semiconductor process technology

References

- 1) *Multithreaded Technologies Disclosed at MPF*, Microprocessor Forum/Microprocessor Report, Marcus Levy, November 10, 2003, www.MPRonline.com
- 2) *Sun Microsystems' UltraSPARC IV Processors*, Quinn Jacobson, Presented at Microprocessor Forum 2003, www.MPRonline.com
- 3) *Transformation of Loops Containing Induction Variables*, K.S. Serebrainy, <http://www.informika.ru/text/magaz/it/2003/09/contents.html>
- 4) *Index-Association Based Dependence Analysis and its Application in Automatic Parallelization*, Sun Microsystems, Yonghong Song and Xiangyum Kong, <http://parasol.tamu.edu/lpc03/informal-proceedings/Abstracts/2.pdf> or <http://parasol.tamu.edu/lpc03/informal-proceedings/Papers/2.pdf>
- 5) *Throughput Computing*, Greg Papadopoulos, EVP, Chief Technology Officer, Sun Microsystems, Presented at Microprocessor Forum 2003, www.MPRonline.com
- 6) *Processor Aware Anticipatory Prefetching*, Partha Tirumalai, Yonghong Song, Spiros Kalogeropoulos, Vikram Rao, and Raja Mahadevan, Proceedings of the 10th International Symposium on High Performance Computer Architecture, Madrid, Spain, February 14-18, 2004.