Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

FREESCALE'S MULTICORE STRATEGY

Key Components: Optimized CPU Core, Accelerators, and Interconnects

By Tom R. Halfhill {8/27/07-01}

If anyone still thinks multicore chips are merely the latest technology fad, banish such impure thoughts immediately. It has become clear that chip-level multiprocessing is the only visible path toward significantly higher performance, and every leading-edge processor company has a multicore strategy. The latest company to revamp its strategy is Freescale Semiconductor.

Freescale is a good case study, because the company has been selling multicore chips—of a sort—since the mid-1990s. Freescale's popular PowerQUICC communications chips are asymmetric multiprocessors that integrate a general-purpose Power Architecture core with a special-purpose networking-acceleration engine. The first of these acceleration engines appeared in 1995 as the PowerQUICC Communications Processor Module (CPM), which was based on a proprietary RISC architecture. In 2005, Freescale superseded the CPM with the QUICC Engine, which itself contains multiple RISC cores and is backward compatible with the CPM.

One catch is that the CPM wasn't fully user programmable. Users could program it only by calling a limited number of prewritten functions through Freescale's application programming interface (API). The newer QUICC Engine adopts an open programming model. So, depending on the looseness of the definition, pre-2005 PowerQUICC chips are either heterogeneous multicore designs (Power core plus CPM) or conventional single-core designs augmented with an application-specific accelerator. Nevertheless, the nature of the PowerQUICC chip architecture—distributed processing on multiple heterogeneous processing units—forced Freescale to confront the challenges of asymmetric multiprocessing and sophisticated on-chip interconnects a long time ago. Freescale is also a pioneer in symmetric processing on multicore DSPs, having introduced the quad-core MSC8102 in 2001. (The MSC8102 and its successors are based on the StarCore DSP architecture, not the general-purpose Power Architecture.)

In October 2004, Freescale announced the MPC8641D, a Power-based host controller with two 32-bit Power e600 cores. At that time, the MPC8641D was scheduled to begin sampling in 2H05 and enter production in 1H06. (See *MPR 10/25/04-01*, "Embedded CPUs Zoom at FPF.") Unfortunately, the MPC8641D has been delayed more than a year and isn't expected to enter volume production until a revision appears in 4Q07. Meanwhile, Freescale is trying to fill demand with preproduction parts.

Freescale learned from that difficult experience when designing its next multicore processor. In June 2006, Freescale announced the PowerQUICC III MPC8572E, a highly integrated communications processor with dual Power e500 cores. This homogenous multicore design supports symmetric or asymmetric processing and is packed with hardware accelerators and I/O controllers. The MPC8572E sampled on schedule last June and is slated for volume production in 2Q08.

Given Freescale's history, it's probably no coincidence that the company is outlining a new multicore strategy at this time. Despite long experience designing high-performance, highly integrated processors, Freescale has still encountered difficulties bringing a dual-core Power chip (the MPC8641D) to market. What does this difficulty forebode for future PowerQUICC designs that must integrate four or more Power

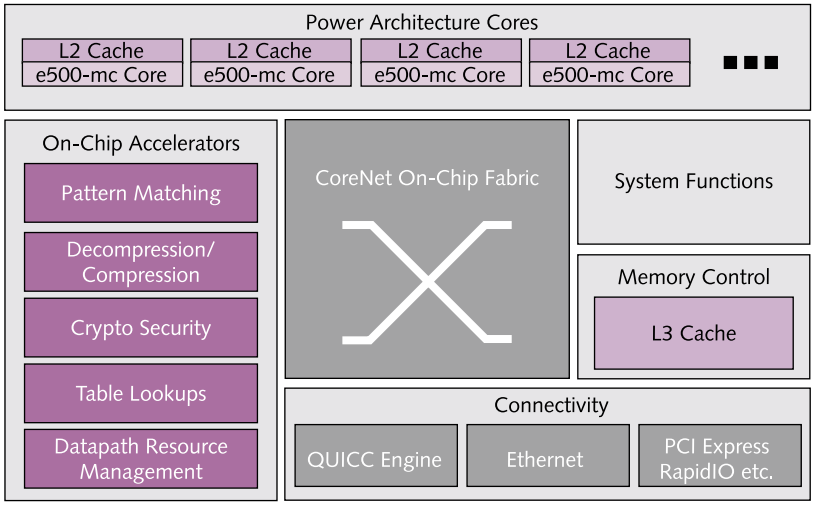

Figure 1. Freescale's multicore platform block diagram. Although the CoreNet block depicts the fabric as a common crossbar bus, it's actually more complex. Freescale describes CoreNet as a "scalable crossbar with multiple address arbiters" and says it will function more like a mesh fabric, but without the denser wiring of a true mesh with global point-to-point connections.

cores? Now is the right moment for Freescale to reassure customers by unveiling a new multicore platform. It shows that Freescale is thinking hard about future processing requirements and is developing the technology needed to implement a long-term multicore strategy.

Multicore Platform Has Multiple Components

Freescale's strategy centers on a communications-oriented technology platform encompassing all the components required for future multicore chip designs. The primary components are a 32-bit general-purpose processor core, application-specific acceleration engines, a multicore-capable hybrid simulation environment, multicore software-development tools, and a new on-chip interconnect fabric for tying the cores, I/O controllers, acceleration engines, and other resources together. The only wholly new component of this platform is CoreNet, the on-chip fabric. Other components are improved versions of existing products.

For instance, the basic processor core is an enhanced version of the existing Power e500 core. Dubbed the Power e500-mc, the enhanced core has its own L2 cache on a back-side bus, with hooks to coherently share an L3 cache with other e500-mc cores. Otherwise, the e500-mc is virtually identical to the e500 core introduced in 2001. (See *MPR 7/16/01-01*, "Speedier Book E Encore.") Private L2 caches help reduce intercore bus traffic, in contrast with a shared L2 cache. By adding a common L3 cache to this hierarchy, Freescale is trying to combine the advantages of private and pooled caches. The L3 cache will be multiple megabytes in size.

Note that the Power e500-mc is more powerful than the processor cores in some other embedded multicore designs. For example, PicoChip's basic building block is a simple 16-bit integer processor, roughly comparable to an ARM9. (See *MPR 10/14/03-03*, "PicoChip Makes a Big MAC.") Elixent—now owned by Matsushita Electronics—uses hundreds of tiny 4-bit ALUs for integer processing. (See *MPR 6/27/05-02*, "Elixent Improves D-Fabrix.") In contrast, the Power e500-mc is a two-way superscalar 32-bit processor. Freescale is aiming for clock frequencies in the 1.8–2.0GHz range for chips fabricated in a 45nm CMOS process with silicon-on-insulator (SOI) transistors. Those high clock speeds should keep the e500-mc competitive with the four-way superscalar 64-bit MIPS-compatible cores in Cavium Networks' communications processors, which currently reach a top speed of 1.0GHz.

Freescale's choice of the Power e500-mc reflects a trade-off that all multicore designers must make: either use a smaller number of more powerful cores or a larger number of less powerful cores. A major factor in Freescale's choice was backward compatibility with existing PowerQUICC chips. To maintain software compatibility, a Power core is mandatory, and there are no 4-, 8-, or 16-bit implementations of the Power Architecture. Freescale could have chosen a simpler Power core, such as the e200, but the e500 has the extra muscle needed for high-performance networking and communications. Nevertheless, Freescale's multicore platform doesn't rule out using other Power cores instead of (or in addition to) the e500-mc, including less powerful cores like the e200.

The CoreNet interconnect is capable of linking more than 32 Power e500-mc processor cores in a fully coherent on-chip network. Such a large number of cores actually goes beyond the "multicore" class and enters a loosely defined class recently dubbed "manycore"—though it still falls short of the supreme "massively parallel" class. Freescale's plans may seem overly ambitious for a company struggling to ship its first homogenous dual-core Power chip. However, many-core designs will be necessary to compete in future high-performance networking and communications applications. Freescale's strategy reflects the experience (both positive and negative) gained while developing the tardy MPC8641D.

Future Freescale chips will supplement the Power e500-mc cores with multiple hardware accelerators, such as the QUICC Engine, compression/decompression engines, pattern-matching engines, and crypto engines. MPR deems it highly unlikely that Freescale will cram as many as 32 Power cores on a single chip anytime soon, even with 45nm process technology. However, CoreNet's expansive capacity shows that Freescale is planning ahead for 32nm technology and beyond. Of course, CoreNet supports smaller designs, too. We expect Freescale to start conservatively with a homogenous dual-core design and proceed from there. The

dual-core design may have provisions for coupling two chips together in a quad-core processor complex that fits within a 30W power envelope.

On-Chip Networks Must Be Scalable

CoreNet is the most crucial component of Freescale's multicore platform and its prospects for success. Slapping down multiple processor cores on a chip is easier than making them play together like a symphony orchestra. As the number of cores grows, conventional multidrop buses quickly become saturated with intercore bus traffic. The limitations of a conventional bus are one reason Freescale invented a proprietary fabric instead of using IBM's CoreConnect, a freely licensable on-chip bus for Power Architecture chips. If CoreNet succeeds, it will help differentiate Freescale's multicore chips from others. If CoreNet becomes a bottleneck, it will jeopardize Freescale's entire multicore strategy.

Figure 1 is a highly abstract block diagram of a future multicore communications chip using CoreNet. Freescale isn't publicly disclosing detailed technical specifications at this time. However, CoreNet has several important characteristics. First, it supports homogenous or heterogeneous multicore designs—very important for PowerQUICC-type communications chips. Second, it's a high-bandwidth interconnect that supports multiple simultaneous transactions across the fabric. Third, it maintains coherency among the L2 caches attached to each processor core and with the shared L3 cache. Fourth, it allows multiple external-memory controllers to access the fabric simultaneously, without blocking each other. Fifth, CoreNet has multiple address arbiters and automatic buffering.Freescale describes CoreNet as "self-routing." That description implies a packet-based on-chip network in which the processor cores, accelerators, and other on-chip resources have internal network addresses. (Not to be confused with the external network's data packets, which the chip would repackage into internal data packets for travel over the on-chip network.)

Offloading traffic management is an important feature. Unlike some other multicore chips, Freescale's devices won't have to reserve a general-purpose processor core to be the traffic cop. All Power e500-mc cores will be available for application processing. A separate datapath resource manager offloads the traffic control.

Freescale says the datapath resource manager is flexible enough to be tuned for different applications. This statement implies some level of user programmability, at least through an API. One possibility is to equally distribute workloads across all the processor cores and accelerators, to maximize resource utilization. Another possibility is to distribute the resources unequally, to grant higher priority to some data traffic (flow classification). In addition, the datapath manager can adapt its traffic control to varying workloads at run time, providing intelligent load balancing.

The networking industry is buzzing about virtualization technology, which can dramatically slash costs and power consumption by running multiple instances of operating systems on a single system. (See our three-part series beginning with *MPR 3/5/07-01*, "The gHost in the Machine.") Freescale says its multicore platform will have configurable memory regions, so multiple application programs and operating systems can run at the same time without colliding. Virtualized memory management can also prevent caches from thrashing as control passes from one context to another.

Hybrid Simulation for Software Development

It should be obvious that Freescale's new multicore communications platform is a complex piece of work. The big question is whether mere mortals will actually be able to design workable systems around the chips and write the software. It's a difficult question that all multicore chip vendors are trying to answer in one way or another.

Freescale's future multicore PowerQUICC chips will provide some degree of backward software compatibility with today's PowerQUICC chips. But to allow developers to leverage the higher-scale integration to come, better development tools are also a key part of Freescale's multicore platform. One of the most important software-development tools is Simics, a simulation environment from a third-party company, Virtutech. Simics allows developers to run an accurate full-system model of the entire multicore design, including all the Power cores, hardware accelerators, and I/O interfaces.

Traditionally, software developers use simulators to begin writing their programs before the hardware design is complete. However, Freescale expects programmers to use Simics not only for early development but also for the whole development process. Simics provides a full-system functional model and a detailed view of the chip. It has advanced debugging features, such as checkpointing, reverse execution, and full determinism, and it can stop all the processor cores at the same instant. Of course, the drawback of any software simulator is slow execution. Even on a fast workstation, Simics is about 50 times slower than the hardware it simulates—the equivalent of about 40 native mips.

Even so, the Simics functional model is much faster than a cycle-accurate model, which Freescale will also provide. In Freescale's hybrid simulation environment, the Simics functional model will be the first choice, unless developers need cycle accuracy. In that case, developers can switch to the Freescale model and focus on the section of code for which they need cycle accuracy—without leaving Simics. This hybrid approach lets developers reserve cycle-accurate analysis for small portions of code, thus minimizing the loss of performance that normally makes cycle-accurate simulators so tedious to use. Both simulation models can be assisted by sophisticated on-chip debug units and instrumentation built into the chips. (Simics is already available for Freescale's MPC8641D and MPC8572E, and it can model two to eight cores.)

In any event, multicore programming should be easier on PowerQUICC-type chips than on multicore processors

For More Information

Freescale plans to begin sampling the first chips based on its multicore communications platform in late 2008. The Virtutech Simics simulator is available now for Freescale's MPC8572E and MPC8641D processors. Freescale plans to introduce the hybrid simulation environment for the new multicore platform in 4Q07, allowing developers to begin work on their designs before the first devices ship.

To learn more about Freescale's multicore communications platform, see:

www.freescale.com/webapp/sps/site/overview.jsp?nodeId=0162468rH3bTdG25E4

A Freescale white paper on the multicore communications platform is available here:

www.freescale.com/files/32bit/doc/white_paper/MULTICOREFTFWP.pdf

for desktop PCs. It's easier to extract parallelism from tasks like packet processing than it is from general-purpose productivity applications. Networking and communications processors with a dozen or more cores are already available from companies like Cavium and PicoChip, whereas the mainstream PC market is just getting its first quad-core microprocessors.

Preparing for the Manycore Future

Freescale's new multicore communications platform has all the components needed to succeed. Executing the strategy is the challenge. To remain competitive in the high-performance networking and communications markets, Freescale simply must make the strategy work. Clock-frequency scaling with a single core is a dead end, as AMD and Intel ruefully discovered with their PC processors. Multicore designs are the future, and manycore designs already are making serious inroads in the embedded market.

Over the past two years, Cavium has introduced an extensive line of MIPS-compatible networking processors with up to sixteen 64-bit cores per chip, augmented by numerous hardware accelerators and high-speed I/O interfaces. (For a report on Cavium's latest move, see *MPR 7/16/07-01*, "Cavium Stalks Storage.") PicoChip has been shipping its massively parallel processors for years, too. (See *MPR 7/28/03-02*, "PicoChip Preaches Parallelism.") AMCC recently announced a new 32-bit Power Architecture core that could intrude on Freescale's business. (See *MPR 7/23/07-01*, "AMCC's Titan Core.") ARC International, ARM, MIPS Technologies, and Tensilica are all licensing their 32-bit processor cores to customers building multicore chips for networking and communications. To keep up, PowerQUICC must get quicker.

But merely integrating numerous processor cores on a chip isn't enough. Without an efficient on-chip network, bottlenecks will stall the processors, acceleration engines, caches, and other resources that must work together in concert to deliver high performance. Memory coherency becomes a maddening problem as the number of processor cores increases. Providing enough I/O bandwidth to keep the chip fed is yet another challenge. And looming over all the complex hardware is the dark shadow of multicore software development, which requires new programming tools and techniques. This is a time for big visions, not piecemeal solutions to point-source problems.

MPR believes that Freescale has developed a solid multicore platform for a vital product line. Freescale's strategy includes a critical component—the CoreNet on-chip fabric—that could make or break the strategy. The ultimate test will be shipping timely silicon that lives up to the promises.