

The UltraSPARC® T2 Processor

Multicore and multithreaded processor with integrated networking and cryptography

Overview

The UltraSPARC T2 processor is a highly integrated multicore and multithreaded system on a chip (SoC) processor featuring Sun's unique chip multithreading technology. Because 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. The UltraSPARC T2 processor incorporates four, six, or eight independent 64-bit

Each core integrates 16 KB, eight-way associative instruction Level 1 cache (L1) with 8 KB, four-way associative Level 1 data cache. The eight SPARC processor cores share an eight-banked, 4 MB 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 cache enable eight simultaneous accesses to support the high bandwidth requirements of the UltraSPARC T2 processor.



Highlights

- Chip multithreading (CMT) technology for massive throughput
- Integrated cryptographic acceleration and 10 Gigabit Ethernet (GbE) enables secure computing at wire speed
- Integrated networking reduces board space, power, and system cost
- Scalable, multithreaded solution enables equipment suppliers to deliver software-based upgrades to improve time to market
- Designed with virtualization in mind to provide hardware-based domain protection

Key Features	Benefits
Integrated dual 10 GbE networking with XAUI	<ul style="list-style-type: none"> • High integration, smaller footprint versus multichip solutions • Reduces components and power, and improves processor performance • Reduces latency and provides higher memory bandwidth
Chip multithreading	<ul style="list-style-type: none"> • Optimizes throughput by increasing CPU utilization
Multicore/multithread	<ul style="list-style-type: none"> • Scalable solution at the thread and core levels
Highest memory bandwidth	<ul style="list-style-type: none"> • Supports high-throughput applications
Per-processor cryptography	<ul style="list-style-type: none"> • Optimizes cryptography performance
Control Plane and Data Plane consolidation	<ul style="list-style-type: none"> • Dramatic reduction of components, power, and board space

SPARC® cores, with each core having full hardware support for execution of eight independent threads. A thread can be a process that is part of a parallel program, or a program by itself. Each SPARC processor core consists of two integer execution units, a floating-point and graphics unit, and a cryptographic stream processing unit. Threads are statistically assigned to each execution unit, four threads per integer execution unit.

The processor cores communicate to the L2 cache through a nonblocking crossbar switch. The crossbar connects the eight SPARC processor cores to the eight banks of the L2 cache and the I/O port. The L2 cache connects to four on-chip DRAM controllers, which interface directly with a pair of fully buffered (FB-DIMM) channels.

In addition, the SoC includes an on-chip PCI Express (PCIe) controller and two native 1 Gigabit and 10 Gigabit Ethernet (GbE) media access controllers (MACs). The PCIe implements PCI Express Base Specification 1.0a and supports x1, x4, and x8 configurations. The network interface unit (NIU) implements two on-chip 10 GbE ports with XAUI interfaces. The PCIe and XAUI interfaces are programmable to drive short and long distances to traverse a backplane for blade applications.

The high level of system integration on the UltraSPARC T2 processor reduces overall system complexity, component count, and power consumption. The UltraSPARC T2 processor combines the benefits of general-purpose processing and flexible programming, along with the scalability of chip multithreading. This offers the ability to efficiently consolidate Control Plane and Data Plane operations — as well as security functions — on a single chip.

Feature List

PCI Express interface unit

- Operates at 2.5 GHz per lane per direction, differential signals with variable output levels to optimize for long- or short-reach interfaces
- Implements the root complex behavior of the PCI Express Base Specification 1.0a
- Implements the transaction layer, data link layer, and logical subblock of the physical layer
- Supports x1, x4, and x8 configurations at the data rate of 2.5 Gb/sec
- Supports lane reversal

Network interface unit

- Two 10 Gb/sec line interfaces
- Two XAUI interfaces
 - Each XAUI interface consists of four-lane CML, 3.125 GHz differential signals, with variable output levels to optimize for long- or short-reach interfaces. Four transmit and receive pairs with integrated CDR are provided per interface
 - Supports two-wire MDIO control signals
- Packet distribution and order sequencing in hardware
- Line rate packet classification based on layer 1/2/3/4 protocol stack
- Multiple DMA engines with DMA port binding
- Hardware virtualization support

10 GbE MAC

- Two dual-speed, full-duplex MACs (1 Gb/10 Gb)
 - 16 unique MAC addresses per port
 - Jumbo frame support to 9,216 bytes
 - PAUSE frame support per IEEE 802.3
 - Optional PAUSE generation via control registers
 - Option for buffering of entire egress packets before transmitting them to the line port, enabling discard of errored packets before transmission
- Store-and-forward operation
- Error checking
 - Frame integrity and frame length checks

- CRC checking/generation and optional bypass mode
- Option for discard of packets with CRC errors in receive direction
- Remote/local fault signaling at reconciliation sublayer (RS)
- Automatic padding of transmitted packets of less than minimum frame size
- Support for big-endian data formats
- Programmable inter-frame gap (IFG)
- Capable of variable idle insertion to support WAN interconnect sublayer (WIS) data pacing

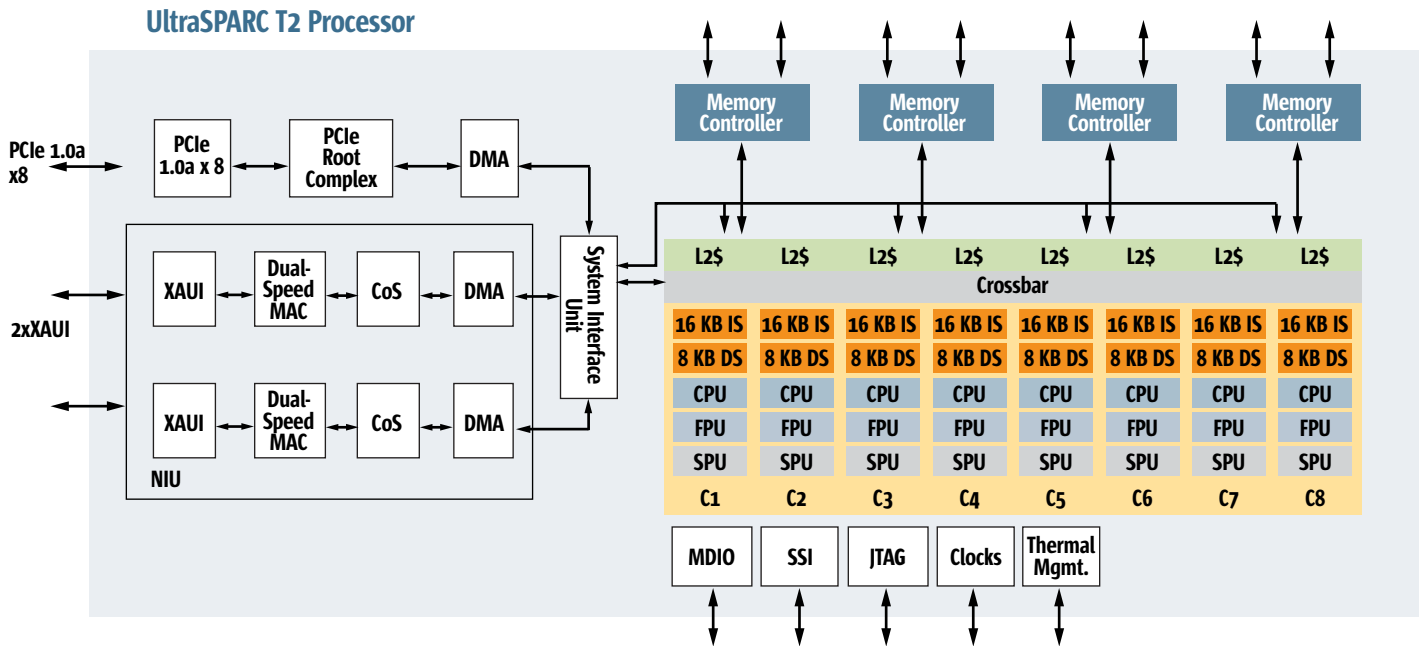


Table 1. Transmit and Receive Statistics Counters

Receive Statistics	Transmit Statistics
Total octets	Total octets
Received bytes	Good frames
Broadcast frames	Link fault
Multicast frames	
Fragmented frames	
Frames with 64 octets	
Frames with 65 to 127 octets	
Frames with 128 to 255 octets	
Frames with 256 to 511 octets	
Frames with 512 to 1,023 octets	
Frames with 1,024 to 1,522 octets	
Frames with 1,523 to max. octets	
CRC errors	
Oversized frames	
Errored packet 1 Gb link or 10 Gb link	

Packet classification and CoS

- Classification
 - Port
 - Destination MAC address
 - EtherType
 - Virtual local area network (VLAN)
 - TOS traffic class
 - Proportional integral derivative (PID) controller/next header
 - IP source address
 - IP destination address
 - L4 port/serial peripheral interface (SPI)
- Hash compute across port, destination MAC, VLAN, PID/next header, IP source address, IP destination address, L4 port
- Ternary content addressable memory (TCAM) (tuple) across TOS/traffic class, PID/next header, IP source address, IP destination address, L4 port. TCAM size 200 bits wide x 128 entries deep
- Congestion management used to prevent overloading of the CPU and to fend off potential SYN attacks when resources are low

- Uses weighted random early discard (WRED) where a threshold and window are defined
- Rx packets are dropped when congestion is detected

DMA engines

- 16 independent transmit (Tx) and receive (Rx) DMA engines
- Match DMAs to threads
 - Binding flexibility between DMAs and ports
- Bandwidth is allocated according to a deficit round-robin scheduler
 - Rx packets are dropped when congestion is detected
- TCP/UDP/IP checksum offload
- Virtualization support
 - DMA resource separation by logical groups
 - Descriptor address relocation

World's first 64-thread, highly integrated multicore system on a chip (SoC)

Interrupts

- Interrupt coalescing
- Mailbox
- INTX, MSI, MSI-X

SPARC core

- SPARC V9-compliant CMT with eight threads/core at 1.0 GHz and 1.2 GHz
- The eight threads share instruction and data caches and are divided into two groups per core
- 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
- Up to two instructions can issue per cycle, one per thread group
- The UltraSPARC T2 processor's fine-grained multithreading scheme minimizes throughput performance losses arising from branch miss predictions and cache misses
- Two integer execution pipelines/core
- Fully pipelined floating-point/graphics unit
- Stream processing unit for cryptographic acceleration (see advanced crypto/security engine)
- Cache
 - L1 caches and the translation lookaside buffers (TLBs) are shared by all eight threads
 - L1 instruction cache size: 16 KB
 - L1 data cache size: 8 KB
 - L2 cache size: 4 MB, eight-way banked with a set associativity of 16 in order to meet the bandwidth demands of eight cores, each having eight threads per core
 - Full L2 cache coherency — any processor — any bank

Advanced crypto/security engine

- Full duplex encryption/decryption
- Supports bulk encryption, authentication operations, and public-key cryptography
- High-speed crypto engine per core supporting RC4, DES/3DES, AES (128-, 192-, and 256-bit key lengths), AES-GCM
- AES modes of operation supported: ECB, CBC, and CTR
- DES modes of operation supported: ECB, CBC, and CFB
- MD5, SHA-1, SHA-256 authentication functions including keyed hashing (HMACs)
- Key exchange, key generation, and authentication acceleration for security protocols such as IPsec and SSL/TLS
- Provides direct hardware support for RSA, DSA, and Diffie-Hellman operations for up to 2,048 bit keys via modular exponentiation
- Supports key exchange and signature generation and verification using elliptic curve cryptography
- Additional cypher support available on a per-thread basis; contact factory for details

Memory controller

- Four on-chip memory controllers
- Dual-channel FB-DIMM ports are associated with the adjacent pair of L2 banks
- 10-bit southbound and 14-bit northbound FB-DIMM channel protocols
- Supports 256 Mb, 512 Mb, 1 Gb, and 2 Gb DRAM components
- Supports 128 bits of write data and 16 bits ECC per SDRAM cycle, and 256 bits of read data and 32 bits ECC per SDRAM cycle
- ECC generation, check, correction

Test and diagnostic support

- IEEE 1149.1 port with memory BIST, scan, and JTAG boundary scan

Power management

- Minimal speculation
- Extensive clock gating on data path, control, and array structures
- Core shutoff
- Memory throttling
- External power throttling. Based on the state of the three power throttle pins, stall cycles are injected into the processor core pipeline to reduce overall dynamic power consumption
- On-chip thermal diodes enable the system to regulate the die temperature by controlling the instruction issue rate and by disabling threads

Core reliability and serviceability

- Major hardware structures are protected with ECC or parity as appropriate
- Extensive RAS features
- Combination hardware and software correction flows
- ECC on integer, floating point, store buffer data, trap stack, and other internal arrays

General specifications

- Packaging: 45mm x 45mm, 1,831 ball flip-chip glass ceramic BGA
- RoHS 6
- 1mm ball pitch
- Technology: 11 metal, 65nm, ltriple-Vt CMOS 1.1 V core; 1.1 V/1.5 V I/Os
- Ambient operating conditions: 0° C to 40° C

- SSI ROM interface at 50 Mbps serial interface connects to an external FPGA to interface with BOOT ROM
 - Supports PIO accesses, thus supporting optional CSRs or other interfaces within the FPGA
- External clocks required
 - XAUI clock
 - PCIe clock
 - Core clock
 - FB-DIMM clock

Standards compliance

IEEE 802.3-2005

IEEE 802.1q for VLAN support

IEEE 802.3x for flow control and PAUSE frame generation

IEEE 802.3ae-2005 for 10 GbE MAC

Typical applications

- Application delivery systems
- Intrusion detection systems
- Multiservice Ethernet switches
- Routers
- Security appliances
- Video streaming and IPTV
- WAN optimization
- Wireless LAN controllers
- Wireless Control Plane and Data Plane processors for CDMA, UMTS, WiMax, and LTE


Table 2. UltraSPARC T2 Processor Part Numbers

Part Number	Core	Frequency (GHz)
SME1908ABGA-1170-8	8	1.2
SME1908ABGA-1000-6	6	1.0
SME1908ABGA-1000-4	4	1.0

Table 3. UltraSPARC T2 Processor Maximum Power at 1.1 V

Number of Cores	Frequency (GHz)	Power (Max. Watts)
8	1.2	91
8	0.9	65
6	0.9	57
4	0.9	46

10 GbE appliance with redundant I/O, based on UltraSPARC T2 processor



UltraSPARC T2 processor integrates key components to minimize power, space, and cost.

Figure 1. 10 GbE appliances: application delivery systems, intrusion detection systems, security, wireless LAN controllers.

12x/24x GbE appliance based on UltraSPARC T2 processor





Figure 2. 12x or 24x 10 GbE appliances: application delivery systems, intrusion detection systems, security, wireless LAN controllers.

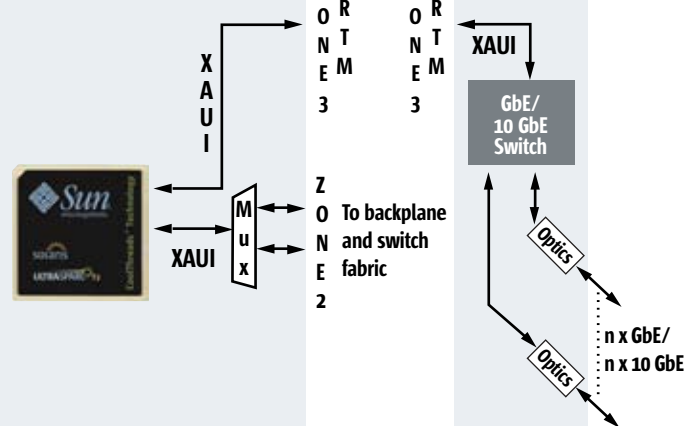
Processor blade



UltraSPARC T2 processor integrates processing power and key components into a single device, often replacing up to three devices, to minimize power, space, and cost in a high-throughput solution.

Figure 3. ATCA processor blade with redundant switch fabric interfaces.

Processor blade



Dual 10 GbE XAUI I/Os enable straightforward configuration changes with minimal new components and complexity.

Figure 4. ATCA blade with n x GbE or n x 10 GbE interfaces on an RTM.