



Yen and the Art of Microprocessor Design (Sun Launches Its Sun Fire T1000/2000)

Nathan Brookwood

Nathan@Insight64.com

(408)257-5124

December 7, 2005

Yesterday, a few months earlier than it had projected three years ago, Sun launched its first systems based on its UltraSPARC T1 (nee Niagara). These systems, the 1U, \$2,995 T1000 and the 2U, \$7,795 T2000, contain only a single processor chip, but that chip contains eight CPU cores, each with four threads, and looks to the software like 32 discrete (single-core) CPUs. The new Sun systems handily outperform comparably configured, albeit more expensive Power 5+, Xeon and Itanium configurations on standard benchmarks (SPECweb2005, SPECjAppServer2004, and SPECjbb2005, along with others). But performance is only half the story; the T1000 and T2000 consume far less power and dissipate far less heat than those Power 5+, Xeon and Itanium systems. This gives Sun a dramatic advantage in terms of performance per Watt, a metric of increasing importance to IT managers struggling to handle growing workloads in datacenters that have already maxed out their power and HVAC resources.¹ Sun's systems beat competitive offerings by factors of four or five on this key metric, a truly stunning improvement over the prior state of the art. If these new systems cannot reignite growth of Sun's SPARC-based systems business, it's hard for Insight 64 to imagine what could.

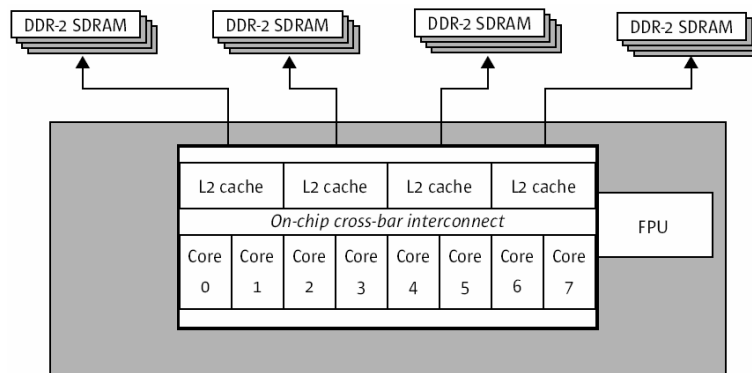
Sun's advantages all stem from the architecture of its new UltraSPARC T1 processor. Most processor architects agree that DRAM latency (the time it takes to move data from the system's main memory to its on-chip caches) is usually the principal impediment to improved system performance. It takes anywhere from 50 to 100 nanoseconds for a typical DRAM subsystem to deliver data to the CPU, and during this interval, a 2GHz processor like Intel's Xeon or AMD's Opteron loses the opportunity to execute 200 to 400 machine instructions. CPU designers employ a variety of techniques to minimize the impact of cache misses. Large on-chip caches reduce the likelihood of a miss somewhat, but consume large amounts of chip real estate and add to cost. Out-of-order execution (OOO) allows a CPU to process instructions that don't depend on missing data while stalled instructions wait for the data they need to trickle in from memory. This approach complicates CPU design and provides at best a partial solution, since the most advanced CPUs can juggle about 100 instructions in this manner, but the DRAM delays force the loss of 200 to 400 instruction execution slots. Sun's engineers identified (correctly, in our opinion) that the mismatch between CPU and DRAM speed would only worsen over time, and that an entirely different approach was needed.

David Yen, the former CPU architect who now heads Sun's SPARC systems business, decided to set out in an entirely new and radical direction². Rather than fighting DRAM latency, he designed a chip that accepts latency as a fact of life, and operates well in spite of latency. Instead of adding esoteric features to minimize the impact of cache misses, his design merely switches to a ready to run thread whenever an executing thread stalls due to a cache miss. The processor works on this second thread until it too stalls, and then switches to a third, and even a fourth thread. As long as at least one of the four threads remains unblocked, the CPU does useful work that directly contributes to the final calculations. This in turn simplifies the overall execution pipeline, and eliminates the need for fancy branch prediction and OOO hardware. The chip operates on the philosophy that "if we stall, we stall; there's always something else to do." This in turn allows Sun to shrink the size of each core and allows them to fit more cores on the chip (eight) than any other general purpose processor manufactured on a 90nm process. This in turn allows Sun to run their CPU at the lowly speed of 1.2GHz and reduces the chip's power requirements. The T1 processors announced today consume 70 Watts, a little more than half the power needed to run the fastest Xeons in Intel's line. (Sun missed its original 60W power target, but few will notice and the product still beats all its competitors in this regard.)

¹ Paul Otellini recently indicated that Intel believed performance per watt would become a key system purchasing criterion, and asserted that Intel intended to lead the industry with regard to this metric.

² Although this report cites Yen as the chip's designer, in practice a large team of Sun engineers, including a cadre added via its Afara Websystems acquisition, slaved over this design for more than two years. Yen, however, was smart enough to recognize the idea's merit, and brave enough to support it organizationally although it flew in the face of the conventional wisdom regarding CPU design.

Fast processors tend to run slowly unless they are mated with memory subsystems that can provide the bandwidth needed to feed the processor's voracious appetite for data. To satiate the UltraSPARC T1's appetite, Sun includes four DDR2 DRAM controllers directly on the chip. These controllers can move up to 20GB/second between the DRAM banks and the processor. Memory bandwidth should not constrain this processor in most applications. The on-chip location of these controllers minimizes memory access latency. (Even though the T1 embraces DRAM latency, less is always better in this regard.) The memory controllers connect to the unified 12-way associative on-chip level two cache, which in turn is connected to a crossbar switch that moves data between all eight cores and the cache. All of these pieces fit together as shown below:



Source: Sun

The diagram above highlights one of the new chip's key limitations, namely its lackluster floating point performance. All eight cores share a single floating point execution unit. When any thread in any core encounters a floating point instruction, it schedules that instruction's execution on the FPU, and stalls until the operation completes. The wrong mix of integer and floating point instructions could slow the processor down to a crawl. Sun argues that such sequences occur rarely in the applications they targeted for the T1, but Insight 64 anticipates that it won't take IBM and Intel long to find an optimal mix of integer and floating code that invokes this pathological behavior.

The diagram also highlights a second limitation – this 8 core/32 thread processor lacks the ability to be used in dual- or multi-processor arrangements. Users seeking more than 32 hardware threads in their systems will have to wait for Niagara 2, the 65nm T1 follow-on slated for 2007. That chip will feature 8 cores and 8 threads per core (64 threads in all), and a dual-processor capability that will allow up to 128 cache-coherent threads in a system.

Conclusion

We're (obviously) impressed with the approach Sun took with its T1-based systems, and with the performance and performance per watt results they published yesterday. Before Sun embarked on its throughput computing adventure, we (like many others) questioned whether it made sense for Sun to develop its own processors. Its UltraSPARC line offered few advantages with regard to performance or price/performance. Their chips' primary virtue was the ability to run Solaris without the need to recompile and/or reacquire software applications (a painful process many Macintosh users will be forced to undergo when Apple moves from PowerPC to x86 processors next year). The launch of Sun's new products demonstrates that there is still a place for innovation in the systems business, and that customers are best served when many suppliers are allowed to bring their unique perspectives to the market. It's unlikely that an Intel or an AMD would have pursued the kind of extreme multi-core/multi-thread approach Sun pursued, since those suppliers try to leverage their designs across both client and server markets; the fit between fat clients and 32-thread processors remains an unknown at this point.

Now that Sun has brought its new products to market, it remains to be seen whether prospective customers will adopt them. We're confident that Sun's installed base will find them irresistible, but will the company be able to sway customers that have grown used to purchasing x86-based systems from a multitude of system suppliers? Although Insight 64 generally views the rise of (so called) industry-standard based systems as an ineluctable force, we also accept that systems based on proprietary approaches can have merit, especially when there is no industry-standard alternative. We believe the advantages offered by Sun's new approach are sufficiently compelling that buyers should set aside their understandable bias toward systems based on industry standards and carefully evaluate these new Sun systems.