

The UltraSPARC T1 Processor - Power Efficient Throughput Computing

Rick Hetherington
Chief Architect and Distinguished Engineer

December 2005



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

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, California 95045 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology described in this document. In particular, and without limitation, these intellectual property rights may include one or more patents or pending patent applications in the U.S. or other countries.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, and Solaris 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 US and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95045 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company Ltd.

Sun, Sun Microsystems, le logo Sun, Java, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPONDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.

While several vendors have attempted to mitigate the power issues using traditional methods, Sun's UltraSPARC T1 processor offers the most innovative and effective solution. Designed from the ground-up, the UltraSPARC T1 processor employs multiple architectural techniques to maximize performance per power while minimizing power density. This fundamentally new approach to processor design enables the UltraSPARC T1 processor to be more suitable for dense rack-mount implementations in the data center than other contemporary processors.

FIGURE 2 in the paper (replicated as an icon to the right of the paper's title), demonstrates a significant differentiation of the UltraSPARC T1 processor. Compared with the power density trend of other processors, the UltraSPARC T1 processor is in a class of its own. The UltraSPARC T1 processor's power density introduces a new paradigm for computing.

The Challenge

As the capabilities of computers have evolved to meet ever increasing computational needs, so has the infrastructure required to support these vital systems. Data centers are reaching a state where power and cooling requirements are no longer secondary concerns, but must be accounted for up front whenever compute investments are considered. In some enterprise environments, the power and cooling requirements are pushing the limits of what is attainable, not to mention the budgetary strains—both for initial upgrades of power and cooling systems and for their ongoing energy costs. In some regions, governmental policies are putting limitations on energy use. At the same time, the next generation of computers are expected to deliver higher performance to support growing compute demands.

The evolution of microprocessor design has led to processors with higher clock frequencies to improve single-thread performance. These processors exploit instruction-level parallelism (ILP) to speed up single-threaded applications. ILP attempts to increase performance by determining, in real time, instructions that can be executed in parallel. The trade-off is that ILP extraction requires highly complex microprocessors that consume a significant amount of power. Voltage and frequency scaling are techniques used to scale ILP-focused microprocessors by sacrificing performance to meet power budgets.

Another technology that results in lowering processor power demands is chip multiprocessors (CMPs)—multiple processors on a single silicon die. CMPs inherently reduce processor power by eliminating the long power-hungry wires present in large monolithic ILP-focused microprocessors. While all major microprocessor vendors are taking advantage of CMP technology, most are doing so while ignoring the main characteristic of commercial applications—an abundance of threads. Commercial applications are rich in threads but poor in instruction-level parallelism, so the ILP exploitation, scaled or not, doesn't offer any performance advantage and continues down the path of increasing power demands.

At first glance it appears that improving processor performance and lowering processor power demands are at odds, but this doesn't have to be the case. Multithreading (MT)—running multiple threads per processor core—hides the frequent high-latency events and exploits the thread-level parallelism (TLP) common in commercial applications. Combining simple core CMPs with MT results in a lower power, higher throughput processor. We call this *Chip Multithreading (CMT)*, and it is an ideal solution for commercial applications.

The UltraSPARC T1 processor – CMT Realized

The first instance of a radical CMT processor is the UltraSPARC T1 processor. The UltraSPARC T1 processor is a SPARC® V9, single-chip processor designed to exploit TLP by employing fine-grain multithreading.

The UltraSPARC T1 processor has eight independent 64-bit execution pipelines (cores). Each core is capable of selecting from four active threads. The result is a processor that provides for up to 32 different threads or processes to be executing simultaneously on a single chip (FIGURE 1). The aggregate throughput is approximately 5 to 15 times better than contemporary processors.

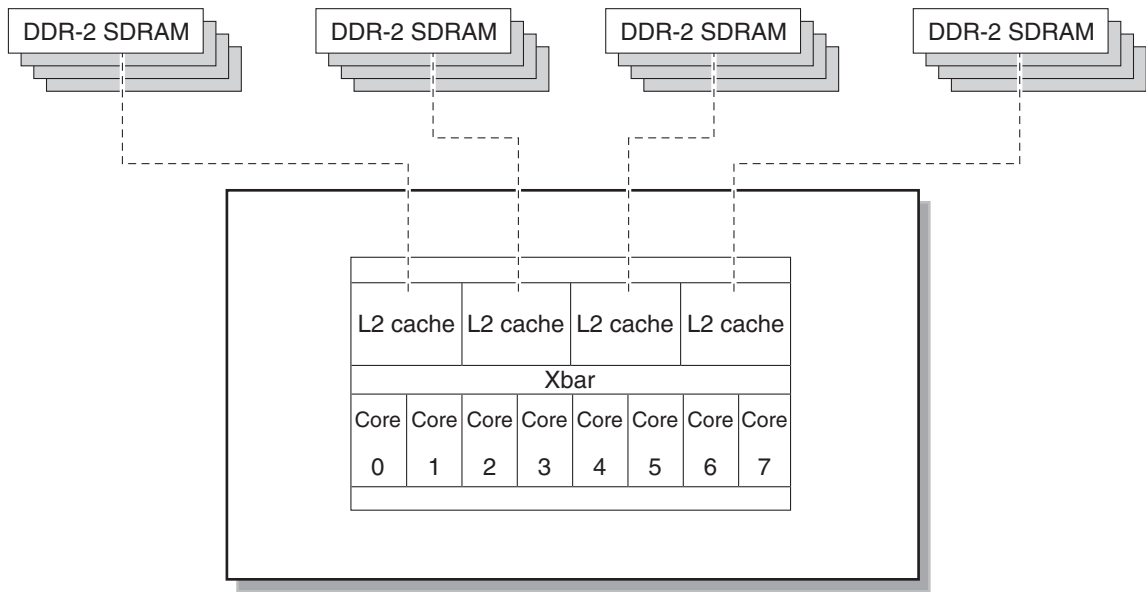


FIGURE 1 UltraSPARC T1 Processor Block Diagram

In addition, as shown in FIGURE 2, the power density for the UltraSPARC T1 processor is several times less than other contemporary processors, making the UltraSPARC T1 processor more suitable for dense rack-mount installations in the data center. You can see the power density increasing over time with improved transistor density and increased processor complexity.




FIGURE 2 Power Density (W/cm^2) Over Time

The UltraSPARC T1 processor design goal was to achieve maximum throughput at the lowest possible power consumption. In addition, the UltraSPARC T1 processor required appropriate controls to guard against the pathological peak current draw. Energy management in the UltraSPARC T1 processor has three distinct aspects:

- Inherent architecture of simple-core chip multiprocessors with fine-grain multithreading
- Specific hardware features targeted at power throttling
- Energy management assistance from software

Power Efficiency Derived From the UltraSPARC T1 Processor-Style CMT

Nearly all of the UltraSPARC T1 processor's competition in the low-end horizontally-scaled market employ pipeline designs primarily targeted at the desktop (single-threaded or ILP-focused) applications. Designers then modify cache and interface architectures to support a coherent shared memory across two, four, and eight processor systems, and reduce the operating frequency and voltage for low-power consumption. These pipelines attempt to execute instructions from a single thread in parallel (ILP) and typically have the following characteristics:

- Super-scalar designs that attempt to increase instructions per clock (IPC) and require multiple independent execution units with dependency checking, scoreboarding, complex bypassing of data paths, and complex register files.
- Very high clock frequency (2.2 GHz to 3.6 GHz) that demands hot, fast, and leaky transistors as well as dynamic circuit designs.
- Deep pipelines that have as many as 30 pipeline stages from instruction fetch to retirement, requiring aggressive branch prediction and a vast number of state devices.
- Out-of-order execution with wide instruction issue windows, reorder buffers, and a variety of memory speculation techniques and apparatus.
- Huge multilevel caches, three and four levels in some cases, to further combat memory latency.

The result is increased single-threaded performance on applications with high ILP. But performance on the typical commercial application is limited by a serial and highly dependent instruction stream. The features mentioned above become a power consuming burden when executing code void of ILP and high in cache misses.

The UltraSPARC T1 processor employs a fine grain multithreaded pipeline design that is optimized for TLP applications, thereby achieving high throughput. This unique architecture is a result of matching processor design to the targeted

application. There is no speculation at all, which means there is virtually nothing that is discarded, flushed, or inaccurately predicted. The UltraSPARC T1 processor has:

- Eight independent single-issue, in-order pipelines – no extra transistors and power for multiple-issue or out-of-order execution.
- No branch prediction – a thread switch occurs on a branch instruction and it is resolved in two pipeline stages.
- A single issue, in-order pipeline with a relaxed clock frequency – enables more robust and simpler processor designs using a standard cell approach from a library of static cells.
- The core pipeline has only six stages.
- Large, highly banked, highly associative, on-chip, unified secondary cache – handles most memory requests to avoid incurring external DRAM power.
- Integrated I/O and DRAM controllers – avoids having to power external components and associated busses.

Hardware Design Features Control Power Consumption

To further manage energy consumption, the UltraSPARC T1 processor has several specific hardware features. Designers concentrated on two areas that are sources of high power consumption: controlling the issue rates of the SPARC cores and limiting activity in main memory.

The UltraSPARC T1 processor has a mechanism to throttle issue rates within the cores by putting specific threads in a *sleep* mode. When power and thermal conditions return to normal, threads resume executing instructions. This operation requires an external thermal diode that produces a voltage level into a temperature control device. This device then asserts a level on the Temp_Trig pin (an input to the UltraSPARC T1 processor). This signal, when asserted, delivers idle interrupts to software designated threads. When conditions return to within the normal range, Temp_Trig de-asserts and a *resume* type interrupt is issued to the idle threads.

At the end of the UltraSPARC T1 processor pipeline is the second source of high power consumption, and that is the DDR2 memory controller. The UltraSPARC T1 processor has four DDR2 memory channels with each channel supporting four DIMM slots. Under peak periods, there are software enabled control registers that limit the number of open pages in the DDR memories. This throttles the rate at which reads and writes are issued resulting in a reduction of the interface power as well as the peak operating power with the DRAM devices.

The Operating System's Role in Reducing Power Consumption

Another vital piece in the UltraSPARC T1 processor power efficient architecture is the operating system. Idle loops occur more frequently than one might expect, especially on a processor with 32 threads. When a thread enters the idle loop, the Solaris OS halts that thread and resumes execution when it is ready to schedule work. This not only helps with power reduction but also limits the interaction with other active threads on that core.

The thread scheduling algorithm of the Solaris OS also takes into consideration the minimization of power. Although it has no way to gauge how a workload might behave, it can follow guidelines as to how best to spread limited workloads across the eight cores. The Solaris OS is encouraged to pack threads unless there is a performance cost. If there are 16 active light-weight processes (LWPs), it is beneficial from a power perspective to pack those threads on four cores, as opposed to spreading that load across eight cores.

The Solution

The UltraSPARC T1 processor enables an era of cool computing at Sun. The combination of high throughput and low power consumption is unprecedented in the industry. The UltraSPARC T1 processor is roughly eight times the UltraSPARC IV processor throughput measured on a number of workloads while consuming 50 percent of the power of a single UltraSPARC IV processor. When compared to offerings from IBM, Intel, and AMD, the UltraSPARC T1 processor has a significant power and performance advantage that translates into a dramatic reduction in total cost of ownership (TCO).

The UltraSPARC T1 processor is the first in a series of processors that employ fine-grain multithreading. Subsequent designs will continue to focus on industry leading throughput performance as well as power and performance. The goal is to double throughput performance and hold total power constant. This is accomplished with greater efficiency in pipeline designs, fine granularity clock gating, and advanced circuit techniques.

Competitors who pursue higher performance with the traditional ILP techniques will fall further behind the power and performance curve.

The Results

The performance and associated power consumption of the Sun Fire T2000 server with one UltraSPARC T1 processor is shown in TABLE 1.

TABLE 1 Sun Fire T2000 Server Performance and Power Consumption

Benchmark	Benchmark Score ¹	Power Consumption
SPECweb2005	14,001	334 W
SPECjbb2005	63,378 bops	306 W
SAP SD 2-Tier	950 users	311 W
SPECjAppServer2004 (two-node server)	615.64 JOPS	328 W

1. See disclosure for each of the benchmark results in the section titled "Benchmark Disclosures".

About the Author

Rick Hetherington is a Sun Distinguished Engineer and Chief Architect in the Horizontal Systems group. Rick joined SUN in 1996 as co-architect of the Millennium project and spent a brief period as VP of Engineering with a Network start-up company. Prior to SUN, Rick spent 16 years at Digital Equipment Corporation. His most recent position there was as a system architect on the Alpha processor team. Rick has nearly 30 years of experience with over 40 granted patents.

Benchmark Disclosures

1. SPECweb2005 Sun Fire T2000 (8 cores, 1 chip) 14,001 SPECweb2005, submitted to SPEC for review. SPEC, SPECweb reg tm of Standard Performance Evaluation Corporation. Results from www.spec.org as of 12/05/05.

2. SPECjbb2005 Sun Fire T2000 (1 chip, 8 cores) 63,378 bops, 15,845 bops/JVM submitted for review. SPEC, SPECjbb reg tm of Standard Performance Evaluation Corporation. Results from www.spec.org as of 12/05/05.

3. Two-tier SAP ECC 5.0 Standard Sales and Distribution (SD) benchmark Sun Fire T2000 (1 processor, 8 cores, 32 threads) 1.2 GHz UltraSPARC T1, 32 GB mem, 950 SD benchmark users, 1.91 sec avg resp, MaxDB 7.5 database, Solaris 10. SAP certification number was not available at press time, please see: www.sap.com/benchmark. Benchmark data submitted for approval: 950 SD Users (Sales & Distribution), Ave. dialog resp. time: 1.91 seconds, Throughput: Fully processed order line items/hour: 95,670, Dialog steps/hour: 287,000, SAPS: 4,780, Average DB req. time (dia/upd): 0.080 sec / 0.157 sec, CPU utilization of central server: 99%, central server OS: Solaris 10, RDBMS: MaxDB 7.5, SAP ECC Release: 5.0, Configuration: Sun Fire Model T2000, 1 processor / 8 cores / 32 threads, UltraSPARC T1, 1200 MHz, 64 KB(D) + 128 KB(I) L1 cache, 3 MB L2 cache, 32 GB main memory. SAP, R/3, mySAP reg TM of SAP AG in Germany and other countries. More info www.sap.com/benchmark.

4. SPECjAppServer2004 Sun Fire T2000 (8 cores, 1 chip) 615.64 JOPS@Standard submitted for review. SPEC, SPECjAppServer reg tm of Standard Performance Evaluation Corporation. Results from www.spec.org as of 12/05/05.