Power Savings in the UltraSPARC T1 Processor

David Greenhill, Distinguished Engineer
Jerome Alabado, Staff Engineer
Sun Microsystems Inc.

December 2005

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 U.S.A.
www.sun.com

1.0 Introduction

Power and power density limits constitute the primary constraints for next-generation processor and system designs.

Power and power density limits constitute the primary constraints for next-generation microprocessor and system designs. As a result of ever increasing transistor budgets and demand for computation, server processors today are dissipating well over 100 W and continue to push the limits of what is both physically and cost-effectively attainable. If this alarming trend continues, then contemporary chips will produce as much heat, for their proportional size, as a nuclear reactor.

In order to address this issue, Sun designed the UltraSPARC T1 processor from the ground-up, employing multiple architectural and circuit techniques to maximize performance per Watt while minimizing power density. Ranging from architecture to thermal monitoring and many other techniques such as clock gating, power-efficient L2 cache accesses, and a balanced clock distribution system, the UltraSPARC T1 processor's power savings is a fundamentally new approach that enables:

The UltraSPARC T1 processor's power and thermal management lowers the cost of power, space, and cooling in data centers.

- **Lower cost of power, space, and cooling**
  Reducing power consumption of processors and associated components directly translates to lower power and cooling costs. Another effect is lower real estate costs. Reducing the amount needed to cool systems reduces the need for additional air conditioner units, which consume valuable space on the data center raised floor area.

The UltraSPARC T1 processor's low power and power density allows full utilization of rack space.

- **Higher utilization of rack space**
  Most data centers today cannot support fully-populated racks of high power servers. The UltraSPARC T1 processor's low power and power density allows full utilization of the available rack space.

Power and thermal management increases the UltraSPARC T1 processor's reliability and extends its useful lifetime.

- **Higher levels of reliability**
  Managing power effectively reduces the natural wearout and damage of the processor, thereby increasing its reliability and extending its useful lifetime. At the International Reliability Physics Symposium[1], Sun showed that implementing power and thermal management features can dramatically increase both the lifetime and reliability of the device by up to 24 times while maintaining or improving device performance.

1. International Reliability Physics Symposium, April 2005, Sun Microsystems, Inc., http://www.irps.org/05-43rd/IRPS_Keynote_Yen.pdf

2.0 UltraSPARC T1 Processor Power Savings

Power savings in the UltraSPARC T1 processor are attributed to:

- Architecture
- Circuit Techniques
- Thermal Monitoring
- Memory Throttling

The UltraSPARC T1 processor architecture is very power efficient by operating at low clock frequencies with high utilization of processor resources.

Architecture

For several decades, Moore's Law has provided processor architects with an ever increasing transistor budget. The result of many designs has led to faster, but power-hungry processors. Because of the changing attributes of commercial server workloads and the memory bottleneck, complex superscalar processors today spend a significant portion of their resources underutilized while consuming power, illustrated in Figure 1.

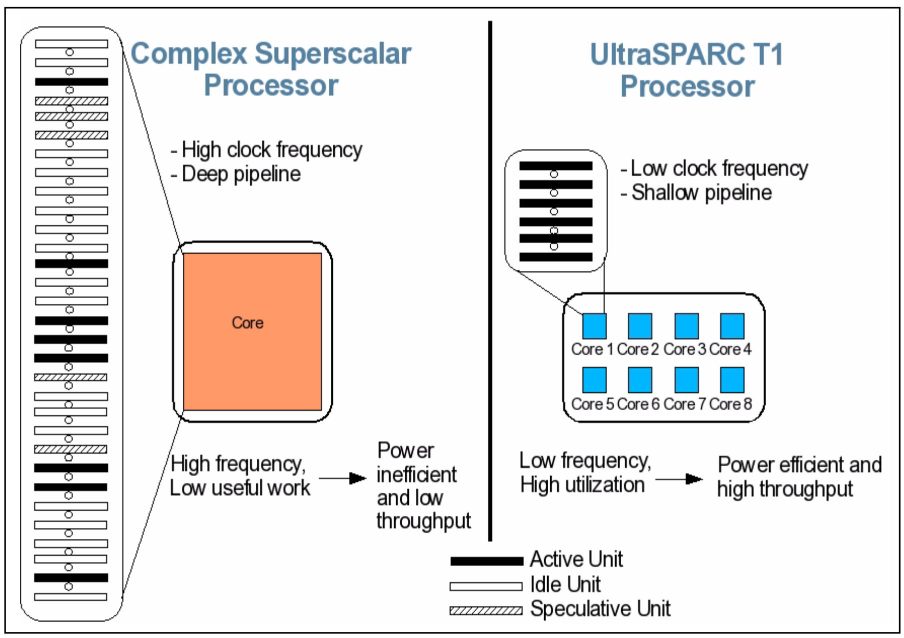

FIGURE 1 Power Efficiency: Complex Superscalar Processor vs. UltraSPARC T1 Processor

In the case of the complex superscalar processor, which today operates at a high clock frequency and implemented with a deep pipeline, many power-hungry units are clocked even when the units remain idle, or will have their results discarded due to bad speculation. This design approach results in very low efficiency and high power. In sharp contrast, the UltraSPARC T1 processor is much more power efficient than the complex superscalar processor by fully utilizing its pipeline and clocking it at a low frequency. The end result of this new approach is a processor that delivers an order of magnitude better performance per Watt than complex superscalar processors of today.

The UltraSPARC T1 processor uses clock gating and a balanced clock distribution scheme to reduce power consumption.

Circuit Techniques

Circuit techniques implemented in the UltraSPARC T1 processor to reduce power consumption include:

- *Clock Gating*
  Clock gating effectively turns off certain parts of the processor that are not in use to conserve power, with no performance impact. In the UltraSPARC T1 processor, clock gating techniques include coarse-grain, to disable selective cores, and fine-grain, to disable pieces of circuitry to reduce power consumption.
- *Clock Distribution*
  The processor clock system, composed of the clock interconnection network and timing elements, is one of the main contributors to processor power. A balanced clock distribution scheme (H-tree) is used to distribute the processor global clock.
- *Power-Efficient Timing Elements*
  Timing elements such as master-slave flip-flops have been enhanced for power-efficiency.
- *Power-Efficient L2 Cache Accesses*
  The UltraSPARC T1 processor's on-chip L2 cache is divided into 4 independent banks, each consisting of 12 panels. When the L2 cache is accessed for requested data, only the panel that contains the requested data is activated. Clearly, this technique is more power-efficient than activating the entire 48 panels to access only a small segment of the cache.

The UltraSPARC T1 processor throttles its power consumption by disabling individual threads or disabling cores.

Thermal Monitoring

The UltraSPARC T1 processor utilizes two on-chip thermal sensors, located near the cores (the parts of a processor that generate the most heat), enabling system hardware to measure the processor temperature and trigger one of the following "throttling" measures:

- Disable individual threads within a core, effectively forcing the threads to enter an idle state
- Disable entire cores

By reducing processor utilization, throttling reduces power consumption. This active temperature monitor and control allow the UltraSPARC T1 processor to limit its temperature in the case of system failures such as fan failure or other external factors. When conditions return to within the normal range, the idle threads are activated and resume execution.

2.1 Results

For a high-performance processor like the UltraSPARC T1 processor to dissipate merely 72 W and have low power density is unprecedented.

On top of the performance gains the UltraSPARC T1 processor delivers, the processor typically dissipates merely 72 W (at 1.3 V, 1.2 GHz). For complex superscalar processors to match the UltraSPARC T1 processor's performance, several processors would be needed, but at a significant power cost. Another result of the UltraSPARC T1 processor's power savings is a reduction in power density. Illustrated in Figure 2, the power density of a single-core processor is compared to that of an UltraSPARC T1 processor. In the case of the single-core processor, several "hot spots" (junction temperatures exceeding 100°C) are located in many parts of the processor. In contrast, the junction temperature of the UltraSPARC T1 processor is less than 75°C[1] and is much more distributed than that of the single-core processor, thereby eliminating hot spots.

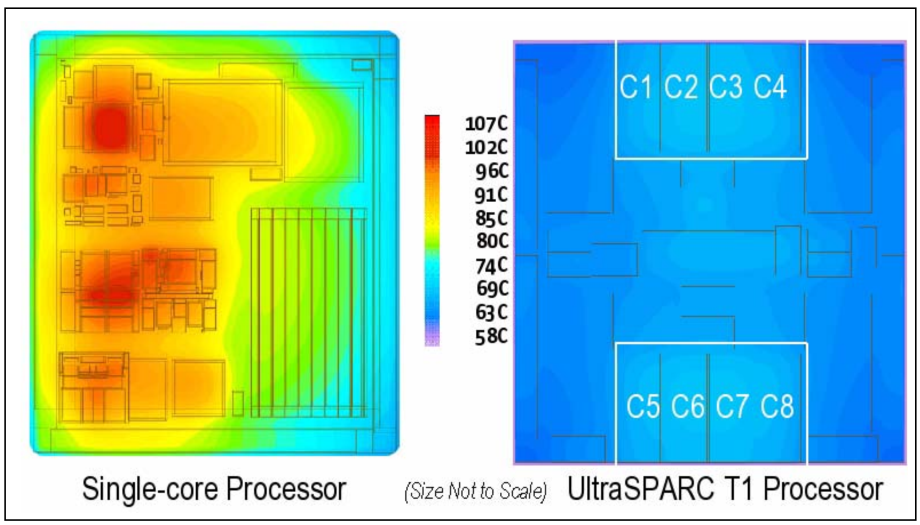

FIGURE 2 Power Density: Single-Core Processor vs. UltraSPARC T1 Processor

1. Compared to a typical junction temperature of 105°C, the UltraSPARC T1 processor's reliability improves by 5 times.

3.0 Processor and Memory Power: UltraSPARC T1 Processor vs. Xeon

The processor and memory power constitute a significant portion of the overall system power. In fact, these two components alone can contribute to more than half of the total system power. With the UltraSPARC T1 processor's highly-integrated design, the processor has many advantages over the competition.

The differences in processor and memory architecture between the UltraSPARC T1 processor and Xeon consequently leads to differences in power consumption.

Consider the difference in processor and memory power between a system based on the UltraSPARC T1 processor and a system based on Xeon. Unlike the UltraSPARC T1 processor, which has 4 memory controllers integrated on-chip, Xeon processors require external chipsets to connect to main memory. Commonly known as the "Northbridge" chipset, these off-chip processors serve as the Xeon's memory controller. The Northbridge chipset architecture not only becomes a major source of bottleneck, but also contributes to the overall system power. The differences between an UltraSPARC T1 processor-based system and a typical two-processor Xeon system is illustrated in Figure 3.

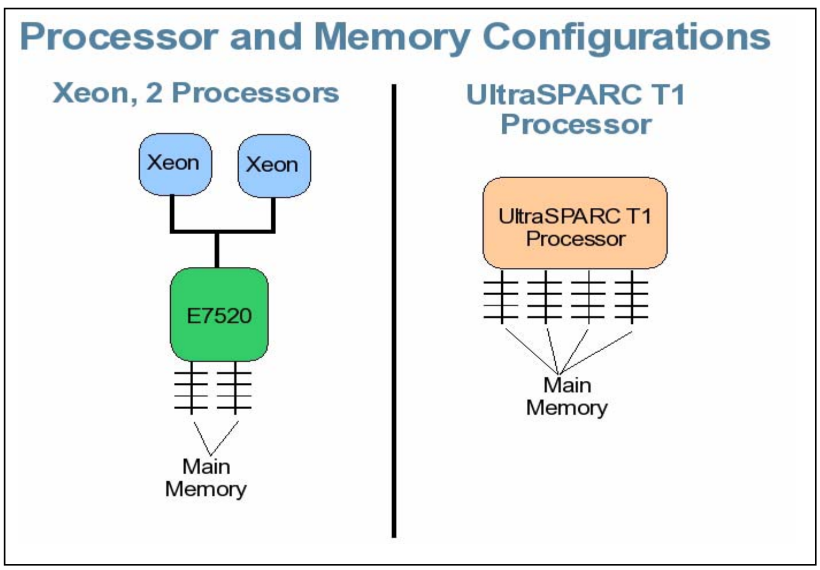

FIGURE 3 Processor and Memory Power: UltraSPARC T1 Processor vs. Xeon (2 Processors)

Two Xeon processors connect to the E7520 chipset (memory controller) to access 16 GB of main memory. In the case of the UltraSPARC T1 processor, the memory controllers are integrated on-chip so that the processor can directly access main memory. The breakdown of power consumption in these two configurations is shown in Table 1.

TABLE 1 Power Contributions

System	Xeon	UltraSPARC T1 Processor
Processor	103 W[a] x 2 = 206 W	72 W
Memory Controller	9 W[b]	0 W (included in processor)
16 GB Memory	54 W	54 W
Total Power	269 W	126 W
Relative Power	2.1X	1.0X

a.Intel Xeon Datasheet, ftp://download.intel.com/design/Xeon/datashts/30235501.pdf

b.Intel® E7520/E7320/E7525 Memory Controller Hub (MCH) Thermal/Mechanical Design Guide, ftp://download.intel.com/design/chipsets/designex/30240303.pdf

The UltraSPARC T1 processor configuration consumes over 2x less power than a 2-processor Xeon configuration for the same memory capacity.

In addition to the performance gains the UltraSPARC T1 processor delivers, an UltraSPARC T1 processor-based system consumes over 2X less power than a 2-processor Xeon system for the same memory capacity.

The difference in power is even more significant in 4-processor Xeon systems. When configuring 4 Xeon processors, a different version of Xeon is used (Xeon MP) that dissipates even higher power (129 W compared to 103 W) and requires more external chipsets. The differences between an

UltraSPARC T1 processor-based system and a 4-processor Xeon system is illustrated in Figure 4.

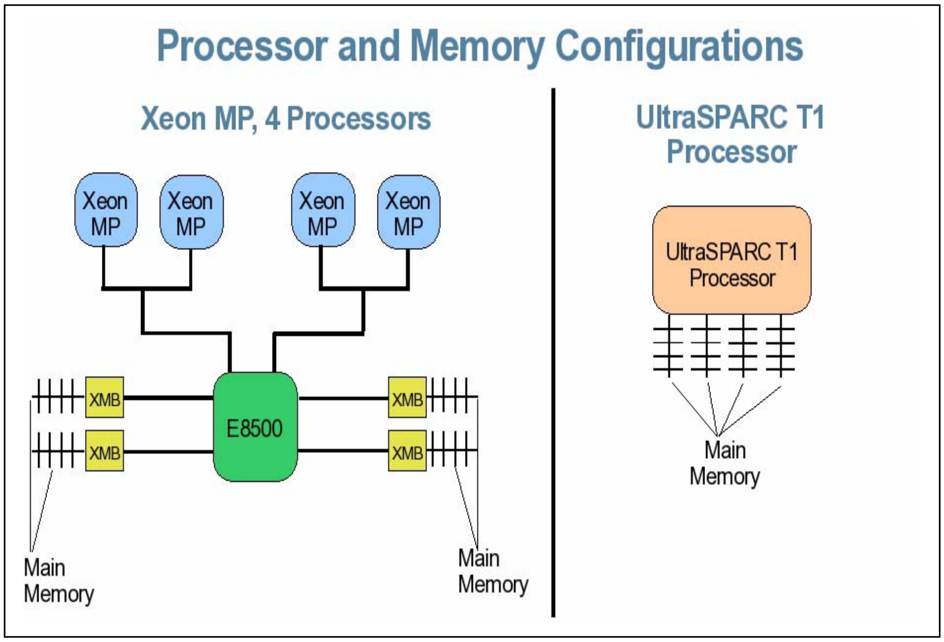

FIGURE 4 Processor and Memory Power: UltraSPARC T1 processor vs. Xeon (4 Processors)

The differences between the Xeon and the UltraSPARC T1 processor configurations still apply to those from the previous example. In Figure 4, each system now consists of 32 GB of main memory, and the 4-processor Xeon configuration requires additional external memory bridge (noted as "XMB" in Figure 4) chipsets for each memory interface. The breakdown of power consumption in these two configurations is shown in Table 2.

The UltraSPARC T1 processor configuration consumes almost 4X less power than a 4-processor Xeon configuration for the same memory capacity.

TABLE 2 Power Contributions

System	Xeon	UltraSPARC T1 Processor
Processor	129 W[a] x 4 = 516 W	72 W
Memory Controller	24.5W[b] (includes accompanying XMB chipset)	0 (included in processor)
32 GB Memory	91 W	91 W
Total Power	631.5 W	163 W
Relative Power	3.9X	1.0X

a.Intel Xeon MP Datasheet, ftp://download.intel.com/design/Xeon/datashts/30675401.pdf

b.Intel® E8500 Chipset North Bridge (NB) and eXternal Memory Bridge (XMB) Thermal/Mechanical Design Guide, ftp://download.intel.com/design/chipsets/e8500/guides/30674901.pdf

4.0 Conclusion

Many opportunities to reduce power consumption still exist at the system level.

Next-generation CMT processors will continue to deliver performance per Watt advantages.

Power savings in the UltraSPARC T1 processor is attributed to circuit techniques, thermal monitoring, and most importantly, its underlying architecture. With its power-efficient architecture, the UltraSPARC T1 processor is suitable for dense, rack-mount solutions. Moving forward, many opportunities still exist at the system level to reduce power consumption. Other components, such as memory, disk drives, I/O slots, and power supply, are being considered for the most power-efficient use of these components. Additionally, next-generation CMT processors will continue to improve on the performance per Watt advantages that the UltraSPARC T1 processor already achieves.