

UltraSPARC T2: A Highly-Threaded, Power-Efficient, SPARC SOC

Manish Shah, Jama Barreh, Jeff Brooks, Robert Golla, Gregory Grohoski, Nils Gura, Rick Hetherington, Paul Jordan, Mark Luttrell, Christopher Olson, Bikram Saha, Denis Sheahan, Lawrence Spracklen, Aaron Wynn

Sun Microsystems

5300 Riata Park Court, Austin, Texas 78727

Abstract—UltraSPARC T2 is Sun Microsystems' second generation multi-core, multi-threaded SPARC System-on-a-chip. It delivers twice the throughput performance of the first generation UltraSPARC T1 processor in essentially the same power envelope. UltraSPARC T2 supports concurrent execution of 64 threads by utilizing eight SPARC cores, each with eight hardware threads. The cores communicate via a high bandwidth crossbar and share a 4MB, eight bank, L2 cache. Each SPARC core includes two integer execution units and a dedicated floating point and graphics unit, which delivers a peak floating point throughput of 11.2 GFLOPS/sec at 1.4 GHz. Each core also has a cryptographic unit. For I/O, UltraSPARC T2 has an integrated x8 PCI-Express channel and two 10G Ethernet ports with XAUI interfaces. Memory is accessed via four on-chip controllers each controlling 2 FBDIMM channels for a peak memory bandwidth in excess of 60 GB/sec. UltraSPARC T2 is fabricated in an 11 metal, 1.1V, triple-Vt CMOS process. The chip has ~500M transistors on a 342 mm² die with a power consumption of 84 W at 1.4 GHz. The high level of system integration along with high throughput, floating point, and cryptographic performance makes UltraSPARC T2 an ideal choice for a range of applications including web servers, database and applications servers, High Performance Computing, secure networking, campus backbones, and file servers.

I. INTRODUCTION

UltraSPARC T2 is Sun Microsystems' second generation chip threaded processor. Historically, microprocessors have been designed to improve the execution performance of single thread programs by exploiting instruction level parallelism (ILP). Common techniques used to improve single thread performance are deep pipelines, multiple instruction issue, speculation, and out-of-order instruction execution. Recently, these techniques have reached a point of diminishing returns because of inherently low or hard to exploit application ILP [1]. Techniques used to improve single thread performance often give rise to complex processor designs with poor pipeline efficiencies and high power consumption.

UltraSPARC T2 has been designed for commercial workloads which exhibit large amounts of thread level parallelism (TLP). Unlike desktop workloads, which often have a small number of threads running concurrently, most commercial workloads achieve high scalability by employing large numbers of concurrent threads. A thread can be a process which is part of a parallel program or a program by itself. Due to large working sets, commercial applications generally exhibit low ILP, and high cache miss rates. As the

speeds of traditional processors have increased much faster than the memory speed, traditional processors spend a significant amount of time waiting for data from memory. UltraSPARC T2 pipelines load misses to memory on a thread basis in order to minimize the effect of memory latency on overall throughput performance. UltraSPARC T2 employs chip multi-threading (CMT) technology which utilizes up to 64 hardware threads to achieve high throughput on commercial workloads by taking advantage of the TLP inherent to such workloads.

This paper describes the design of the UltraSPARC T2 processor. Section II gives an overview of UltraSPARC T2's architecture and its design goals. Section III describes the SPARC processor core. Section IV details the cryptographic functions supported by UltraSPARC T2. The reliability and serviceability features of UltraSPARC T2 are described in section V. Section VI describes the power management features of UltraSPARC T2. Section VII describes the performance and applications of UltraSPARC T2.

II. UltraSPARC T2 ARCHITECTURE

The block diagram of UltraSPARC T2 is shown in Fig. 1. UltraSPARC T2 is a single chip, multi-threaded processor consisting of eight 64-bit SPARC processor cores. Each SPARC processor core has full hardware support for execution of eight independent threads. Each SPARC processor core consists of two integer execution units, a floating point and graphics unit, and a cryptographic stream processing unit. The eight SPARC processor cores share an 8-banked, 4MB Level 2 cache (L2). Each bank of the L2 cache is 16-way set associative with a line size of 64 bytes. The eight banks of L2 allow eight simultaneous accesses to support the high bandwidth requirements of UltraSPARC T2. The SPARC processor cores communicate to the L2 cache through a high bandwidth, non-blocking, pipelined crossbar switch. This crossbar connects the eight SPARC processor cores to the eight banks of the L2 cache and the I/O port, and has a total write bandwidth of 90 GB/sec, and a total read bandwidth of 180 GB/sec. The L2 cache connects to four on-chip DRAM controllers which interface directly to a pair of fully buffered (FBDIMM) channels with a peak link rate of 4.8 Gb/sec. FBDIMM was selected since it provides ~2x the bandwidth at less than half the pin count of DDR2. The peak memory bandwidth of UltraSPARC T2 is 50 GBytes/sec for reads and 26 GBytes/sec for writes.

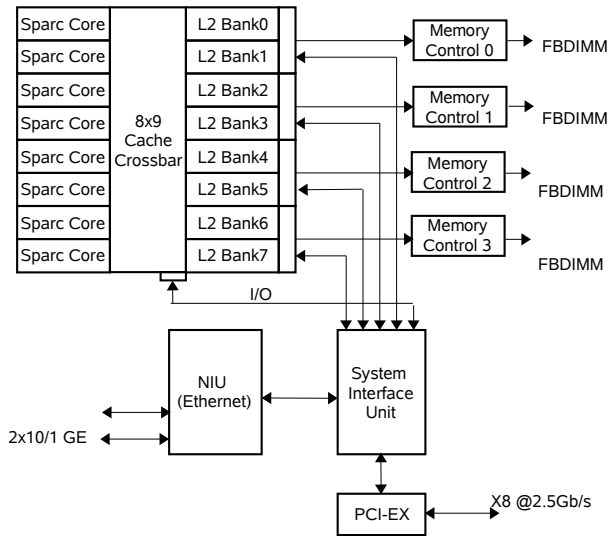


Fig. 1. UltraSPARC T2 Block Diagram

UltraSPARC T2 is a system-on-a-chip featuring a PCI-Express unit which implements the root complex behavior of the PCI-Express Base Spec 1.0A. The PCI-Express unit supports x1, x2, x4, and x8 configurations at a data rate of 2.5 Gb/sec. The network interface unit on UltraSPARC T2 implements two on-chip 10 Gb/sec Ethernet ports with XAUI interfaces. All the SerDes are on-chip, providing an aggregate bandwidth greater than 1 Tb/sec. The high level of system integration on UltraSPARC T2 reduces overall system complexity, component count, and power consumption.

UltraSPARC T2 shown in Fig. 2 is fabricated in Texas Instruments' 11 Metal, 1.1V, 65nm triple-Vt CMOS process, and has ~500M transistors. UltraSPARC T2 has a die size of 342mm². It is packaged in a flip-chip, glass ceramic package with 1831 pins.

The primary goal of UltraSPARC T2 is to double the throughput and throughput/watt performance of its predecessor, the UltraSPARC T1. One method to double UltraSPARC T1's throughput would have been to simply double the UltraSPARC T1 cores from 8 to 16. While simple,




Fig. 2. UltraSPARC T2 Chip Micrograph

this approach is area intensive and power-inefficient. Instead, UltraSPARC T2 doubles UltraSPARC T1's throughput performance by doubling the number of threads per SPARC processor core from 4 to 8, and doubling the number of execution units per core from 1 to 2. Threads are statically assigned to each execution unit, 4 threads per unit. The level 1 caches (L1) and the translation lookaside buffers (TLB) are shared by all eight threads. UltraSPARC T2's L2 cache size is 4MB, 8-way banked with a set associativity of 16 in order to meet the bandwidth demands of 8 cores each having 8 threads per core.

UltraSPARC T2 improves upon the UltraSPARC T1's floating point throughput by at least a factor of 8. This gain is achieved by incorporating a dedicated floating point and graphics unit on each SPARC processor core. UltraSPARC T2 extends cryptographic support and improves integer performance over the UltraSPARC T1.

III. SPARC PROCESSOR CORE

Fig. 3 shows the block diagram of a UltraSPARC T2 SPARC processor core. Each SPARC processor core on UltraSPARC T2 has full hardware support for execution of eight independent threads. The SPARC processor core consists of an instruction fetch unit (IFU), two integer execution units (EXU0 and EXU1), one floating point and graphics unit (FGU), one load-store unit (LSU), one stream processing cryptographic unit (SPU), a trap logic unit (TLU), a memory management unit (MMU), and a gasket unit [2].

The eight threads are statically partitioned into two thread groups of four, TG0 and TG1. The four threads within a thread group share one EXU. The execution engines in the FGU and SPU are shared by all eight threads. The IFU holds a 16KB, 8 way set associative level 1 instruction cache with a 32 byte line size, and a 64 entry instruction TLB. The instruction cache and the instruction TLB are shared by all eight threads. The LSU holds a 4 way set associative level 1 data cache with a 16 byte line size, and a 128 entry data TLB. The data cache and the data TLB are shared by all eight threads. The MMU supports 8KB, 64KB, 4MB, and 256MB pages, and implements a hardware tablewalk engine for fast reloading of the TLBs. The hardware tablewalk engine is

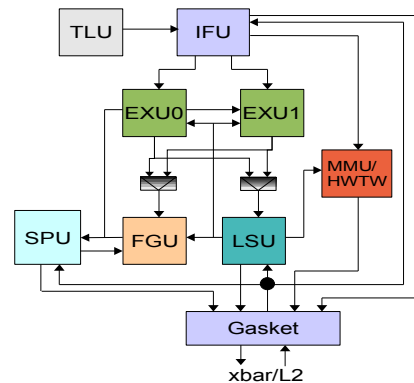


Fig. 3. UltraSPARC T2 SPARC Processor Core Block Diagram

shared by all eight threads. The TLU handles exceptions and interrupts, and is responsible for maintaining the machine state for each thread. The gasket unit sends memory and I/O requests from the processor core to the crossbar, receives the responses from L2 and I/O port, and sends the responses to the appropriate SPARC cores.

UltraSPARC T2 implements a fine-grained multi-threading scheme where the threads are switched on a cycle-by-cycle basis between the available threads within the statically partitioned thread groups TG0 and TG1, using a least recently issued priority scheme. When a thread encounters a long-latency event such as a cache miss, it is marked unavailable, and instructions will not be issued from that thread until the long-latency event is completed. Each cycle up to two instructions can issue, one from each thread group. These instructions may consist of two integer operations, one integer and one memory operation, one integer and one floating point operation, or one floating point and one memory operation. UltraSPARC T2's fine-grained multi-threading scheme serves to minimize throughput performance losses arising from branch mispredictions and cache misses.

UltraSPARC T2 implements an eight stage pipeline for integer operations, and has a three cycle load use penalty. The floating point pipeline has a six cycle latency for dependent floating point operations.

IV. CRYPTOGRAPHIC UNIT

The cryptographic stream processing unit (SPU) is designed to accelerate generic encryption and authentication operations frequently used in security protocols such as SSL, TLS, and IPsec. The cipher group (CG) subunit supports bulk encryption and authentication operations. The modular arithmetic (MA) subunit supports public-key cryptography. The CG and MA subunits operate independently.

The CG subunit is designed to support wire-speed cryptographic operation of bulk cipher and hash functions across the two 10 Gb Ethernet ports. This equates to 5 GB/sec of full-duplex data bandwidth. The CG subunit contains dedicated hardware engines for RC4, DES/3DES, and all variants of AES (128, 192, and 256-bit key lengths). Various modes of operation are supported for AES (ECB, CBC, CTR) and DES (ECB, CBC, and CFB). Hardware also supports MD5, SHA-1, and SHA-256 authentication functions, including keyed hashing (HMACs). Each of the 8 SPUs on UltraSPARC T2 provides 1/8th of the necessary processing to achieve an aggregate data rate of 5 GB/sec at 1.4 GHz.

The MA subunit accelerates key exchange, key generation, and authentication functions for security protocols such as IPsec and SSL/TLS. It provides direct hardware support for RSA, DSA, and Diffie-Hellman operations for up to 2048 bit keys via modular exponentiation. Additionally, the MA subunit supports key exchange and signature generation and verification using elliptic curve cryptography. More specifically, the MA subunit accelerates point multiplication for binary and integer fields for all standardized curves and key lengths exceeding 600 bits.

V. RELIABILITY AND SERVICEABILITY FEATURES

UltraSPARC T2 brings a new level of throughput performance to low-end and mid-range servers. Major hardware structures are protected with ECC or parity, as appropriate. Table 1 below summarizes reliability and serviceability (RAS) features of major UltraSPARC T2 chip components.

If an error occurs, the CMT nature of UltraSPARC T2 allows independent workloads to continue operation unaffected until the error condition is logged and, if necessary, corrected and cleared.

TABLE 1
RAS FEATURES OF SALIENT UltraSPARC T2 HARDWARE STRUCTURES

HW Structure	Protection	Recovery
Instruction Cache Tag	Parity	Hardware
Instruction Cache Data	Parity	Hardware
Windowed Integer Register File	SEC/DED ECC	Software
Floating Point Register File	SEC/DED ECC	Software
Data Cache Tag	Parity	Hardware
Data Cache Data	Parity	Hardware
Store Buffer Data	SEC/DED ECC	Software
Store Buffer Tag	Parity	Software
L2 Cache Data	SEC/DED ECC	Hardware
L2 Cache Tag	Parity	Hardware
L2 Cache Coherence	SEC ECC	Software

VI. POWER MANAGEMENT

Over the last two decades, the performance of traditional processors has improved significantly by increasing clock frequencies and the number of transistors per die. This performance improvement has come at the expense of increased processor power consumption and heat dissipation. This has led to higher costs for packaging, electricity, and cooling. UltraSPARC T2 seeks to minimize these costs by utilizing different power management techniques.

Traditional processors often utilize instruction speculation to increase execution performance, but incorrect speculation can result in increased power consumption. UltraSPARC T2 implements very limited execution speculation in order to minimize overall power consumption. To further reduce power, UltraSPARC T2 transitions a thread to a not-ready state whenever a thread encounters a long latency operation. Clock power is further reduced on UltraSPARC T2 by utilizing clock gating on datapath, control, and array structures. Power throttling is supported through 3 external power throttle pins. Based on the state of the power throttle pins, stall cycles are injected into the processor core pipeline to reduce overall dynamic power consumption.

In order to reduce static or leakage power, logic gates with sufficient timing, noise, and slew margins are replaced with footprint-compatible gates with longer channel length [3]. UltraSPARC T2 includes on-chip thermal diodes which enable the system to regulate the die temperature by controlling the instruction issue rate, and by disabling threads.

UltraSPARC T2 achieves its high throughput performance with a low power consumption of 84W at 1.4 GHz and at less than 2W per thread.

VII. APPLICATIONS AND PERFORMANCE

Like its predecessor, the UltraSPARC T1 [4], the UltraSPARC T2 based systems are oriented towards a wide variety of applications, including web servers, database and application servers, Java Applications, search, streaming video, and telco applications. With twice the threads, the UltraSPARC T2 can achieve a much higher throughput for these applications. Additionally, with enhanced floating point, integrated 10 GbE networking, and cryptographic coprocessors, UltraSPARC T2 greatly extends its application space to cover areas such as financial modeling, High Performance Computing (HPC), secure networking, and network file servers.

The high throughput performance of UltraSPARC T2 makes it valuable for consolidation and virtualization projects. Using Sun's Logical Domains (LDOMs) virtualization technology, a single UltraSPARC T2 processor can support up to 64 independent Solaris images. Combining LDOMs with Solaris Containers can achieve even higher levels of consolidation.

We have tested UltraSPARC T2 based systems on a wide variety of Applications and benchmarks. As part of the UltraSPARC T2 processor launch on August 7, 2007 we published SpecCPU2006 numbers. SpecCPU2006 is the SPEC's cpu intensive benchmark suite, stressing a system's processor, memory subsystem and compiler.

There is both an integer and floating point portion to the benchmark. Each section is further subdivided into a single instance (Ratio) and multiple instances (Rate) versions. Due to the throughput nature of CMT processors we have published Rate numbers for both the Integer and Floating Point benchmarks in Table 2.

TABLE 2
UltraSPARC T2 SPEC CPU2006 PERFORMANCE^a

Workload	Peak	Base
SPECint_rate2006	78.3 estimated	72.9 estimated
SPECfp_rate2006	62.3 estimated	57.9 estimated

^aAll SPEC CPU metrics are quoted from full "reportable" runs, but are designated as "estimates" because they used pre-production systems. SPEC, SPECint, SPECfp reg tm of Standard Performance Evaluation Corporation. Results from www.spec.org as of 8/6/2007. Sun UltraSPARC T2 @ 1.4 GHz (64 threads, 8 cores, 1 chip) 78.3 est. SPECint_rate2006, 62.3 est. SPECfp_rate2006.

A number of customers and Universities have also tested the performance of the UltraSPARC T2 processor. Williams, et al [5] use sparse matrix-vector multiply (SpMV) – one of the most heavily used kernels in scientific computing – to compare a 1.4 GHz UltraSPARC T2, with 2.2 GHz AMD dual-core Opteron and 2.33 GHz Intel quad-core Clovertown. Table 3 shows the memory bandwidth in GB/s, percentage of peak bandwidth, performance in GFlop/s and percentage of peak computation for all 3 designs. The results show that the combination of high memory bandwidth and 8 floating point units deliver superior performance on the UltraSPARC T2.

TABLE 3
PEAK MEMORY BANDWIDTH AND GFLOPS FOR spMV

Processor	Memory BW (% of peak)	GFlops (% of peak)
UltraSPARC T2	23.3 GB/s (54.6%)	5.8 GF/s (51.8%)
AMD Dual Core	13.4 GB/s (62.6%)	3.3GF/s (18.8%)
Intel Quad Core	11.2 GB/s (52.7%)	2.8 GF/s (3.7%)

The eight integrated cryptographic units give unprecedented performance when compared to traditional processors which implement ciphers and hashes in software. Table 4 highlights this difference in performance for the RSA 1024 sign operations and the AES 128 Bulk Cipher on 1.4 GHz UltraSPARC T2, 2.2 GHz AMD dual core Opteron, and 2.66 GHz Intel Quad Core Clovertown.

TABLE 4
CRYPTOGRAPHIC PERFORMANCE

Processor	RSA 1024	AES128
UltraSPARC T2	37k Operations/s	36 Gbits/s
AMD Dual Core	2.3k Operations/s	1.6 Gbits/s
Intel Quad Core	4.8k Operations/s	4.2 Gbits/s

ACKNOWLEDGEMENTS

The authors acknowledge the contributions from the UltraSPARC T2 CPU development team at Sun Microsystems.

REFERENCES

- [1] J. L. Hennessy, and D. A. Patterson, *Computer Architecture: A Quantitative Approach*, Fourth Edition, Morgan Kaufmann, San Francisco, CA. 2007.
- [2] G. Grohoski, "Niagara-2: A Highly Threaded Server-on-a-Chip," 18th Hot Chips Symposium, Aug., 2006.
- [3] U. G. Nawathe, et al., "An 8-Core 64 Thread 64b Power-Efficient SPARC SoC," in IEEE International Solid-State Circuits Conference (ISSCC), Dig. Tech Papers, Feb. 2007, pp. 108-110.
- [4] A. S. Leon, et al., "A Power-Efficient High Throughput 32-Thread SPARC Processor," in IEEE International Solid-State Circuits Conference (ISSCC) Dig. Tech Papers, Feb. 2006, pp. 98-99.
- [5] Samuel Williams, Leonid Oliker, Richard Vuduc, John Shalf, Katherine Yelick, James Demmel "Optimization of Sparse Matrix-Vector Multiplication on Emerging Multicore Platforms," unpublished.