UltraSPARC® T2 Processor with Wind River Platform for Network Equipment, Linux Edition

The future for embedded systems

Today's network infrastructure must deliver far more than just packet routing. Deep packet inspection, XML processing, different qualities of service based on packet type, and more are all in demand. The challenge for embedded systems designers is to deliver the necessary functionality while coping with the triple threat of flattening processor performance, increasing heat output, and increasing memory latency. To date, most designers have resorted to a combination of general-purpose processors (GPUs), for compute-intensive tasks, and specialized networking processors (NPUs), for networking functions. Now, with the UltraSPARC T2 CMT processor, plus Wind River Platform for Network Equipment, Linux Edition, embedded systems designers can build their systems around a single, integrated processor and CGL distribution, producing systems that are simpler to design, more reliable, more powerful, and more economical to manufacture.

Network growth is changing embedded processor requirements

The skyrocketing popularity of the Internet coupled with widespread wireless adoption has resulted in a dramatic increase in the requirement for access bandwidth. Over three million subscribers per week join on Internet-capable handsets and desktops on which voice is just one of the many forms of content — which can include images, music, video, and even live television.

Traditional telephony offerings therefore have to be merged with IP-based voice, video, and data services. Packets and sessions are rapidly replacing circuits and calls. Packet processing, meanwhile, is no longer confined to core network switching. It is migrating to the edge of the network and governing access decisions in areas such as content-oriented routing, wire speed security, and metering of varying qualities of service based on the customer profile. This movement therefore comes hand-in-hand with an increasing requirement that the edge processors be able to handle the heavy computing needed for deep packet inspection.

Highlights

- Multicore/multithread processor designs are the only answer to the triple threat of flattening processor performance, increasing heat output, and increasing memory latency
- Wind River Platform for Network Equipment, Linux Edition, running on the UltraSPARC® T2 processor gives you the industry's best CGL on the industry's best chip multithreading (CMT) processor
- With eight cores, 64 threads, integrated FPUs, crypto accelerators, dual 10 GbE, and PCIe I/O, the UltraSPARC T2 processor is a true "system on a chip"
- The UltraSPARC T2 processor combines the processing power of a general-purpose processor (GPU) with the throughput of a networking processor (NPU)
- Support for Wind River Platform for Network Equipment and Workbench Development Suite on the UltraSPARC T2 processor means systems designers can continue to use the same development environment they already know
- CMT systems designers now have full access to the entire Wind River ecosystem — on CMT

Intelligent networks today demand more processor capability

These packet streams must be processed at up to 10 Gb per second, a requirement for handling such immense throughput. In the past, networks have largely been "dumb" conduits, with the sole metric of efficiency being how many bits could be ferried across them in the shortest time. With the proliferation of intelligent networks, however, much more is being asked of the embedded processors in terms of deep inspection of the data packets. Complicated routing decisions are made based on analysis of the data carried within these packets. Most specialized NPUs are not able to handle the computing needed for intelligent packet routing at the high throughput speeds at which modern-day networks transfer information, and most GPUs don't have sufficient throughput for the network.

Today's solutions must address both of these needs: efficient processing for intelligent packet inspection, as well as high throughput.

Increasing clock speed as a solution has hit a wall

To meet today's demanding business requirements, systems designers are constantly looking for ways to improve time to market and deliver greater performance. Historically, the most successful way to achieve better processor performance has been to rely on higher clock speed, but that tactic is near its end.

Moore's law is commonly cited, "processor speeds double every eighteen months." But higher frequencies lead to overwhelming design complexity. Raising clock speed is showing diminishing returns, as processor performance has flattened in spite of huge increases in frequency.

In addition, temperature rises in proportion to increasing processor frequency, which leads to high heat dissipation and wasted power. The added cooling that is required adds wasteful bulk and cost to the systems into which they are designed, and limits their application suitability.

Memory latency is the new obstacle

Today's biggest gating factor for efficient processing is memory latency. DRAM speeds have not increased at anywhere near the same rate as that for CPUs, which means that processors are often idle while waiting for data to be delivered from memory.

Embedded systems in today's networking arena require a high-performance processor that can successfully address the twin problems of latency and high temperature. The multicore/multithread UltraSPARC T2 processor provides an ideal solution.

The case for chip multithreading

Conventional single-threaded CPUs suffer from the fact that memory latency gets worse as processors get faster. Much time is spent waiting for data to be retrieved from memory and made available to the CPU to be processed. Whenever the data is not ready, there is a cache miss and the core freezes until the next successful data access occurs. This "hurry up and wait" limitation leads to poor processor utilization, typically between 15 and 25 percent.

This underutilization is even more pronounced in the case of multicore, single-threaded CPUs, since there are now many more cores (up to 64 in the case of some architectures) all waiting on memory. Adding to the problem is the fact that single-threaded designs do not have shared memory caches. This places a huge burden on the memory bandwidth, to keep the individual caches filled. A memory-starved single-threaded core is no more useful than one that is stalled on a cache miss. The net effect is increased heat generation from a core that is not doing useful work.

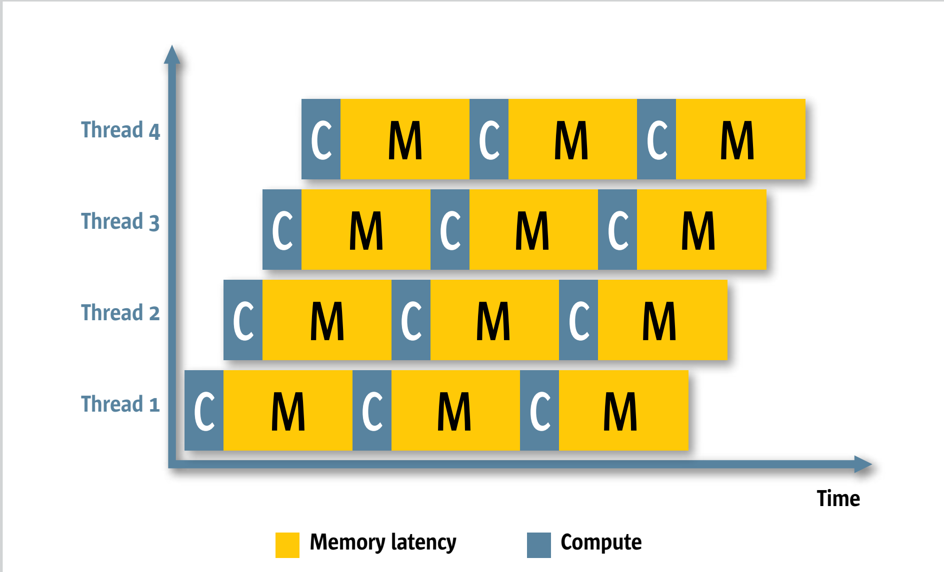

Figure 1. Chip multithreading is designed to fill the processor with data for maximum utilization.

Chip multithreading offers an elegant solution to the need for high-speed throughput combined with the computing requirements for network processing tasks such as deep packet inspection. The UltraSPARC T2 processor ships with up to eight cores and eight threads per core. As shown in Figure 1, if a particular thread stalls while waiting on memory, the processor simply launches another thread, so that the cores are almost always performing useful work, and memory latency is masked, virtually to the point of elimination.

“The UltraSPARC T2 processor delivers remarkable improvement in network performance”

Industry-leading performance through CMT

The efficient parallel threads that are the integral building blocks of CMT in the UltraSPARC T2 processor enable compute efficiency to mask memory latency. The processor’s execution pipeline remains active and is doing useful work while waiting for the stalled threads to overcome their cache misses. The ability to scale with threads rather than frequency enables overall processor utilization to be dramatically increased, up to 85 percent.

The net result is that the UltraSPARC T2 processor delivers remarkable improvement in network performance. In the industry-standard SNORT benchmark, the UltraSPARC T2 processor was shown to deliver more than three times the network throughput of a 2.2 GHz dual-core x86 chip.

The chip’s full support of CMT makes it ideal for embedded systems. The ability to execute 64 threads on a single processor leads to lower part count and cost, along with the potential for a dramatic reduction in the power required to run any system into which it is designed.

For example, in the case of a popular hardware device for email security, a single UltraSPARC T2 processor could replace the six processors that are currently used for computation, networking, and cryptography. As a result, the power drawn for the device could be reduced from 268 W down to 94 W, a reduction in energy cost of almost two-thirds.

Sun Microsystems and Wind River

A fair question that might be asked is whether gaining access to the power, simplicity, and economy of the UltraSPARC T2 processor has any cost as far as programming complexity in order to take full advantage of the processor. The answer is that there is no additional learning overhead because the industry-standard Wind River Platform for Network Equipment and its integrated Workbench Development Suite are being released on this versatile chip architecture.

The Wind River Platform for Network Equipment is based on the company's industry-leading Carrier Grade Linux platform, which enables device manufacturers to develop, run, and deploy all kinds of network devices. It is the platform of choice for many of the leading networking and communications equipment suppliers such as Nortel, Motorola, and Avaya. It is very well-suited for developing system control software in wireless infrastructure systems, fixed-mobile convergence elements, and multiservice switches.

In February 2008, Wind River became the first commercial vendor to be certified as Carrier Grade Linux 4.0 compliant, as specified by the Linux Foundation. Its software meets all the mandatory requirements outlined by this specification in seven key categories for equipment providers: availability, clustering, serviceability, performance, standards, hardware, and security.

The award-winning Workbench Development Suite is based on the open-source Eclipse development environment. It accelerates the design and development of powerful systems by supporting multiple target connections and process/task/thread debugging. It also supports 64-bit processors and includes specific tools to debug and tune multicore, multithreaded, symmetric multiprocessing systems such as those based on the UltraSPARC T2 processor. These tools provide developers with visibility into the entire platform: application code, third-party libraries, and the operating system. Additional tools for debugging, diagnostics, and testing are available.

A giant leap forward with Sun Microsystems and Wind River

The complementary strengths of Wind River's software development framework for embedded systems and the innovative, efficient architecture of the UltraSPARC T2 processor enable designers to leverage their years of experience in system design without missing a beat. You can seamlessly adopt the new processor while designing future network equipment or other demanding embedded solutions that can take advantage of CMT to deliver high throughput and high-performance computing.

Learn More

For more information about UltraSPARC T2 processors, go to: sun.com/ultrasparcT2. For a complete UltraSPARC T2 Reference Design Kit, with everything you need to start designing, go to: sun.com/products/microelectronics/support.jsp.

For more on Wind River's support for the UltraSPARC T2 processor, go to: http://www.sun.com/products/microelectronics/partners.jsp and http://www.sun.com/products/microelectronics/os.jsp.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

SunWIN #531478 4/08