Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

FREESCALE'S MULTICORE MAKEOVER

New QorIQ Processors Will Eventually Supersede PowerQUICC Chips

By Tom R. Halfhill {7/7/08-01}

They will be powerful and quick, but they won't be PowerQUICC. Instead of using the brand name that has been a household word since 1995—in the households of network engineers, that is—Freescale Semiconductor has unveiled a new name for its future communications processors. The new brand is QorIQ (pronounced "Core IQ"). Although the name doesn't seem like an upgrade, the chips look good.

Among the first six QorIQ devices announced is the P4080, the first eight-processor multicore chip from Freescale. And it's not even the top of the line. Some future QorIQ chips will have at least 16 cores, matching the complexity of the biggest multicore processors from Cavium Networks. Freescale claims its eight-core P4080 will process 18 million network packets per second—fast enough to challenge Cavium's Octeon Plus processors. The catch is that Octeon Plus chips are available now, whereas the P4080 isn't due in volume until 2010.

The PowerQUICC brand and product line aren't going away soon. Freescale says PowerQUICC products could persist for 10 years, and a few more devices might even join the line. But the vast majority of Freescale's new networking and communications processors will be QorIQ devices.

Today, PowerQUICC chips are the industry's most popular communications processors, found in everything from home gateways to back-office routers. Freescale plans to introduce QorIQ chips spanning the same range. To that end, Freescale has created five initial product families within the QorIQ line. Numbered P1 through P5—low end to high end—these families will include low-power, low-cost, single-core chips as well as large-scale multicore devices. Future QorIQ product families may have different prefixes. (The new part numbering brings welcome order to Freescale's product-naming scheme.)

All P-family QorIQ chips are based on the Power Architecture e500 and e500mc 32-bit processor cores. The QorIQ devices announced so far include three members of the P1 family (two single cores and a dual core); two members of the P2 family (with single and dual cores); and the aforementioned eight-core P4080. The P3 and P5 families will be populated later. To get a jump on the competition, all first-generation QorIQ chips will be fabricated in a 45nm silicon-on-insulator (SOI) process. Currently, Freescale's best PowerQUICC devices are fabricated at 90nm, so QorIQ is leapfrogging the 65nm node altogether.

Like their PowerQUICC forebears, QorIQ processors are highly integrated. They have numerous I/O controllers, hardware accelerators, and peripherals on chip. Some future QorIQ chips may be heterogeneous multicore designs that supplement the Power Architecture cores with DSP cores. Another important high-end feature is Freescale's CoreNet Coherency Fabric, a fast on-chip interconnect that links together all the processor cores and internal components. CoreNet is a key part of Freescale's QorIQ strategy, which aims to propel the product line toward the manycore future. (See *MPR 8/27/07-01*, "Freescale's Multicore Strategy.")

Reaching Beyond Dual Cores

QorIQ has big shoes to fill. The PowerQUICC line is broad because it's been around so long and has pioneered the concept of highly integrated networking/communications processors. The first QUICC device, Motorola's 68360, was based on the 68K (68000) architecture and debuted in 1993.

Price & Availability

Some devices in Freescale Semiconductor's new QorIQ product line of networking and communication processors are scheduled to begin sampling this year, with additional parts sampling in 2009. Samples of the single-core P2010 and dual-core P2020 chips are expected to be available in 4Q08; samples of the single-core P1010, single-core P1011, and dual-core P1020 chips are expected to be available sometime in 2009; and samples of the eight-core P4080 chip are expected to be available in mid-2009. Freescale hasn't announced a schedule for volume production and has quoted only two prices so far: $23 for the P1010 and $50 for the P2010 (both in 10,000-unit quantities).

For more information, visit: *www.freescale.com/multicore*.

(See *MPR 5/10/93*, page 13, "68360 Provides Sophisticated Communications.") This chip built on earlier, less specialized 68K-based processors. QUICC, rarely spelled out, stands for "Quad Integration Communication Controller," indicating the addition of dedicated logic for accelerating communication tasks. The first PowerPC-based PowerQUICC processor was Motorola's MPC860, announced in 1995. (See *MPR 9/11/95-02*, "New PowerPCs Aimed at Consumer Devices.")

Freescale is a Motorola spinoff that inherited the PowerQUICC line. Although PowerQUICC chips remain popular, they compete in a world radically changed from 10 years ago. Networking and communications have become huge markets that demand extreme performance at the high end and very low cost at the bottom. Integrated devices like PowerQUICC are now the norm. They are widely available from industry behemoths like Broadcom and from fast-rising newcomers like Cavium. General-purpose microprocessors like the x86 are also making serious inroads. Although Freescale remains the market leader, PowerQUICC no longer has a corner on the market.

In particular, Freescale needs better multicore processors to compete at the lucrative high end. Although counting cores is as misleading as counting clock cycles, multicore integration is the clear path toward higher performance, and the PowerQUICC line has been stalled at dual cores for years. The most powerful PowerQUICC is the dual-core MPC8641D, which was actually designed with the Macintosh in mind, before Apple switched to the Intel x86. (See *MPR 10/25/04-01*, "Embedded CPUs Zoom at FPF.") The MPC8641D shipped about two years late. Meanwhile, Cavium began shipping parts with as many as 16 cores. (Freescale does make two quad-core chips—the MSC8102 and MSC8144—but they are StarCore-based DSPs, not PowerQUICC processors.)

By announcing the eight-core P4080 among the first batch of QorIQ chips, Freescale is signaling that it's back in the fight at the high end. To land punches, however, Freescale must deliver the P4080 on time and within spec. The P4080 is scheduled to begin sampling in mid-2009, but Freescale isn't saying when volume production will begin. Usually, sampling leads to production 6 to 12 months later.

Leapfrogging from 90nm to 45nm while launching an entirely new product line won't be easy. Much depends on Freescale's engineering relationships with fab partners IBM and Chartered. Like almost all semiconductor companies these days, Freescale must rely on partners to share the enormous capital investments in fabrication technology and manufacturing.

QorIQ Spans the Gamut of Applications

The P4080 will anchor the high end of the QorIQ line, at least until Freescale introduces chips in the P5 family, which will be even more powerful. Freescale says the P4080 is suitable for metro-scale routers, integrated management system (IMS) controllers, wireless-network controllers, and routers for General Packet Radio Service (GPRS) and GPRS Support Nodes (GSN).

At the low end of the QorIQ line, P1- and P2-family processors are intended for home gateways, network-attached storage (NAS), wireless base stations, unified threat managers, VoIP gateways, and similar applications near the edges of networks. Most, if not all, QorIQ chips in the P1 and P2 families will have one or two processor cores—sufficient horsepower for their target applications.

Freescale hasn't announced any parts in the P3 family, but some will have more cores, a better on-chip network fabric, and additional hardware acceleration for packet routing and filtering. They are intended for converged media gateways, firewalls, and systems that need hardware acceleration for Secure Sockets Layer (SSL) and Internet Protocol Security (IPsec) transactions.

Likewise, Freescale hasn't announced any P5-family parts, but they will be the most powerful QorIQ chips of all. Target applications include large routers operated by Internet service providers, network admission controllers, and enterprise storage networks. Their estimated maximum power consumption of 30W should be welcomed by data centers struggling with exorbitant electric bills and cooling problems.

Over time, some PowerQUICC parts will be replaced by QorIQ parts. Although some QorIQ parts will be pin compatible with each other, Freescale hasn't said if any will be pin compatible with existing PowerQUICC chips. As the industry moves toward DDR3 memory and other new I/O standards, pin compatibility with older chips becomes more difficult to justify. Software compatibility shouldn't be a problem, though. Both product lines support the Power Architecture and use the same, or very similar, processor cores.

Comparing QorIQ Processors

Table 1 compares the features of the six QorIQ chips announced so far. Distinguishing features include the number of cores, amounts of cache, maximum clock frequencies,

Feature	Freescale QorIQ P1010	Freescale QorIQ P1011	Freescale QorIQ P1020	Freescale QorIQ P2010	Freescale QorIQ P2020	Freescale QorIQ P4080
CPU Core	Power e500	Power e500	Power e500	Power e500	Power e500	Power e500mc
Arch. Width	32 bits	32 bits	32 bits	32 bits	32 bits	32 bits
Pipeline	7 stages	7 stages	7 stages	7 stages	7 stages	7 stages
Cores Per Chip	1	1	2	1	2	8
Core Freq (Max)	800MHz	800MHz	800MHz	1.2GHz	1.2GHz	1.5GHz
L1 Cache (I / D)	32K / 32K per core	32K / 32K per core	32K / 32K per core	32K / 32K per core	32K / 32K per core	32K / 32K per core
L2 Cache	—	256K	256K shared	512K	512K shared	128K per core
L3 Cache	—	—	—	—	—	1,024K x 2
Coherent On-Chip Bus	PowerQUICC type	PowerQUICC type	PowerQUICC type	PowerQUICC type	PowerQUICC type	CoreNet fabric
Memory Controller	DDR2 / DDR3 1 channel	DDR2 / DDR3 1 channel	DDR2 / DDR3 1 channel	DDR2 / DDR3 1 channel	DDR2 / DDR3 1 channel	DDR2 / DDR3 2 channels
DMA Controller	2 x 4 channels	2 x 4 channels	2 x 4 channels	2 x 4 channels	2 x 4 channels	2 x 4 channels
PCI Express	2	2	2	3	3	3
Ethernet Controllers	3 x 1Gb	3 x 1Gb	3 x 1Gb	3 x 1Gb	3 x 1Gb	2 x 10Gb 8 x 1Gb
SerDes	4 lanes	4 lanes	4 lanes	4 lanes	4 lanes	18 lanes
Serial RapidIO	—	—	—	2	2	2
Frame Mgr	—	—	—	—	—	2
TDM	—	Yes	Yes	—	—	—
UART	1 DUART	1 DUART	1 DUART	1 DUART	1 DUART	2 DUART
I^2C	2	2	2	2	2	4
SPI	Yes	Yes	Yes	Yes	Yes	Yes
USB 2.0	1 (ULPI*)	1 (ULPI*)	1 (ULPI*)	1 (ULPI*)	1 (ULPI*)	2 (ULPI*)
SD / MMC	Yes	Yes	Yes	Yes	Yes	Yes
Crypto Accel.	Yes	Yes	Yes	Yes	Yes	Yes
XOR Accel. †	Yes	Yes	Yes	Yes	Yes	Yes
RegEx Engine‡	—	—	—	—	—	Yes
Real-Time Debug	—	—	—	—	—	Yes
IC Process	45nm SOI	45nm SOI	45nm SOI	45nm SOI	45nm SOI	45nm SOI
Voltage (Core)	0.9–1.1V	0.9–1.1V	0.9–1.1V	0.9–1.1V	0.9–1.1V	0.9–1.1V
Power (Max)	<4.0W	<4.0W	5.0W	7.0W	7.0W	<30W
Sampling	2009	2009	2009	4Q08	4Q08	mid 2009
Production	n/a	n/a	n/a	n/a	n/a	n/a
Price (10KU)	$23	n/a	n/a	$50	n/a	n/a

Table 1. Feature comparison of the six Freescale QorIQ chips announced to date. They introduce the P1, P2, and P4 families, to be followed by chips in the P3 and P5 families. All are 32-bit implementations of the Power Architecture, and all but the P4080 use the Power e500 processor core. The P4080 uses the Power e500mc core, an enhanced version of the e500. In addition, the P4080 is the only device in this group with Freescale's CoreNet Coherency Fabric, a new on-chip interconnect. In Freescale's part-numbering scheme, the first two digits identify the QorIQ family, the third digit indicates the number of processor cores, and the last digit identifies variations within a QorIQ family. *ULPI: a triple-nested abbreviation that stands for UTMI+ Low-Pin Interface, where UTMI+ stands for USB 2.0 Transceiver Macrocell Interface Plus, where USB stands for Universal Serial Bus. Essentially, ULPI reduces the pin count of the interface between the USB controller and the USB physical-layer (PHY) chip. †XOR Accelerator: special logic that speeds up exclusive-OR operations, mainly for cryptography. ‡RegEx Engine: special logic that speeds up the regular-expression evaluations in pattern-matching algorithms, mainly for malware filtering. (n/a: data not available.)

integrated I/O controllers, hardware accelerators, and power consumption. Price, of course, is another important differentiation, but Freescale hasn't yet announced pricing or volume availability for all the parts.

The most interesting QorIQ device is the P4080. With eight processor cores and numerous controllers and accelerators, it's by far the most complex chip ever designed by Freescale. It also introduces features that will distinguish higher-end QorIQ devices from existing PowerQUICC chips. Two examples are the Power e500mc processor core and the CoreNet Coherency Fabric.

The Power e500mc is based on the seven-year-old Power e500. Both are 32-bit processor cores with seven-stage pipelines and two-way superscalar execution. However, the Power e500mc has several improved features. One is a private back-side bus for an L2 cache plus additional support logic for a front-side L3 cache. Under this arrangement, each core can have its own L2 cache that doesn't share bus bandwidth with other cores. In other words, the L2 caches behave much like L1 caches—except they are unified, not split into instruction and data memories. The L3 cache, however, is shared among all the processors. It connects to

The New, Improved Power e500mc Processor Core

Freescale has improved the Power e500 processor core in several ways to produce the new Power e500mc core. In addition to support for a three-level cache hierarchy, as the main article describes, the e500mc has a new privileged-execution mode, extended virtual addressing, privileged interrupts, a better load/store unit, a more compatible FPU, larger cache lines, and optional cache stashing. Yet Freescale says the e500mc is software compatible with the e500—including the ability to run existing operating systems that know nothing about the embedded hypervisor.

Enabling the hypervisor is optional. Whether it's enabled or not, it's transparent to other software, thanks to a new privileged-execution mode.The e500 already had two privilege levels for user tasks and system-supervisor tasks. The e500mc has a third privilege level called "guest state." Two bits in the e500mc's state register (PR and GS) define the privilege level for a process. The hypervisor runs at the lowest level (relative to the hardware), with the operating system and user software running in guest state. Within guest state, the operating system still runs at a lower level than user tasks do, so the existing user/supervisor relationship doesn't change. There's simply another privilege level (the hypervisor) running beneath them. The accompanying figure illustrates the difference between an e500mc-based chip running with and without the embedded hypervisor enabled.

Extended virtual addressing allows the hypervisor to access more memory and to partition the system into logical blocks. Processor cores, I/O controllers, and memory can all be partitioned for exclusive access, but they can also share some resources under controlled conditions. A new memory-management unit called the peripheral-access management unit (PAMU) can supervise the memory references of all on-chip I/O controllers that use DMA. When enabled, the PAMU checks a programmable table of physical memory addresses to ensure that a particular I/O controller is allowed to access a particular region of memory. This mechanism helps preserve the integrity of the various memory partitions.

One drawback of a hypervisor is the overhead of performing more checks on memory accesses, register references, and other critical operations. To minimize the performance penalty, the e500mc adds a few shadow registers that duplicate the contents of existing registers. The hypervisor has exclusive access to the shadow registers, allowing guest-level tasks to access the regular registers as usual. Also, the processor can directly handle some interrupts in guest state without invoking the hypervisor.

The improved load/store unit in the Power e500mc core is capable of "fire and forget" atomic operations. A read-modify-write sequence can execute in order as virtually a single operation, sometimes executing in one clock cycle. In some cases, these operations may include a compare instruction. Freescale has added a few instructions to provide these capabilities.

To enable better software compatibility across the QorIQ product line, Freescale has reverted the Power e500mc's FPU to the same type found in the Power e300. It's a classic FPU with single- and double-precision floating-point math. In contrast, the Power e500 executes floating-point math on a signal-processing engine (SPE) derived from the AltiVec vector-processing engine (although neither the e500 nor the e500mc supports AltiVec extensions). Swapping the SPE for a classic FPU does introduce a difference between the e500 and e500mc, which would seem to worsen software compatibility. But the implication is that future QorIQ processors will have a classic FPU, so Freescale is laying the groundwork for them.

Cache lines in the Power e500mc core are 64 bytes long—twice as long as the e500's cache lines. The main reason is that DDR3 memory transfers 64 bytes of data in

(continued on page 5)

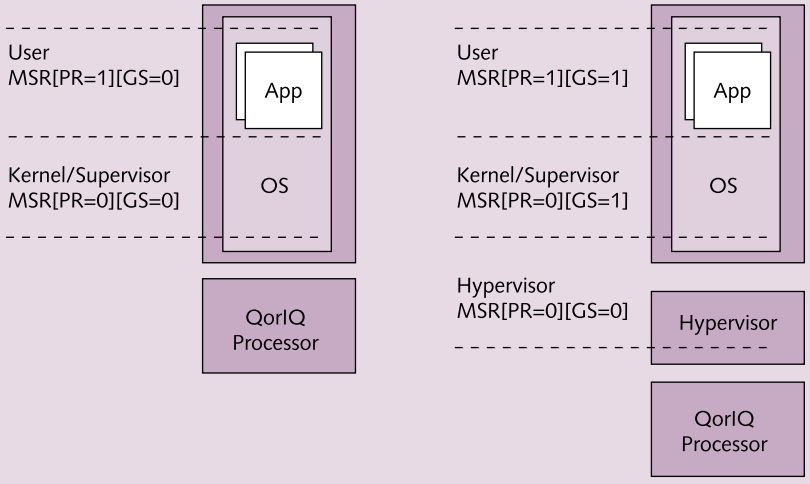

The Power Architecture embedded hypervisor is truly a forklift upgrade. At left, a QorIQ chip with an e500mc processor core is running without the embedded hypervisor enabled. This mode is functionally identical to the e500's existing model of separate user and supervisor modes. At right, with the hypervisor enabled, the user and supervisor modes continue running on separate privilege levels, but they are encapsulated within the new guest mode. Wedged underneath, below guest mode, the hypervisor runs in its own privileged-execution mode. The PR and GS bits in the e500mc's state register define these modes for each level of software.

The New, Improved Power e500mc Processor Core (cont. from p.4)

burst mode, so the e500mc can cache an entire burst in a single line. Another cache improvement in the e500mc is "stashing." The processor can prefetch data into the L1, L2, or L3 caches more intelligently than the usual cache-fill methods that are based on memory locality. For instance, if the processor anticipates that it will need a certain block of data from memory, it can prefetch the data and stash it in a cache, without waiting for a load instruction to reference an address in that memory region. Stashing can be turned on or off under program control.

Freescale has posted a white paper about the embedded hypervisor and Power e500mc processor core at *www.freescale.com/files/32bit/doc/white_paper/EMBEDDED_HYPERVISOR.pdf*.

the CoreNet fabric, which in turn connects to the front-side bus of each core.

This three-level cache hierarchy allots generous amounts of private cache to each core while providing additional shared cache at the highest level of the hierarchy. Shared memory and coherency are vital for symmetric multiprocessing, but local caches keep critical data close to each core. Cache sizes will vary among QorIQ parts. In the P4080, each Power e500mc core has 32KB L1 instruction and data caches, plus a 128KB unified L2 cache. And there are two L3 memories (1MB each), because each DDR channel to main memory has its own L3 cache.

In addition to the three-level cache, the Power e500mc has more improvements over the Power e500. Several of these improvements are related to multicore processing and the embedded hypervisor. (See the sidebar, "The New, Improved Power e500mc Processor Core.") The older, simpler e500 core is sufficiently powerful for single- and dual-core QorIQ devices in the lower range of the product line. (See *MPR 7/16/01-01*, "Speedier Book E Encore.")

More Diversity in the Future

Freescale plans to use the Power e500mc more broadly in future P-series QorIQ chips. In other QorIQ series, Freescale will use different Power processors—such as the Power e300 core, entirely new Power cores, and DSP cores. In future Power cores, deeper instruction pipelining would be advantageous. The seven-stage pipeline in the Power e500 and e500mc just isn't deep enough to wring out the most performance that clock-frequency scaling can deliver. This limitation nudges Freescale toward a larger multicore design before it's absolutely necessary.

Freescale will probably stick with 32-bit implementations of the Power Architecture for most or all QorIQ products. Interestingly, the networking and communications market is split between 32- and 64-bit processors. Two of Freescale's competitors—Broadcom and Cavium—use 64-bit MIPS-compatible processor cores in their chips. On the other hand, all of IBM's licensable Power cores are 32-bit processors, and AMCC's brand-new Titan core is a 32-bit implementation of the Power Architecture. (See *MPR 7/23/07-01*, "AMCC's Titan Core.") Marvell's Feroceon core is a 32-bit custom implementation of the ARM architecture. (See *MPR 5/23/05-01*, "Marvell Puts ARM Out of Order.") Numerous custom chips in these markets are based on 32-bit licensable processor cores from ARM, ARC International, MIPS Technologies, and Tensilica.

Although a 64-bit processor operates on larger chunks of data, which is especially useful in packet processing, the additional transistors required for wider datapaths and registers tend to enlarge the die and use more power. Advocates of 64-bit processing argue that a 32-bit design must compensate for its narrower datapaths and registers in other ways that can use more transistors than a 64-bit design does.

The widespread use of both 32- and 64-bit processors in this market indicates that CPU architects are deeply divided on this issue. In fact, there is plenty of support for both arguments. Although some unannounced QorIQ families might introduce a 64-bit Power core to match Broadcom and Cavium, *MPR* expects Freescale to stick with 32-bit Power cores for most QorIQ devices.

CoreNet Fabric Knits It Together

Figure 1 is a block diagram of the QorIQ P4080. With eight cores, it's a big leap beyond the dual-core MPC8641D, and it doesn't skimp on other components, either. The P4080 has two 10Gb/s Ethernet ports, eight 1.0Gb/s Ethernet ports, three PCI Express interfaces, and two Serial RapidIO interfaces. Two memory channels work with DDR2 or DDR3 SDRAM, each with its own coherent L3 cache. A number of lesser I/O interfaces support USB 2.0, SPI, I^2C, Secure Digital (SD), MultiMediaCard (MMC), and the usual UARTs. Hardware accelerators include a regular-expression (RegEx) engine and a cryptography unit with exclusive-OR (XOR) logic. RegEx engines are a common feature in networking chips. They speed up the regular-expression evaluations in pattern-matching algorithms, which are especially useful when screening network packets for malware signatures. XOR acceleration is useful for applying bit-masks in cryptography.

The CoreNet Coherency Fabric knits everything together and helps maintain coherency among the eight processors and shared memory. It links all the processor cores, L3 caches, high-bandwidth I/O controllers, and accelerators. Although CoreNet is a key enabling technology

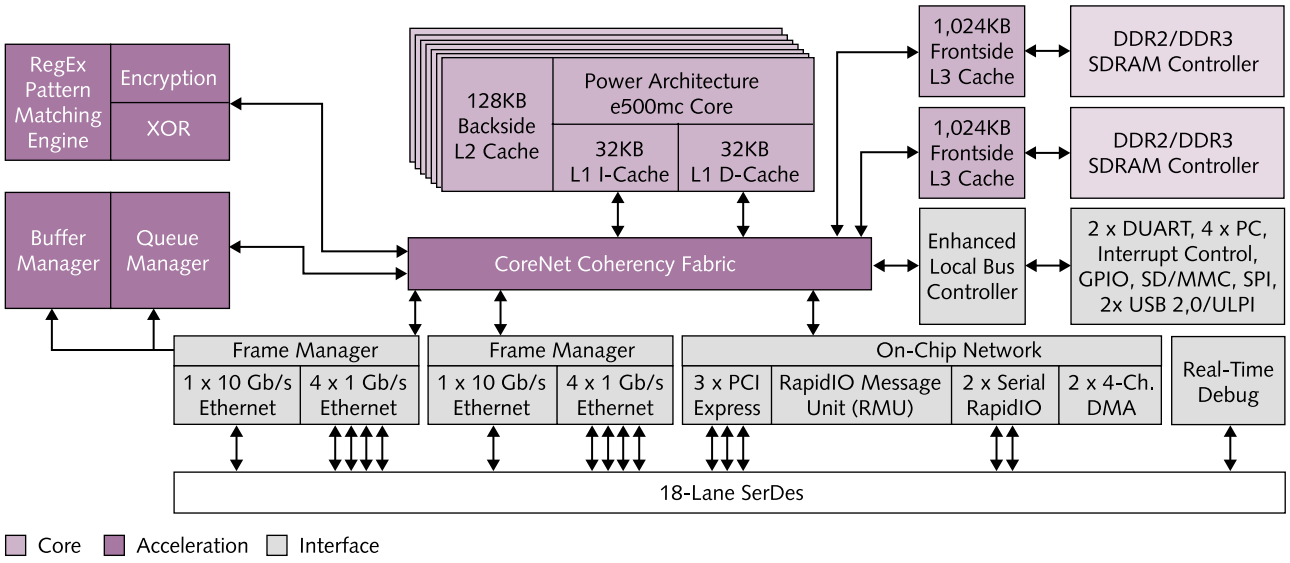

Figure 1. Freescale QorIQ P4080 block diagram. This eight-core networking and communication processor is by far the most complex multicore design from Freescale to date. The new CoreNet Coherency Fabric enables on-chip symmetric multiprocessing (SMP) or asymmetric multiprocessing (AMP). Each processor has two levels of private cache, allowing the cores to perform local processing while reducing accesses to shared memory. Prodigious I/O resources keep the cores fed with network data.

in the QorIQ product line, Freescale hasn't disclosed a great deal of technical detail about it. It's a true fabric, not a multidrop bus, so it can handle multiple I/O transactions simultaneously. Any processor core can send or receive data packets through any network I/O interface without interfering with other cores. Programmers needn't write synchronization code—the fabric manages traffic. Multiple address arbiters and automatic buffers prevent blocking. CoreNet is Freescale's proprietary technology, unrelated to IBM's CoreConnect bus, which is popular in other Power Architecture SoCs.

Freescale describes CoreNet as self-routing, which implies an intelligent packet-based on-chip network. It supports homogenous or heterogeneous multicore designs. That feature is especially important for Freescale, because some PowerQUICC and QorIQ accelerators are proprietary RISC engines with a degree of programmability. In other words, the seemingly homogenous multicore design illustrated in the P4080 block diagram can be construed as a heterogeneous multicore design.

CoreNet works hand in glove with another Freescale technology called the Data Path Acceleration Architecture (DPAA). The DPAA is implemented as multiple blocks of control logic attached to a CoreNet fabric. As with CoreNet, Freescale hasn't disclosed the DPAA in detail, but it's essentially a collection of controllers for internal data-management tasks. Freescale says the DPAA handles some RegEx operations, low-level packet-processing tasks, queue management, and buffer management. By offloading these tasks from the operating system and application software, the DPAA lets programmers focus on writing higher-level code that will differentiate their products from those of their competitors.

Additional logic supports virtualization and an embedded hypervisor. Freescale says these features will enable the P4080 to run any combination of symmetric multiprocessing (SMP) or asymmetric multiprocessing (AMP) operating systems simultaneously. However, one limitation is that the hypervisor can't simultaneously run multiple operating systems on the same Power e500mc processor core.

The embedded hypervisor is a work in progress by a technical subcommittee of the Power.org consortium, which includes Freescale and IBM as leading members. (See *MPR 8/21/06-01*, "The New Power Architecture.") Later this year, Power.org plans to release a revision of the Power Instruction-Set Architecture (Power ISA 2.06) that significantly improves the embedded hypervisor. In a multicore chip, the hypervisor can isolate the processor cores and their software processes from each other. Each processor can access only the memories, I/O controllers, accelerators, and other resources assigned to it. Results: better reliability, availability, and security.

P1, P2 Families Are Poorer Relations

This article focuses on the QorIQ P4080 because it is Freescale's most advanced multicore design and has the greatest number of genuinely new features. In contrast, the first QorIQ chips that Freescale has announced in the P1 and P2 families have few improvements over existing PowerQUICC devices, except for 45nm fabrication. They lack the enhanced Power e500mc processor core, L3 caches, CoreNet fabric, DPAA, RegEx engine, real-time debug module, virtualization extensions, and embedded hypervisor. Indeed, the initial P1 and P2 chips probably began life as

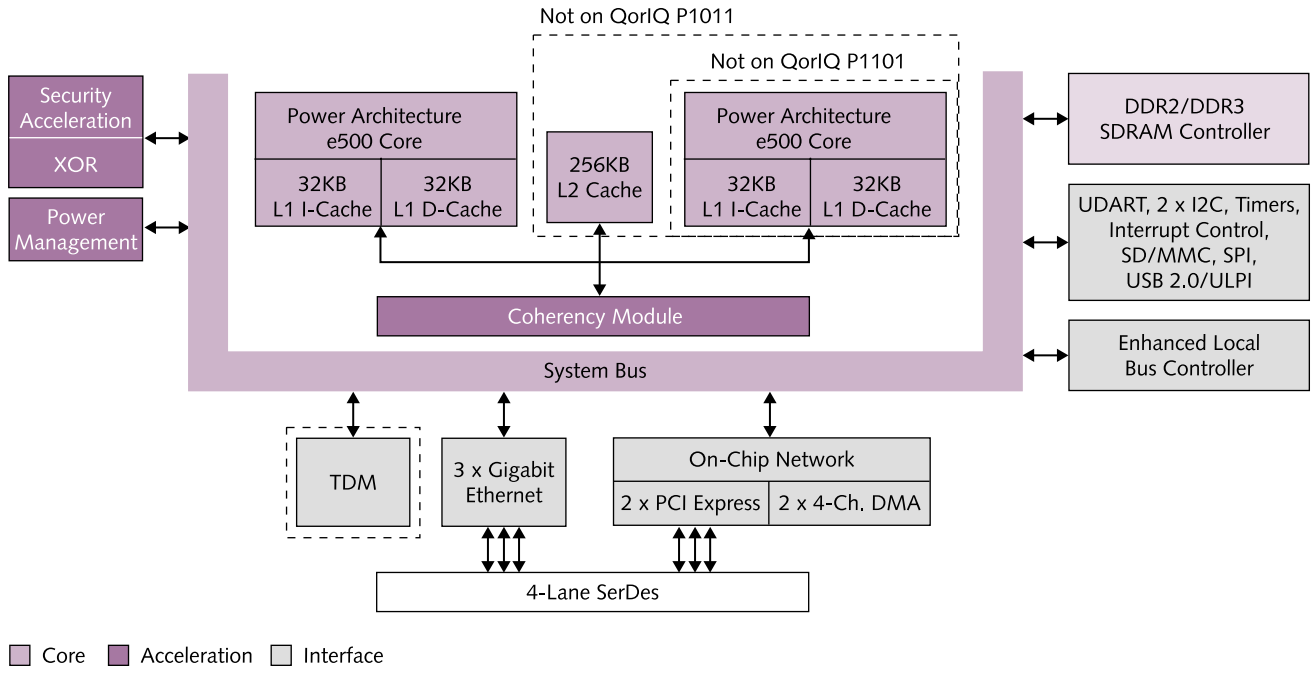

Figure 2. Freescale QorIQ P1-family block diagram. Dotted lines surround features omitted from some members of this family. P1 chips closely resemble existing PowerQUICC designs, lacking most new features introduced with the QorIQ P4 family. However, all QorIQ chips will be fabricated in a 45nm SOI process, greatly improving their power/performance characteristics. The XOR block has special logic for accelerating exclusive-OR operations, often used in cryptographic algorithms.

PowerQUICC designs, then were renamed to anchor the lower end of the QorIQ product line.

Nevertheless, both the P1 and P2 families will benefit mightily from QorIQ's leap to 45nm—two generations beyond the best PowerQUICC chips. Die sizes and power consumption will shrink by about 50%. Freescale says P1 chips will consume less than 4.0W, and P2 chips will use less than 8.0W. Those are maximum power envelopes, when all on-chip peripherals and buses are running. Typical power will be lower.

The first members of the P1 family are the P1010, P1011, and P1020. If we decode Freescale's part-numbering scheme, the P1010 is a single-core processor, the P1011 is an enhanced single-core version of the P1010, and the P1020 is a dual-core processor. Their announced clock speeds range from 400MHz (P1010) to 800MHz (P1020). Therefore, the P1020 is about four times faster than the P1010, because it has twice as many cores running at twice the clock frequency. Figure 2 is a block diagram of these P1 devices.

As Table 1 shows, the only differences between the P1010 and P1011 are their clock speeds and caches. The P1010 lacks an L2 cache; the P1011 has 256KB. The P1020 also has a 256KB L2 cache, which both cores share. By their features, P1-family chips resemble existing PowerQUICC I, II, and II Pro devices. However, they won't be pin compatible with those devices, due to upgraded memory and I/O interfaces. Some P1-family chips will be pin compatible with P2-family chips. For instance, the dual-core P1020 will be pin compatible with the dual-core P2020.

P2 Chips Sacrifice Speed to Cut Power

Only two P2-family devices are announced: the single-core P2010 and the dual-core P2020. The former closely resembles the single-core PowerQUICC III MPC8548, and the latter closely resembles the dual-core PowerQUICC III MPC8572E. Actually, both P2 chips are downgrades from their PowerQUICC III cousins in virtually all respects but die size and power consumption.

For example, whereas the MPC8572E runs the Power e500 cores at a maximum 1.5GHz, the P2020 runs its e500 cores at 1.2GHz. Whereas the MPC8572E supports dual-channel DDR2/DDR3 memory, the P2020 supports only a single memory channel. Whereas the MPC8572E has 1MB of shared L2 cache, the P2020 has only 512KB. Whereas the MPC8572E has a pair of table-lookup units for accelerating pattern-matching algorithms, the P2020 has none. Figure 3 is a block diagram of the first P2-family chips.

Although the P2020 drops some power-management logic found in the MPC8572E, the P2020 will use much less power, thanks to fewer features, the 20% reduction in clock speed, and, especially, the 45nm process shrink. Freescale says the P2020 will operate at 0.9V–1.0V and draw only 7.0W, maximum. In contrast, a 90nm MPC8572E running at the same clock speed (1.2GHz) requires 1.1V and draws about 20W, maximum.

Initial QorIQ chips make extensive use of high threshold-voltage (V_t), low-leakage transistors, trading off clock frequency for lower power. Evidently, Freescale is hearing the

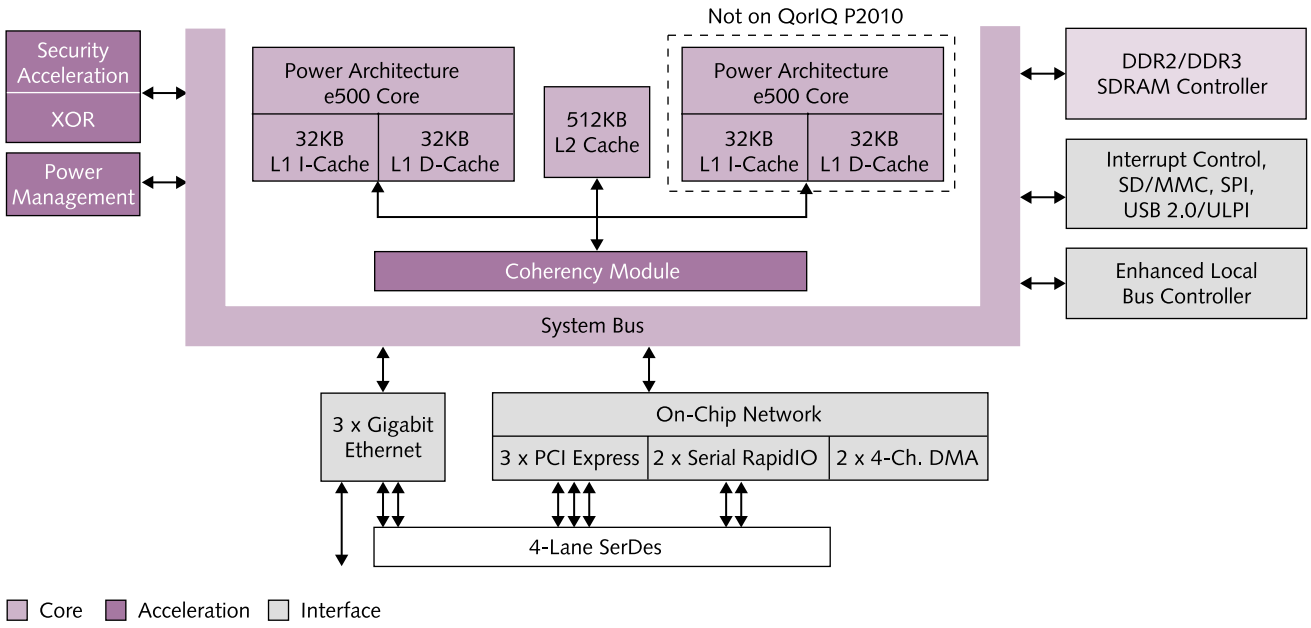

Figure 3. Freescale QorIQ P2-family block diagram. Dotted lines surround the second Power e500 processor core missing from the P2010. Otherwise, the P2010 and P2020 are virtually identical. They are patterned after existing PowerQUICC III MPC8548 and MPC8572E devices, but they are slower and have a few features subtracted.

same customer feedback as other processor vendors: lower power is favored over marginally faster clock speeds. However, Freescale says future P5-family chips in the QorIQ line will compete more aggressively in the clock-speed derby, pushing the Power e500mc core beyond 1.5GHz. To get beyond 2.0GHz, Freescale needs a core with deeper pipelines.

Execution Is Everything

Freescale's QorIQ announcement lays an important foundation for superseding the PowerQUICC product line. In our estimation, the two most important facets of Freescale's strategy are the CoreNet Coherency Fabric and the new hardware support for hypervisors and virtualization.

CoreNet is the crucial element for enabling larger multicore designs. Without a fast, efficient on-chip network, adding more processor cores is pointless. Of course, there's more to building a better communication processor than merely joining the multicore arms race—more cores aren't always better. However, it has become apparent throughout the industry that larger multicore designs are the only path toward significantly higher performance. The power-consumption penalty of higher clock speeds is simply too burdensome. Besides, Freescale's existing processor cores can't reach high clock speeds.

The big question is whether CoreNet is equal to the task of weaving 8, 16, or more processor cores and other components together into a coherent fabric that keeps them usefully busy. Until QorIQ chips are designed into systems and thoroughly tested, there's no way to be certain. Before the first silicon arrives, early developers can test a P4080 design using the Virtutech Simics Hybrid Simulator, an excellent product that creates a fast, functional model of the device. In functional mode, Simics can simulate the P4080 at a speed of about 300 MIPS, according to Virtutech's early estimates. (Green Hills Software used Simics to rapidly port its MULTI development tools and Integrity real-time operating system to the P4080.) However, cycle-accurate simulation is slower, especially for an eight-core processor. Developers will need sample chips before they can perform thorough system-level testing and benchmarking with full confidence.

The embedded hypervisor and virtualization features are easily overlooked improvements. We deem them important for two reasons. First, supporting these functions in hardware should reduce the overhead of virtualization, which taxes performance. Anything that improves virtualization will allow systems to allocate workloads among the cores more efficiently. (See our three-part series on virtualization in *MPR 3/5/07-01*, *MPR 3/12/07-01*, and *MPR 3/26/07-01*, "The gHost in the Machine.")

Second, the task isolation enabled by an embedded hypervisor and virtualization can significantly improve reliability, availability, and security. A common complaint among developers using multicore processors is that a software crash on one core can force the whole chip to reset, bringing down all the processes running on other cores. As the number of cores per chip keeps rising, the penalty in lost clock cycles for restarting those processes becomes more severe. Freescale says that each processor core on a QorIQ chip can recover from a crash by independently

resetting and rebooting, without affecting other cores. That's huge. Task isolation can also guard against security intrusions and malware-induced mischief.

Ultimately, none of these features will matter if Freescale can't deliver the new chips in a timely fashion while hitting the performance targets. The slippage seen with some dual-core PowerQUICC parts could be fatal to the QorIQ line, because Freescale's competitors have become adept in the art of multicore design. The eight-core P4080 is the bellwether. If it's not available in volume by mid-2010 or isn't a real screamer, Freescale's new product line will have trouble competing in the high end of this market.