Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Sun's Throughput Servers: Paradigm Shift in Processor Design Drives Improved Business Value

Sponsored by: Sun Microsystems Inc.

Vernon Turner John Humphreys

December 2005

EXECUTIVE SUMMARY

The business value contribution of IT investments depends in part on the cost of purchasing necessary processing capacity (i.e., throughput) and the cost of operating a datacenter. If more powerful servers were also smaller, consumed less electrical power, and demanded less cooling, then the business value contribution of IT would improve. Power, heat dissipation, and datacenter square footage are significant costs, and IT departments are adding more servers to meet rising demand for computing services.

Sun Microsystems has announced new Sun Fire T1000 and T2000 servers. These servers use the new Sun UltraSPARC T1 processor, which contains up to 8 cores and can manage up to 32 processing threads. Multicore, multithread processors can achieve higher throughput because they provide a better balance between processor clock speeds and memory access latencies. Without multithreading, processor clock cycles are often wasted while the faster processor waits for data to arrive from memory.

The UltraSPARC T1 processor consumes less electrical power than the earlier UltraSPARC IIIi processor, thus reversing a trend across the industry. Lower power, less heat to dissipate, and higher compute density are a powerful combination of factors.

IDC believes that the move to multicore, multithreaded processors is a paradigm shift for Sun and for the industry. Sun has taken a truly different approach to processor design and described a new and better road map for datacenter managers. Extending the life of datacenters without investment in costly refurbishments to accommodate thermal and power demands is yet another cost savings that will be reflected in improved business value.

The Sun Fire T1000 and T2000 servers run the Solaris operating system and are binary compatible with previous Sun processors. As a result, ISVs and existing customers can move software to these new platforms without recompilation.

IDC encourages IT organizations deploying high-density computing solutions for the Web tier to evaluate these new servers based on the SPARC T1 processor. The new servers from Sun should be considered for use outside the normal workload boundaries commonly associated with this form factor.

THE IT INFRASTRUCTURE LANDSCAPE

Demand for server resources continues to grow as organizations move to provide additional capacity for traditional workloads and as they expand the footprint of IT with new workloads, such as Web services, voice over IP, and a multitude of new mobility applications. Additionally, increases in available bandwidth allow customers and staff to call for more information when accessing enterprise Web services, which further drives the need to process even greater amounts of information. Internet-worked computing continues to drive the need to deploy more infrastructure to keep pace with global markets. CIOs and IT directors are considering their options when scaling enterprise infrastructure to meet these demands. At the same time, companies are running up against the practical limits in their datacenter environmentals.

In investigating the potential impact on datacenter costs, IDC research found that power and cooling requirements are rising rapidly when customers shift to emphasize more scale-out configurations using blade servers or clusters of rackmount servers. While datacenter managers are able to meet cooling requirements today, and while average Watts per square foot (W/sqft) are in reasonable ranges, the situation is worsening and a solution is needed. One datacenter manager told IDC that he is exploring water-cooled racks, which will be needed when cold air cannot dissipate the heat fast enough. Meanwhile, another datacenter manager told IDC that rising cost to the customer is inhibiting their choice of compute-dense higher-performance servers. Although expensive, datacenter retrofitting is also underway as datacenter managers increase their investment in both power and cooling systems. Generally, upgrades double the existing Watts per square foot; upgrading from an average of 100W/sqft to 200W/sqft is typical.

Thus, the ability to power and cool the equipment as well as the costs associated with power consumption and heat dissipation are becoming critical factors in a company's IT evolution. Increases in power consumption and the concomitant creation of heat cannot continue or they will become limiting factors for compute density in the datacenter. As a result IT organizations will need to deploy more servers that are better designed to support modern-day workloads with the latest technology. Moreover, as IT organizations continue to consolidate workloads on fewer servers, those servers will need to be highly reliable. The interrelationship among these requirements for more powerful, smaller, and more energy-efficient servers precludes ordinary improvements in yesterday's design. A quantum change in server design is needed soon.

If the trends for power consumption and heat dissipation do not change, then enterprises will face difficult trade-offs. While deploying additional Web services could enlarge the organization's addressable market and bring in new revenue, the cost of fielding such systems may become prohibitive. IT, once the enabler for business expansion and business agility, could become an inhibitor to change.

Power Consumption and Heat Dissipation

IT organizations know that they must support more demanding workloads with more powerful server systems and that, at the same time, datacenters are expensive to build and to operate. According to IDC forecasts, U.S. datacenters alone will house 50% more servers by 2009. IDC estimates that the cost of electrical power for the 14 million servers in use in 2009 will be in excess of $5 billion per year, and this estimate does not take into account the cost of cooling. Reductions in the size, electrical power, and cooling requirements for servers can have a significant impact on both capital expenses for new datacenters and operating expenses for existing facilities.

With each passing year, servers are physically smaller in size (i.e., compute density is rising); however, power consumption and cooling requirements continue to grow. The major culprits are the multiple processors within the servers, which require more electrical power and produce more heat than ever before. Increasing power requirements drive the need for larger power supplies and more in-server cooling.

Today, racks of servers can require 8kW to 12kW of electrical power, which increases the construction costs as well as operating costs of the datacenter. This requirement will only increase as new processors make it into the generation of servers such as server blades. Customers also need to consider the amount of power that their utility provider is able to deliver to the datacenter. There are instances, especially in densely populated business districts, when no more than 1kW can be delivered to a rack.

With these factors in mind, IT organizations are increasingly concerned with the amount of computing throughput they can deliver in a cost-effective power and thermal envelope. Performance per watt is emerging as a key metric when choosing the right servers to meet the workload requirements for the enterprise as well as fit within the datacenter conditions.

Shifting Requirements for Modern-Day Commercial Workloads

Today's typical IT workloads can be categorized in many ways, but they all share at least one element in common: these workloads increasingly depend on multiple concurrent processes or *threads*. A thread is an independent software process consisting of a sequence of instructions to be executed by a processor. Threads occur as multiple independent applications (i.e., two separate user jobs) as well as within a single application (i.e., parallelism inherent to distinct tasks).

- Web workloads illustrate the need for multiple independent applications. As users browse and click, they call on the server to execute a process on their behalf. Thousands of users each need a small amount of the server's resources.

- Database workloads illustrate multiple threads within a single application. Launching parallel processes (i.e., multiple threads) across multiple processors speeds up tasks such as sorting and searching.

Instruction-Level Parallelism Versus Thread-Level Parallelism

Traditionally, servers have been designed to support instruction-level parallelism. Programmers identified logical and mathematical junctions where the software could branch, data could be partitioned, and several routines could be processed in parallel. Thread-level parallelism, in contrast, is finer grained and built into the software development and execution environment. Java technology, for example, was designed to provide a high-quality multithreaded workload execution environment.

While instruction-level parallelism has been an effective technique for improving performance in the past, thread-level parallelism better matches the needs of commercial workloads today (see Table 1). With the exception of data warehouse applications, where instruction-level parallelism has been highly successful in enabling better performance, instruction-level parallelism has not yielded impressive results. In contrast, thread-level parallelism provides a high degree of success across the board. Table 1 also identifies the size of typical applications (i.e., their instructions and data) and the degree to which data is typically shared among users and workloads.

Improved Utilization and Consolidation

Another driver for change is the demand for better IT utilization. Since the late 1990s, customers have embarked on a one server–one application paradigm that, while providing a solid security model and application isolation, has left the equipment sorely underutilitzed. It is commonly believed that today's server utilization rates are a mere 10–20%. As enterprises increasingly demand maximum business value from their IT systems, both in terms of acquisition costs and operating expenses, they are looking for a means to get the most out of the equipment in the datacenter. At the same time, both the density of computing resources and their reliability are important as datacenter managers seek to consolidate server resources, simplify operations, and reduce cost.

TABLE 1

Attributes of Commercial Workloads

	Three-Tier Architecture			Client/Server		Data Warehouse
	Web	Application Server (JBB)	Data	SAP 2T	SAP 3T (db)	Decision Support
Instruction-level parallelism	Low	Low	Low	Medium	Low	High
Thread-level parallelism	High	High	High	High	High	High
Instruction/data working set	Large	Large	Large	Medium	Large	Large
Data sharing	Low	Medium	High	Medium	High	Medium

Source: Sun Microsystems, 2005

Summary

The IT landscape is shifting as organizations contemplate how to best support growing commercial workloads in a cost-effective manner. Expenses once thought to be negligible, such as the cost of electrical power, are looming larger as server density and power consumption continue to increase. Meanwhile, the workloads themselves are placing different demands on servers today. In particular, these workloads are thread-rich. Because commercial workloads provide a high degree of thread-level parallelism, a new breed of servers must recognize and process more granular units of work efficiently. In doing so, IT organizations will be better able to improve utilization and consolidate processing.

PARADIGM SHIFTS IN SCIENCE AND TECHNOLOGY

When analyzing progress in science, philosopher Thomas Kuhn concluded that the growth of knowledge was not well described as a smooth rising function. Rather, a set of ideas would emerge suddenly as a remedy to flaws in the old way of thinking. Kuhn referred to the set of ideas as a *paradigm* and described scientific breakthroughs as *paradigm shifts*. Isaac Newton's calculus and Albert Einstein's theory of relativity are examples of quantum leaps in understanding that enable a host of new problems to be solved. While Kuhn's ideas were rooted in the theoretical sciences, they provide a useful framework for analyzing innovation in information technology as well.

While the microprocessor bears some resemblance to processors built with discrete components, for example, its explosive impact on the IT industry is best described as a paradigm shift. Accelerated by improvements in semiconductor technology and fueled by mass-market economics, the microprocessor triggered the creation of new products and markets. Another paradigm shift occurred when the U.S. Department of Defense research on enabling different network protocols to interoperate changed our thinking of networks and set the stage for the Internet and the World Wide Web.

Kuhn's model raises interesting questions about trends in server and processor technology. Stubborn problems such as increasing power and heat suggest that further development in the same paradigm may not succeed. Moreover, changes in workload characteristics suggest that a different paradigm for parallelism is required. With these ideas in mind, we shall examine Sun's progress with its new processor architecture and new server products.

Sun's Chip Multithreading Processor Architecture

Sun has developed a technology strategy for addressing future server throughput computing requirements. A key component of this strategy is the use of new approaches to processor architectures. The major technology that Sun is implementing at the processor level is chip multithreading (CMT). CMT reduces the "wait time" that is associated with nonproductive processor clock cycles when processors await data from memory. By allowing the processor to support multiple threads at once, the throughput of the entire system is improved. A more detailed explanation of CMT can be found in Appendix A.

Because thread-rich workloads are not created equal, Sun divides the network computing space into network-facing and data-facing workloads.

- Network-facing workloads typically require processing multiple, small applications and databases to support the delivery of Web and transaction services. These workloads are characterized by asynchronous, stateless threads that are I/O driven.

- Data-facing workloads require processing large, multiterabyte multiuser databases. These workloads are characterized by stateful, often synchronized threads and are memory driven or compute driven.

The UltraSPARC T1 Processor

In December of 2005, Sun announced the release of the new UltraSPARC T1 processor, code-named Niagara, and the availability of two servers that utilize this processor — the Sun Fire T1000 and Sun Fire T2000 CoolThreads servers. The UltraSPARC T1 processor is a second-generation CMT processor, and it will be available with 4, 6, and 8 cores. Sun refers to the UltraSPARC T1 as its first "radical CMT processor." As Figure 1 shows, Sun expects a 15-fold improvement in application throughput for servers using the UltraSPARC T1 processor, when compared to servers using the UltraSPARC IIIi processor. The UltraSPARC T1 processor is designed specifically for network-facing workloads because its design emphasizes the ability to handle many threads in parallel at the expense of single-thread performance.

The UltraSPARC T1 processor contains up to eight cores, and each core is capable of executing four threads. The processor can handle 32 threads simultaneously, a 16-fold increase over the previous generation processor. Cores in the UltraSPARC T1 processor are interconnected with a 134GB/sec crossbar switch and there is a 12-way associative 3MB Level 2 cache on the chip. Four DDR2 channels delivering 24Gb/sec of memory bandwidth are available to interconnect with RAM. The UltraSPARC T1 processor comprises approximately 300 million transistors. According to Sun, the UltraSPARC T1 processor typically consumes 73W of electrical power.

FIGURE 1

Sun's CMT Processor Road Map

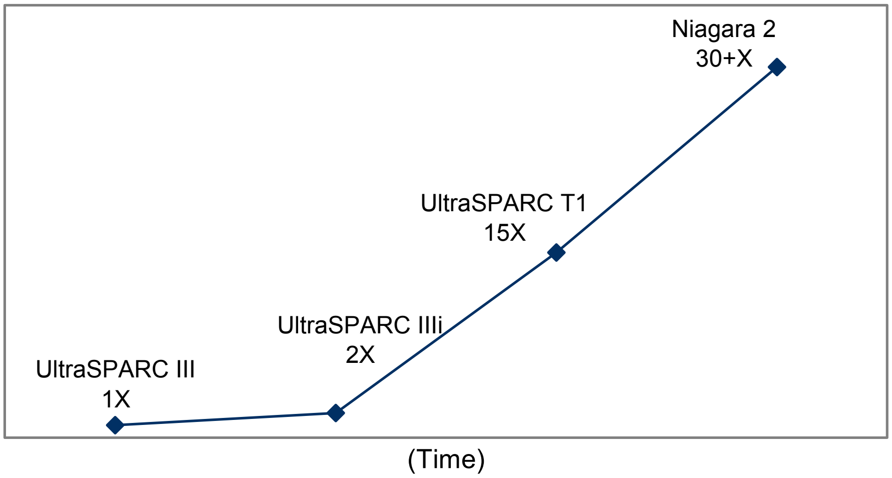

Source: Sun Microsystems, 2005

The UltraSPARC T1 processor is one of two Sun CMT processor lines. The UltraSPARC T1 processor is designed for the network-facing workloads that are typical in Web and transaction services. A second processor, code-named Rock, is designed for CPU-centric workloads, such as database processing. UltraSPARC T1 and Rock will provide new upgrade paths for existing Sun customers and will open up new avenues of use by Sun's new and prospective customers due to binary compatibility with previous generation processors. The data-facing version of the UltraSPARC T1 processor line is scheduled to be released in 2008.

Servers Based on the UltraSPARC T1 Processor

Sun has announced two servers that utilize the UltraSPARC T1 processor. Both servers use a single UltraSPARC T1 processor with 4, 6, or 8 cores and are engineered with RAS capabilities, including chipkill and memory sparing. Both servers have 4Gb Ethernet ports and are designed for lights-out management.

- The Sun Fire T1000 CoolThreads server is packaged in a 1 RU chassis that is 19in. deep. Typical power consumption is 200W.
  - This server can contain up to 16GB of DDR2 memory, a single PCI-E slot, and may contain a single SAS/SATA disk drive or be diskless.
  - Sun calls this server its "network infrastructure workhorse" and suggests that it be put to work as a Web server, an identity and directory server, a portal server, or an application server. The Sun Fire T1000 server could also be a node in a commercial grid.
- The Sun Fire T2000 CoolThreads server is packaged in a 2 RU chassis that is 24.4in. deep. Typical power consumption is 275W.

- This server can contain up to 32GB of DDR-2 memory. There are redundant power supplies and fans that can be hot swapped. Onboard storage options include SAS/SATA 2.5in. disk drives and a DVD drive. In addition to Ethernet ports, the server has 2 USB ports, 3 PCI-E slots, and one or two PCI-X slots for expandability.
- Sun calls the Sun Fire T2000 server its "application engine" and suggests that it is a top candidate for deployment as a Java application server, an enterprise application server for ERP and CRM, and a Web server for Web-tier server consolidation.

Space, Watts, and Performance: SWaP

Sun suggests that enterprises evaluate server choices for Web and application workloads based on three considerations: space, power consumption, and throughput performance (SWaP). SWaP is a composite of system performance, system space, and power consumption. Performance per square foot and power per square foot are the traditional metrics, while SWaP brings all three critical factors together into a single metric.

The SWaP metric is defined as performance/(space * power). For SWaP, performance is measured by an appropriate benchmark, space is measured in rack units (RUs), and power is measured in Watts. The higher the SWaP metric's value, the better.

Migrating to the Sun Fire T1000 and T2000 Servers

The transition to the next-generation Sun Fire T1000 and T2000 servers will be streamlined for existing Sun customers because the new processor is binary compatible with legacy UltraSPARC processors at the instruction set level. These servers run the Sun Solaris operating system. Thus, no recompilation of software will be needed. Customers will have uninterrupted access to Solaris operating system capabilities such as RAS functionality and thread support. The transition will also be streamlined for ISVs and developers who produce the packaged software upon which customers depend.

IDC expects Sun to gain new customers who will migrate to these new platforms. The Sun Fire T1000 and T2000 will appeal to IT organizations seeking servers with pwerformance, power, density, and I/O capabilities that are greater than the normal workload boundaries for existing servers with similar form factors. These servers will be evaluated by customers who would not normally associate this technology with Sun.

Managers of datacenters that are nearing their power and thermal limits are also likely to evaluate Sun's new servers favorably. Migration to the Sun Fire T1000 and T2000 may be a viable alternative to refurbishing the datacenter power and cooling infrastructure.

According to Sun, the best candidate workloads for the Sun Fire T1000 and T2000 servers will be those that are highly threaded and are supported by up to 8-way or 12-way servers today. Consolidating workloads running on 1- or 2-processor servers with Solaris containers is also a specific design point for the platform. Web services, mail, and most business applications are workloads that these servers will support. Applications that depend substantially on single-thread performance, such as batch workloads that take hours to run, are not as ideally suited for migration to a multithreaded approach. Also, because each multicore UltraSPARC T1 processor contains a single floating-point processor, the Sun Fire T1000 and T2000 servers are not good platforms for floating point–intensive applications.

Launching the Sun Fire T1000 and T2000 Servers

Sun is taking a leading role in the move to multicore technology. The company has gone beyond just simply trying to add capability to the processor and servers and has instead focused on designing servers that leverage the UltraSPARC T1 processor. The design of these systems has four key objectives in mind:

- **High throughput.** Achieved primarily by Sun's multicore, multithreaded processor.
- **Optimization for commercial applications.** Workloads with high levels of integer code and minimal floating-point operations. Such applications are typical network infrastructure workloads: Web, mail, and business applications.
- **Better space, power, and cooling requirements.** Largely driven by the multicore design of the UltraSPARC T1 processor with up to 8 simple, low-power cores.
- **Continuity with previous products.** Achieved by providing both instruction-level binary compatibility and supporting the Solaris operating system.

Beyond the technical aspects of the systems, Sun has taken several important steps to ensure a smooth and rapid acceptance of its new throughput servers First, the company has been in close contact with its customers and partners to ensure that these stakeholders are prepared to leverage features such as increased compute density and reduced power consumption.

Second, Sun has deployed T1000 and T2000 servers in its own Sun Solution Centers around the globe. These servers have been put to use by customers and ISVs.

Finally, Sun moved prototype versions of these servers into the hands of ISVs and select customers early in 2005 to make sure that beta testing was thorough and rigorous. Hundreds of T1000 and T2000 servers have been shipped to customers and ISVs. The early access program has confirmed for Sun that these servers provide binary compatibility with legacy software and that they perform within the expected power and thermal parameters.

IDC ANALYSIS

IDC believes that Sun's shift to a multicore processor design is a paradigm shift that addresses serious problems emerging for single-core designs. The old paradigm, which provided two decades of performance improvements, held that the best design was a powerful single-core processor and all of the available transistors were budgeted to this goal. In the new paradigm, simpler cores that manage multiple threads not only reduce power and heat, but also improve overall throughput by addressing the memory latency problem directly. Transistors budgeted for large caches, out-of-order execution, and prefetching have been put to other uses because these functions are unnecessary in a multicore multithread processor architecture.

The paradigm shift to multicore processors is industrywide. Sun and its competitors have embraced the same basic concept. IDC expects that variances in performance and throughput will emerge as the different vendors create different kinds of multi-core processors. That is, within the multicore paradigm, there will be opportunities and challenges that vendors will approach in different ways. Customers should be pleased with the shift in attention to issues that directly improve the business value of their IT investments. In particular, growing energy use has been an increasingly expensive and difficult problem for datacenter managers. The shift to servers using multicore processors will help to reduce operating costs while providing higher performance servers — benefits that should improve the IT organization's ability to support growth in commercial workloads within the power, space, and cooling specifications of current datacenters.

Opportunities

At a time when competitors are offering dual-core processors, Sun is first to market with an eight-core processor. The successful deployment of CMT processors in Sun servers will provide Sun with a competitive edge, both with its current customers expecting improved price/performance on a regular basis.

Sun's new offerings will also be attractive to new customers who need to support Web-based workloads with servers that require a smaller footprint and reduced power consumption. IDC believes that the price/performance advantage may well compensate for the cost of migration for IT organizations who currently use servers based on competing technologies.

Sun's continuing ability to deliver new technology that is compatible with legacy systems is an attractive value proposition for its existing customers. Migration can be a serious issue when technology changes abruptly and as continuity continues to be a differentiator for Sun and a key benefit to Sun's customers.

Sun's position as both a processor designer and server manufacturer provides the potential to tightly integrate CMT processing into Sun's server line. Moreover, Sun's long-term commitment to network computing means that the company supports a broad offering of products that includes systems software and software tools as well as server hardware. Sun brings a systems approach to IT.

Challenges

Sun's CMT will compete in the marketplace with other vendors' approaches to multicore and multithreaded processors. Examples of vendors with commitments to multicore processors include AMD, IBM, and Intel. Others may join over time.

Product pricing depends, in part, on the number of units that can be shipped. Sun will need to ramp up quickly to generate the volume that is needed to reduce cost. Ramping up the fabrication of new processors, while keeping yields in line, has been difficult for many semiconductor suppliers.

The delivery of this new technology is clearly a key initiative in the remaking of the company. While the Sun sales force still grapples with selling volume systems, the technologies represented by the new UltraSPARC T1 and the Sun Fire T1000 and T2000 servers certainly help bolster Sun's reputation as a technology and systems leader.

Finally, Sun must differentiate its product offerings and clearly communicate the value proposition to gain new customers. An important element of the CMT product strategy will be Sun's success when encouraging the ecosystem of ISVs and SIs to embrace this new line of servers.

CONCLUSION

IDC recommends that datacenter managers who are evaluating dense computing solutions should be sure to include Sun Fire T1000 and T2000 servers in their research. Moreover, IT strategists responsible for estimating the return on IT investments should pencil out the implications of using a server platform that provides higher performance in a smaller footprint, with less power consumption, and produces less heat.

APPENDIX

An Introduction to CMT

Moore's Law continues to drive significant improvements in computer system component technologies, but not at the same pace for all components. As a result, computing system architectures continue to flex as they reflect new balance points among key technologies: processing, memory, storage, display, power, network, and software.

The issue of shifting balance points is increasingly important for processors and memory both in workstations and, especially, in servers. As described by Moore's Law, the density of semiconductor devices continues to increase. For processors, smaller transistor geometries have led to higher switching speeds, increased transistor budgets, and the shift to 64-bit processing. For memory, Moore's Law has improved capacity. However, the relationship between processor speed and memory latency (i.e., the time it takes for data from memory to reach the processor) has not followed the same growth path and is out of balance. Namely, processors increasingly outrun the ability of memory to supply the processor with data.

The consequence of the imbalance between processor and memory is illustrated in Figure A1. The time-series plot of processing and memory access shows that Processor A divides its time between processing steps and pipeline stalls while data is written and retrieved from memory. Processor B is a faster processor, and memory latency remains the same. Figure A1 illustrates two important lessons:

- Increases in processor speed do not appreciably improve throughput when memory latency remains fixed and large compared with processing time.

- Fast processors outrun memory and remain in a stalled state for much of the time.

FIGURE A1

Processor Speed and Memory Latency

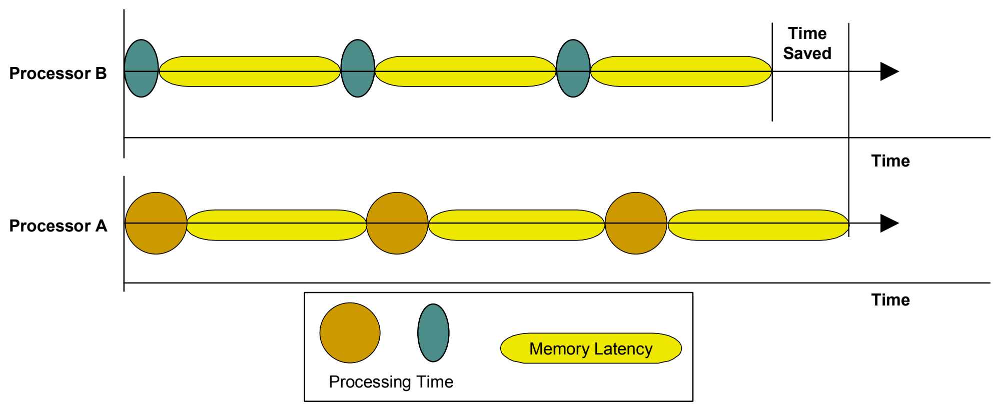

Source: IDC, 2005

Freeing Up Processor Real Estate

When processor designers relax the requirement for the fastest possible single-thread performance, a significant amount of a processor's silicon "real estate" can be opened up for other uses. Based on early laboratory results, Sun told IDC, "A core [i.e., the basic hardware compute engine] providing half the single-thread performance of today's fastest and most complicated processor cores can be built using about 10% of the real estate." Doing so frees up 90% of the silicon for other uses, such as multithread support and integrated network and security functionality, depending on the workload.

Efficient Use of Clock Cycles

CMT processors can use some of the available real estate to implement mechanisms that can reduce the amount of clock cycles wasted in wait states. Specifically, CMT processors that manage two or more threads concurrently can quickly switch between threads when the active thread stalls. The goal of multithreading strategies is "latency hiding." Instead of attempting to reduce memory latency times, multithreading simply allows the processor to jump to a ready-to-run thread as soon as the current or active thread stalls for lack of data. The technical challenge is to provide a *zero cycle context switch*, a processor that can move among threads with the absolute minimum of overhead. Figure A2 illustrates how multithreading works:

- A processor simultaneously holds states for multiple threads with the ability to switch among threads within a single clock cycle.

- When a currently executing thread stalls on a memory reference, the processor switches to the next executable thread and begins processing it. There is no danger of switching to a thread that is stalled because only executable threads are inserted into the queue of threads awaiting processing.

- When the second thread stalls on a memory reference, the processor switches to a third executable thread and then to a fourth executable thread when the third thread stalls.

- By the time the fourth thread stalls, the first thread has likely completed its memory access and is ready to resume executing.

FIGURE A2

Multithreading to Increase Throughput

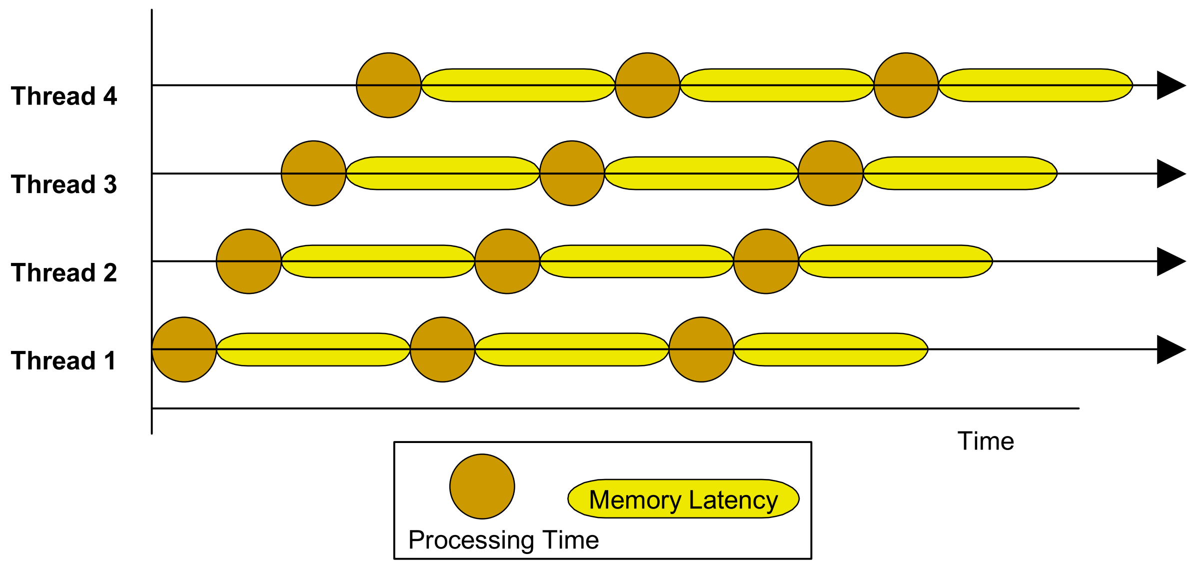

Source: IDC, 2005

While at any given time several threads on the processor will be in a wait state, one or more other threads will likely have the data needed from memory and execute. Ideally, the processor is always doing useful work by switching between threads whenever a running thread is stalled by a memory reference. Thus, although the performance of any given thread is not changed, the total number of threads completed in a given time period is increased.

Multithreaded processor performance is bounded by such factors as:

- The number of threads that a processor can physically manage

- The availability of threads (i.e., how well the applications and systems software are able to identify and manage threads)

- Memory performance (i.e., multithreading acts to hide the effects of memory latency but in doing so places greater demand on processor-to-memory bandwidth)

- Processor speed (i.e., multithreading allows the processor to use as much of its available resources as possible; however, the performance of those resources is ultimately bounded by processor speed)

The multithreaded approach addresses the imbalance between processing speed and memory latency. Together, a server with processors supporting multithreading and an operating environment capable of managing threads could provide significantly greater system throughput. The potential effect of multiple execution of multiple threads is shown in Figure A3. To improve throughput, system designers must meet the following criteria:

- Processor designers will need to allocate semiconductor resources (i.e., the transistor budget) to include support for the queuing and execution of a multitude of threads.

- Support for threads will be needed in the operating system, which will need to be multithreaded and tightly coupled to the processor's multithreading functionality.

- Support for threads will also be needed in the software infrastructure used to develop the applications that carry enterprise workloads (i.e., support for multiple threads must reach from the processor hardware to the top of the application stack and deliver performance improvements to the enterprise).

The Numbers Get Interesting

One important aspect of Sun's approach to throughput architectures is that it gives processor designers three factors to work with when designing future multithreaded processor configurations:

- The number of cores per processor (m)
- The number of threads managed by each core (n)
- The single-thread performance of each core (p)

These three factors combine multiplicatively to yield throughput performance or overall processor capacity (i.e., capacity increase = m*n*p). As transistor budgets increase over time, Sun claims that its processor architects can increase all three factors, thus significantly increasing performance. For example, a processor with eight cores, four threads per core, and each core running at twice the performance of current SPARC chips could provide 64 times the throughput capacity of current processor (i.e., m*n*p=8*4*2=64).

FIGURE A3

CMT Processor

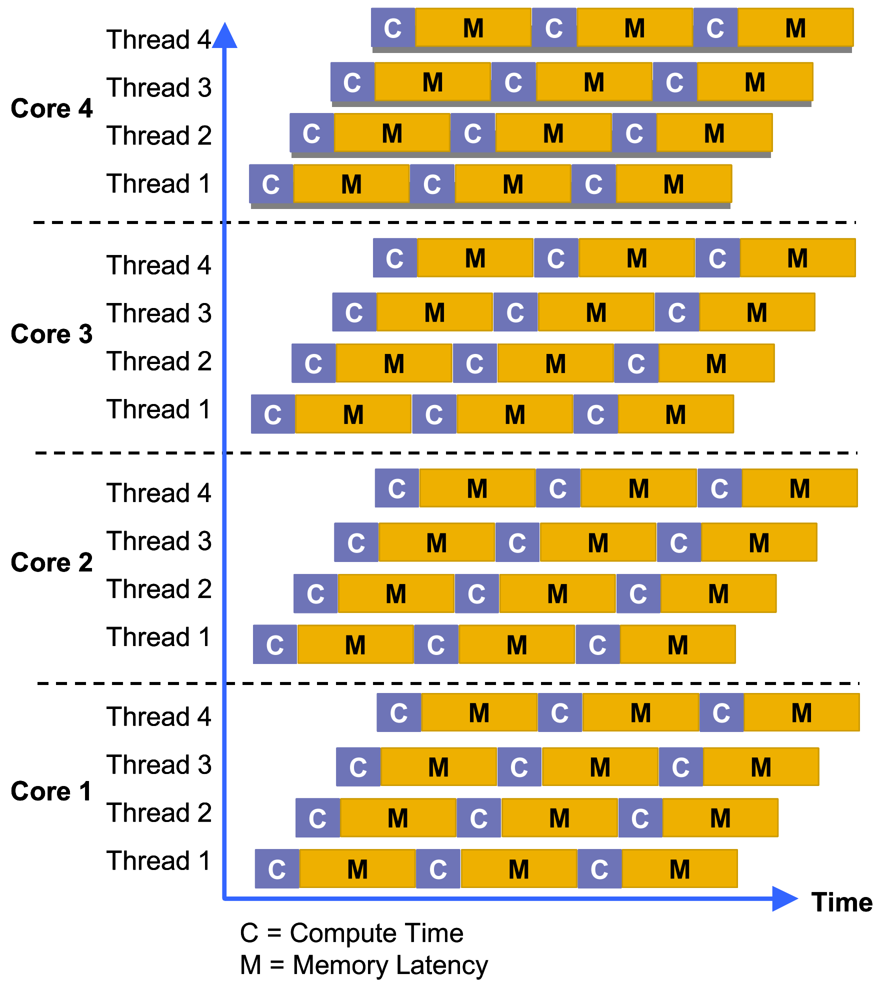

Source: Sun Microsystems, 2005