November 23, 2005

Sun Announces Niagara Processor

Initial Look Shows Convincing Performance And Outstanding Power Efficiency

by Richard Fichera
with Laura Koetzle and Sarah Bernhardt

EXECUTIVE SUMMARY

On November 14, Sun Microsystems announced the arrival of the long-awaited Niagara CPU called UltraSPARC T1, which is the first member of Sun's new multicore highly threaded SPARC processors. Niagara offers a binary-compatible upgrade path for current Sun SPARC users and industry-leading performance and performance per watt compared to alternatives. With this product, Sun hits a sweet spot for its current Solaris customers, and it may also have at least a temporary wedge to expand its Solaris franchise in selected markets. In addition, the fact that the announcement comes slightly earlier than expected bodes well for Sun's ability to execute the rest of the chip multithreaded (CMT) road map.

SUN BETS ON A NEW ARCHITECTURE AS IT SPREADS THE RISK ON THE OLD

While struggling with an ever-slipping CPU development road map for its line of UltraSPARC processors in 2001-2002, Sun embarked on a major technology gamble with its CMT road map, which promised a new and aggressive CPU architecture.[1] Niagara — the first generation of CMT products — and its planned successors represent an abrupt departure from escalating clock speeds and ever-deeper pipelines in favor of more, smaller, and multithreaded cores.[2] Sun even released Niagara early, having originally committed to the first half of 2006.

Marching To A Different Drummer: UltraSPARC T1 Architecture

So why has Sun changed its architecture with CMT? Because:

- **Customers generate more thread-rich workloads.** Customer workloads were becoming increasingly thread-rich across a wide range of applications, many of which (such as Java and Web-facing processing) were expected to grow at above-market rates. Workloads that are thread-rich can benefit from both multiple cores and multiple threads.

- **Implementations have begun to explore the limitations of current designs.** Contemporary deep pipeline CPUs like the UltraSPARC IV, POWER5, x86, and Itanium CPUs were spending an increasing portion of their time waiting on memory. This lent support to the hypothesis that a design that could manage multiple thread states would be an advantage, despite the modest improvements from early implementation of hardware threading. In building such an architecture, Sun was committing to solve several design problems that had plagued earlier threaded

architectures, including the need for extreme memory bandwidth and the requirement that the cores be able to switch thread states with very little overhead, ideally in a single clock cycle.

- **Integration of the technology stack will matter more.** Solaris itself is a highly threaded OS. Sun believed that because it could simultaneously control the hardware architecture and the software stack, it would be able to optimize the combination better than its competitors.

The architecture of the resultant product, the UltraSPARC T1, is fundamentally different from any other general-purpose chip on the market, despite its SPARC binary compatibility. Rather than one or two complex cores running at speeds of two GHz or higher, it has a total of eight cores running at up to 1.2 GHz, each with four threads per core.[3] With sufficiently rapid switching between thread states (on the order of one clock cycle), the design goal of the T1 is to keep the cores busy executing threads rather than waiting for memory. This is important because waiting for memory contributes to considerable overhead and inefficiency in conventional architectures. One prime motivator for the new architecture was the fact that gains from increasing pipeline depth were becoming increasingly expensive in terms of chip real estate — at the same time that the designs were becoming increasingly expensive and power-hungry. A complex of four memory controllers with an aggregate bandwidth of approximately 25.6 GB — higher than the memory bandwidth of any production processor — feeds the UltraSPARC T1's eight cores.

IS NIAGARA PERFORMANCE WORTH THE WAIT?

While Sun has clearly indicated that the T1 is targeted at thread-rich environments, and the architecture seems to address the major potential bottlenecks, the proof is in the pudding. Unfortunately, Sun has failed to serve up any system details and performance pudding with the announcement. However, Forrester has enough information to make some early performance predictions for the T1 processor. Multithreaded benchmarks that stress memory latency and are not cache-size sensitive have the potential to raise the T1 advantage, while benchmarks with single or small thread count, benchmarks with heavy floating point calculations (the cores on the T1 share a single floating point unit), and benchmarks that are optimized for the larger cache sizes of the other CPUs could lower the TI advantage. Specifically, we predict:

- **Worst case.** Assuming equal wait time for memory, a single T1 chip will be substantially slower — probably by a factor of between two and three — than current high-end Xeon or Intel processors on a single-thread benchmark.

- **Best case.** On a multithreaded benchmark, the throughput of the T1 will probably rise to approximately twice that of a dual-core AMD or Xeon processor. This is especially true because of the fact that each core can have up to four ready-to-run threads. This may allow the architecture of the T1 to reduce the memory wait time percentage to approximately 50% of that of a conventional architecture, possibly even lower.

If our assumptions are correct, the UltraSPARC T1 could represent the kind of step up in performance that will reinvigorate Sun's installed SPARC base, especially for Web and Java applications. Our expectation is that Sun will price the T1-based products very close to the price of a high-end single socket x86 server, yielding very aggressive price performance.

NIAGARA POWER CONSUMPTION IS EXTREMELY ATTRACTIVE FOR DENSE DATA CENTERS

When the original Afara design team began to think about a new architecture, there was concern that Internet data centers would be constrained by space and power. From 2001 through 2003, that concern faded as data center growth slowed with the general economy. But by early 2004, heat and power were again surfacing as issues, because each successive generation of x86 processors has gotten hotter as their clock speeds have increased.[4] At the time of the Niagara announcement, with energy prices at levels that would have been a bad joke three years ago, Sun's T1 boasts the lowest power-per-workload metrics of any general-purpose processor.[5] Here's the scoop:

- **Sun's UltraSPARC T1 is the greenest chip around.** The T1 represents a radical departure from other architectures in its power consumption, rated at a typical power dissipation of 72 W for the entire chip, compared with up to 130 W for a high-end x86 CPU. If our estimates of aggregate performance are approximately correct, then on an aggregate throughput-to-power comparison the T1 may be as much as three to five times better than contemporary x86 architecture processors. On a system basis, the comparison is harder to calculate because it involves assumptions about the rest of the system components. But assuming that it takes at least a large dual-socket x86 server to match the throughput of a hypothetical T1 server, the throughput-to-power advantage should still be approximately two to one, possibly higher.[6]

- **For large, dense data centers, the energy efficiency of the T1 has a double benefit.** The T1 reduces actual power costs for the data center, and it defers or eliminates big capital investments in data center power and cooling upgrades. The latter is particularly significant, because changing the power and cooling infrastructure of a server data center with several hundred servers can easily cost more than $500,000.

THE T1 IS THE BEGINNING OF A NEW FAMILY

The T1 is only the first in a series of highly threaded CPUs that Sun intends to deliver. Although details are sketchy, over the last 18 months members of Sun's executive team have mentioned several additional family members:

- **Niagara II.** We assume that the Niagara II — an improved version of the T1 (Niagara I) — will arrive in 2007 with the shift to a 65 nm process. We can expect some combination of more integration of network and I/O, possibly more cores and threads, and moderately increased

cache size. We expect that Sun will continue with integrated memory controllers, probably with significantly higher bandwidth than those of the T1.

- **Rock.** This high-end processor will probably be the successor to both the UltraSPARC and SPARC64 products. Rock will have more cores and threads than the T1, and each core will be a much higher-performance core than the contemporaneous T1 processor. On a per-core basis, our expectation is that Rock will still lag behind both IBM and Intel but will have more cores than competing systems. Rock will probably not appear until at least 2008 or 2009, initially in a 65 nm process.

- **Pebble and Boulder.** Sun Chairman and CEO Scott McNealy publicly mentioned "Pebble" — versus "Boulder" in the Rock family — implying that there would be both smaller and larger versions of the Rock chip. No other details were mentioned.

RECOMMENDATIONS

ULTRASPARC T1 IS A "MUST EVALUATE" FOR PROCESS-RICH WORKLOADS

A number of users should absolutely evaluate T1 performance and economics before making decisions to migrate either to Sun's x86 servers or another operating system and/or system vendor. This includes any user running SPARC-based Web farms, Java application servers, or other process-rich environments, as well as users running smaller general-purpose Solaris workloads. If the actual T1 system performance and performance versus power characteristics meet or exceed our expectations, then any user with power and space constraints for Web and application server workloads should take a look as well.

WHAT IT MEANS

SUN IS POTENTIALLY DISRUPTING THE PRICE/POWER/PERFORMANCE BALANCE

Sun has set a new bar for energy efficiency and processor performance for highly-threaded workloads, which will almost certainly increase as a fraction of both total Internet and total enterprise workloads over time. Competitors that had been complacent about Sun's lack of competitiveness should view this as a wake-up call. Customers who were concerned about Sun's lack of technical leadership can view this as an early return on several years of very high R&D investments and can assume that Sun will stay aggressive and competitive.

ALTERNATIVE VIEW

SUN'S ADVANTAGE IS SUSCEPTIBLE TO FUTURE COMPETITIVE MOVES

Our performance models may be overly optimistic and competition (including Sun's partner AMD) may move more aggressively than we expect in introducing multicore CPUs with lower power and better threading support. In that case, the T1 performance advantage could evaporate to the point where it won't have enough appeal to support any significant new account captures, even if it remains interesting to Sun's installed base.

ENDNOTES

[1] A group of senior Sun engineers left Sun in 1998 and formed Afara, which designed the first version of the UltraSPARC T1 based on UltraSPARC II technology. Sun acquired Afara in 2002 and used it as the base for the T1.

[2] A deep pipeline is one with multiple stages of instruction, data decode, address resolution, and prefetch, usually including some parallel instruction execution capability and post-instruction pipelining of writes to memory. In some cases, advanced hardware branch prediction and the conditional execution of multiple parallel code paths are included. These features, while improving performance, add to the complexity of the CPU and also can force the pipeline to wait for increasingly longer periods of time when it must stop and start with a new context. In Sun's case, its high-end UltraSPARC V was both expensive and increasingly late. Sun dealt with the continual slippage of its UltraSPARC V CPU project by canceling it and entering into a joint agreement with Fujitsu to use Fujitsu's future SPARC64 CPUs in a jointly designed server product line.

[3] For the true geek squad members reading this document, a thread is a set of registers storing the execution state of the core, including the contents of the six-stage pipeline. Adding threads to a processor improves the ability to execute in environments where the processor is stalling due to memory waits. Adding cores improves the ability to execute when there are multiple "ready to run" threads. Each UltraSPARC T1 core has a small L1 cache (64 kB instruction, 16 kB data), and the cores share a global 3 MB L2 cache.

[4] One of the reasons the heat and power debate rose so quickly in the public eye, despite the fact that it had been growing slowly in real data centers, was the introduction of the AMD Opteron. A full 64-bit CPU with performance that seriously challenged Intel, it also operated at a lower clock speed and power than rival CPUs from Intel. This delivered the first tangible proof that performance increases do not have to be coupled with successively higher power and temperatures.

[5] We suspect that the Cell processor jointly designed by Sony, Toshiba, and IBM is also operating on the extreme fringes of power versus workload efficiency, as do many specialized DSP and network processing chips. However, Sun's T1 definitely takes the prize for general-purpose processors.

[6] We are assuming 72 W for the T1 vs. 130 W for the x86 CPU, 6 W/GB for memory (constant across systems), 20 W for the Northbridge controller for Intel systems, and 75 W for fans, disks, power supply loss, and other power requirements.