Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

EMBEDDED CPUS ZOOM AT FPF

Speedy Multicore Chips Dominate Embedded-Processors Session

By Tom R. Halfhill {10/25/04-01}

CPU cores in embedded processors are multiplying like rabbits and sprinting even faster. Four of the six presentations in the high-performance embedded session at **Fall Processor Forum (FPF) 2004** described impressive new multicore designs. One new product family integrates up to 16 cores on a single chip. Clock speeds are soaring to 1.8GHz and beyond.

What's going on? Networking and telecommunications. Although some other embedded applications require high-performance processors, the growing demands of packet routing, control-plane processing, and wireless infrastructures are forcing chip vendors to push their designs farther than ever before. While the advent of dual-core processors is generating enormous excitement in the PC market, most cutting-edge embedded processors now have *at least* two CPU cores. Some chip vendors are pursuing clock frequencies in the 2.0GHz range. Others use coprocessors and application-specific logic to boost performance. Indeed, the latest high-end embedded processors often employ all three of those tactics to some degree.

Six companies delivered presentations in the high-performance embedded-processor session at FPF, and the 13 chips they announced will compete almost head-to-head for the same customers, and frequently for the same applications. Their architectures are diverse but familiar: three product families are MIPS64 compatible, two are based on PowerPC, and one has an ARM-compatible core. The lack of proprietary CPU architectures is a notable departure from the heyday of network processors in the late 1990s. The vendors in this group have found other ways to distinguish themselves from the competition.

Applied Micro Circuits Corp. (AMCC) unveiled its first new PowerPC chip, the 440SPe I/O processor. AMCC recently became a PowerPC vendor by acquiring $227 million worth of chips, licenses, and design teams from IBM Microelectronics. (See the sidebar, "AMCC Strikes a Big Deal for PowerPC," in *MPR 4/26/04-02*, "IBM Loosens Up CPU Licensing.") AMCC is targeting the PowerPC 440SPe almost exclusively at network-storage applications, which sets it apart somewhat from the other new processors described in this session.

Broadcom announced a line of dual- and quad-core SiByte processors based on its efficient SB-1 MIPS64-compatible processor core. The quad-core BCM14xx chips are the latest incarnations of the long-delayed BCM1400, which Broadcom announced at Microprocessor Forum 2002. (See *MPR 10/16/02-01*, "Chip Combines Four 1GHz Cores.") These will be the first Broadcom chips fabricated in a 90nm process, and they're intended for control- and data-plane networking, network storage, wireless infrastructures, and high-density computing.

Cavium Networks revealed the first technical details of its groundbreaking Octeon family, announced a few weeks before FPF. These remarkable devices integrate 2 to 16 MIPS64-compatible processor cores per chip and establish a new data-plane networking category, the network services processor (NSP). For the first time, a single chip has the processing and I/O resources to handle L3–L7 packet routing, deep content filtering, and security. *Microprocessor Report* has already covered the Octeon family in depth. (See *MPR 10/5/04-01*, "Cavium Branches Out.")

Faraday Technology announced its NetComposer-II, a networking processor intended for L4–L7 routing. Unlike the other chips in this FPF session, it's a structured ASIC, not an ASSP. Three optional metal layers provide one million to three million programmable logic gates and 500Kb to 1.5Mb of SRAM. This technology allows designers to customize the chip with application-specific logic in a much shorter turnaround time than is possible with a conventional ASIC. Yet the programmable cells are much faster than FPGA logic, and the chip's manufacturing cost is competitive with ASSPs. Faraday is using a custom-designed ARMv4-compatible processor core in this chip. (See *MPR 5/18/04-02*, "Risk Reduction, Faraday Style.")

Freescale gave a detailed technical presentation on its first dual-core PowerPC chip and a single-core version of the part. The new PowerPC 8641D integrates two e600 CPU

Feature	AMCC PowerPC 440SPe	Broadcom SiByte	Cavium Octeon	Faraday NetComposer-II	Freescale PowerPC 8641/D	PMC-Sierra RM11200
Chip Type	ASSP I/O processor	ASSP BCM12xx (2 CPU) BCM14xx (4 CPU)	ASSP CN34xx (2–4 CPU) CN38xx (8–16 CPU)	Structured ASIC 1–3M prog gates, 0.5–1.5Mb SRAM	ASSP 8641D (2 CPU) 8641 (1 CPU)	ASSP
Applications	SAN, NAS, RAID	NAS, control plane, hi-density computing	L3–L7 routing, balancing, security, content filtering	L4–L7 routing	L2–L4 routing, telecom, storage, gen embedded	Routing, telecom, storage, gen embedded
Architecture	PowerPC	MIPS64	MIPS64-R2	ARMv4	PowerPC	MIPS64
Arch. Width	32 bits	64 bits	64 bits	32 bits	32 bits	64 bits
CPU Core	PowerPC 440	SiByte SB-1	cnMIPS64	FA626	e600 (G4+)	E11K
CPUs	1	2 or 4	2, 4, 8, 16	1	1 or 2	2
Core Frequency	533–667MHz	800MHz–1.2GHz	600MHz	600MHz	>1.5GHz	1.8GHz
Superscalar	3-way	4-way	2-way	—	4-way	2-way
ALU Pipeline	7 stages	9 stages	5 stages	8 stages	7 stages	7 stages
ALU Instr/Sec	1.6–2.0 billion	6.4–19.2 billion	2.4–19.2 billion	600 million	>6–12 billion	7.2 billion
L1 Cache I/D	32K/32K	32K/32K	32K/8K	32K/32K	32K/32K	64K/32K
L2 Cache (Total)	256K	512K or 1MB	Up to 1MB	128K	1MB or 2MB	1,024K
FPU	—	4 or 8	—	—	1 or 2	2
DRAM Controller (MHz)	DDR1 200–333 DDR2 400–667	DDR1 400 DDR2 800	DDR1 400 DDR2 800	DDR1 400	DDR2 667MHz	DDR2 800
DRAM I/O Width	32 or 64 bits	2x64 bits or 4x32 bits	64 or 128 bits	64 bits	2x64 bits	2x64 bits
Max DRAM B/W	5.3GB/s	12.8GB/s	6.4 or 12.8GB/s	3.2GB/s	10.6GB/s	12.8GB/s
Max DRAM	16GB	16GB	16GB	4GB	16GB or 32GB	1TB
Other Memory Interfaces	SRAM, flash, boot ROM	Boot ROM	RLDRAM, FCRAM, TCAM, boot ROM	Flash, SRAM	Local bus for boot ROM, flash	Boot ROM
Ethernet MAC	1xGbE	4xGbE,TOE, FIFO	0, 4, or 8 GbE	2xGbE	4xGbE	4xGbE
PCI/PCI-X (MHz)	32/64b PCI-X (133-266)	64b PCI-X (133) 64b PCI (66)	64b PCI-X (133)	32b PCI-X (64)	—	—
PCI Express	1 x8-lane 2 x4-lane	—	—	Gate-programmable	2 x8-lane	2 x4-lane or 1 x8-lane
HyperTransport	—	0 or 3x16 bits	—	Gate-programmable	—	1x8 bits
SPI-4.2	—	0 or 3	0, 1, or 2	Gate-programmable	—	—
Other I/O	I^2C, UART, GPIO	GMII, SMbus, UART, PCMCIA, GPIO	4 or 8 RGMII, flash, UART, MDIO, GPIO, 2-Wire	Gate-programmable RapidIO, SATA, FC	Serial RapidIO, UARTs, I^2C	n/a
Hardware Acceleration	XOR, XOR DMA, I^2O DMA	Packet DMA, hash, route, checksum	TOE, ZIP compress, reg-exp, malloc	Gate-programmable	TCP/IP checksum, IPv6, TOE, QoS	Direct I/O to L2 cache
Crypto Hardware	—	—	DES, 3DES, AES, RSA, DH, MD5, SHA-1, RC4	Gate-programmable	—	—
True RND	—	—	Yes	—	—	—
Fab Process	IBM 0.13µm	90nm	TSMC 0.13µm	0.13µm	90nm SOI	90nm
Packaging	FC-PBGA 675	BGA-1936	709 or 1,500 pins	SBGA 484	HiTCE 960	n/a
Power (Typical)	4.7–10.8W	13–23W (1GHz)	5–25W (worst)	3W	10–25W	~15W
Price (Units)	n/a	$599–$1,199 (10K)	$125–$750 (10K)	$48 (1K)	n/a	n/a
Production	3Q05	1H05–2H05	2H05	3Q05	1H06	4Q05–1Q06

Table 1. Here's a broad feature comparison of the new chips presented in the high-performance embedded-processors session at Fall Processor Forum 2004. In all, there are 13 new chips from six companies. Four companies introduced multicore designs ranging from two CPU cores to 16 CPUs per chip. Note the growing proliferation of Gigabit Ethernet, PCI Express, DDR2 memory controllers, and 90nm fabrication. All the chips are scheduled to hit the market within about six months of each other, and all are ASSPs, except for Faraday's NetComposer-II, a structured ASIC. n/a = not available.

cores and is designed for L2–L4 routing, telecommunications, network storage, and general-purpose embedded applications. Clock frequencies will exceed 1.5GHz in a 90nm fabrication process, and the dual-core chip will compete directly with the latest multicore offerings from Broadcom and PMC-Sierra.

PMC-Sierra unveiled the RM11200, a dual-core MIPS64-compatible processor for routing, telecommunications, network storage, and general embedded applications. Its new E11K processor core is freshly ported to a 90nm fabrication process and is designed to hit 1.8GHz, the fastest clock speed announced in this session. The RM11200 will compete directly with the new multicore processors from archrivals Broadcom and Freescale.

All these chips are scheduled to enter production in 2H05 or 1H06. They will offer a formidable amount of processing power to embedded-systems designers. Table 1 summarizes their features; with so much variety to choose from, there's something for just about everyone. The only drawback is that the multicore chips may pose an equally formidable programming challenge to operating-system vendors and software developers.

AMCC Targets Network Storage

AMCC's PowerPC 440SPe stands out from the other chips in this session because it's an I/O processor primarily intended for network-storage applications: storage-area networks (SAN), network-attached storage (NAS), and direct-attached storage (DAS), including internal and external RAID subsystems. These are AMCC's traditional target markets. The 440SPe is an incremental but important upgrade over its predecessor, IBM's 440SP, primarily because it adds PCI Express.

In an internal RAID, the 440SPe will manage data traffic between the disk controllers and the system's host processor. In an external RAID, the 440SPe connects the disk controllers to a network processor. The 440SPe can serve multiple roles by providing a bridge among the various I/O interfaces and by offloading tasks from the host processor and disk controllers. By running the RAID software, for example, the 440SPe allows vendors to build a RAID subsystem using low-cost, off-the-shelf disk controllers and adapter cards.

The 440SPe is appropriately equipped for storage applications. Inside is IBM's five-year-old but still competitive PowerPC 440 processor core—a three-way superscalar design with out-of-order execution, seven-stage integer pipelines, dynamic branch prediction, and a single-cycle 32-bit multiplier. The PowerPC 440 also has 41-bit virtual memory addressing; a 64-entry unified translation lookaside buffer (TLB); and a pair of micro-TLBs for instructions and data, all of which make the processor suitable for control tasks. (See *MPR 10/25/99-03*, "IBM PowerPC 440 Hits 1,000 MIPS.") The core runs at 533–667MHz and includes a 32KB instruction cache, a 32KB data cache with parity protection, and a 256KB on-chip L2 cache with parity.

To augment the core, the 440SPe has some application-specific logic: an exclusive-or (XOR) engine. It speeds up parity-bit generation and parity checking by performing bit-wise XOR operations on as many as 16 simultaneous data-streams. The XOR engine has its own DMA controller, too. Of course, parity operations are critical for error checking and correction (ECC), particularly in RAID software, so the XOR engine is a useful feature in a storage processor.

The 440SPe's biggest improvement is its state-of-the-art I/O. It adds three independent PCI Express interfaces to the PCI-X, Gigabit Ethernet (GbE), and miscellaneous I/O interfaces already found in the 440SP. The PCI Express interfaces include an eight-lane primary port, two four-lane secondary ports, internal bridges between the ports, and an internal bridge to PCI-X. Each PCI Express lane can transfer data at 2.5Gb/s bidirectionally. All these features make the 440SPe suitable for the latest board designs that use serial PCI Express instead of parallel PCI as the channels to Serial ATA, Serial Attached SCSI (SAS), and other high-speed disk controllers. Because PCI-X and regular PCI will remain popular for several years, it makes sense to continue supporting those standards, too. The 32/64-bit PCI-X controller transfers data at the conventional 133MHz rate or the PCI-X v2.0 double data rate of 266MHz.

The 440SPe's GbE interface is intended as a management port, not a primary I/O interface. Other I/O resources include three UARTs, two I^2C interfaces, an I^2O messaging unit, 32-bit general-purpose I/O (GPIO) ports, and the usual interfaces for external SRAM, flash memory, and boot ROM.

Is PCI Express Like Thermal Underwear?

Another useful feature in the PowerPC 440SPe is an on-chip SDRAM controller, which AMCC improved over the controller in the 440SP. It has a 32/64-bit data bus with optional ECC that addresses up to 16GB of physical memory in four banks. (The 440SP addresses only 4GB.) The controller supports DDR1 memory at effective bus speeds up to 333MHz (actual clock speed 166MHz) and DDR2 memory at effective bus speeds up to 667MHz (actual clock speed 333MHz). A special "data saver" feature sustains the refresh signal to DRAM during soft resets.

To tie all the memory controllers and I/O interfaces together, AMCC uses a multi-tier two-way crossbar version of IBM's CoreConnect bus. Its four 128-bit-wide datapaths run at clock speeds up to 166MHz, providing up to 10.4GB/s of bandwidth. Figure 1 shows a block diagram of the PowerPC 440SPe.

Using IBM Microelectronics as the foundry, AMCC will manufacture the PowerPC 440SPe in a 0.13-micron copper CMOS process. Samples will be available in 1Q05, with production scheduled to begin in 3Q05. Estimated power consumption is 14.4W maximum, but typical consumption will be about 5W–11W. The "typical" estimate varies so widely because it depends on the way customers use the I/O interfaces. AMCC's presentation at FPF included an interesting

discussion of this point, backed up by some actual power measurements and preliminary estimates. Summary: PCI Express appears to eat power like popcorn.

The PowerPC 440SPe's predecessor, the 440SP, typically consumes 4.7W when using PCI-X at 133MHz and 7.6W when using PCI-X at the doubled 266MHz data rate. When the new PowerPC 440SPe uses two PCI Express x4 lanes instead of PCI-X, power consumption leaps to 10.8W. PCI Express by itself accounts for more than 6W of that total. In contrast, the CPU core uses only 1W at its top speed of 667MHz. Indeed, the CPU core, on-chip SRAM, SDRAM I/O, miscellaneous logic, and static current leakage for a 667MHz PowerPC 440SPe add up to about 5W, which means PCI Express is by far the leading power sink on the chip. And that's using only the four-lane secondary PCI Express interface, not the eight-lane primary interface.

Customers will see only a little extra bandwidth for that jump in power consumption. Two x4 lanes provide 20Gb/s, compared with 17Gb/s for 64-bit PCI-X at 266MHz. However, power consumption isn't really an issue with the PowerPC 440SPe. Even at its estimated maximum of 14.4W, it won't blow the power budget of a network-storage subsystem, and it has power-management features to reduce active power under typical usage patterns. Other storage-subsystem components (like the disk drives) are a more significant power drain.

More important is what the 440SPe says about the effect of adding PCI Express to a processor. The sudden rise in power consumption could become an obstacle as PCI Express filters down into smaller systems with tighter thermal requirements, such as notebook PCs. Although the PowerPC 440SPe isn't designed for those systems, it may be sounding an alarm bell for designers at large.

Broadcom's Latest Fearsome Foursome

At Microprocessor Forum 2002, Broadcom announced the SiByte BCM1400—a quad-core version of its dual-core BCM-1250 processor—and said the product would ship in 2H03. That didn't happen. At FPF this month, Broadcom returned to the stage and promised to ship two improved versions of the BCM1400 and a pair of new dual-core chips in 2005. Despite a two-year production delay, the company's maximum target clock frequency (1.2GHz) is only 20% higher than the target frequency announced at MPF 2002 (1.0GHz).

Broadcom's troubles illustrate that even a company with experience in multicore designs can hit stumbling blocks when moving beyond two CPU cores. (AMD and Intel, take notes.) Broadcom says the BCM1400 was stillborn because customers wanted more I/O options and a dual-core successor to the BCM1250. As a result, Broadcom is adding GbE and PCI-X to all four new chips and is introducing the dual-core BCM1255 and BCM1280. The new quad-core chips are the BCM1455 and BCM1480.

Production delays haven't crippled Broadcom in the marketplace, thanks to the still-available BCM1250. Keep in mind that one of Broadcom's closest competitors for MIPS64-compatible processors, PMC-Sierra, hasn't even announced a quad-core design yet. Another close competitor, Freescale, just announced its first dual-core PowerPC this month. If Broadcom can make its ship dates this time, its new dual- and quad-core processors will probably hit the market shortly before the latest dual-core products from Freescale and PMC-Sierra.

However, all three of those companies face impressive new competition from Cavium Networks. Cavium's Octeons—with 2 to 16 MIPS64-compatible cores per chip—should reach the market within months of Broadcom's latest parts. They will compete head-to-head with processors from Broadcom, Freescale, and PMC-Sierra for L4–L7 packet-processing applications that also require security. On the other hand, the Broadcom, Freescale, and PMC-Sierra

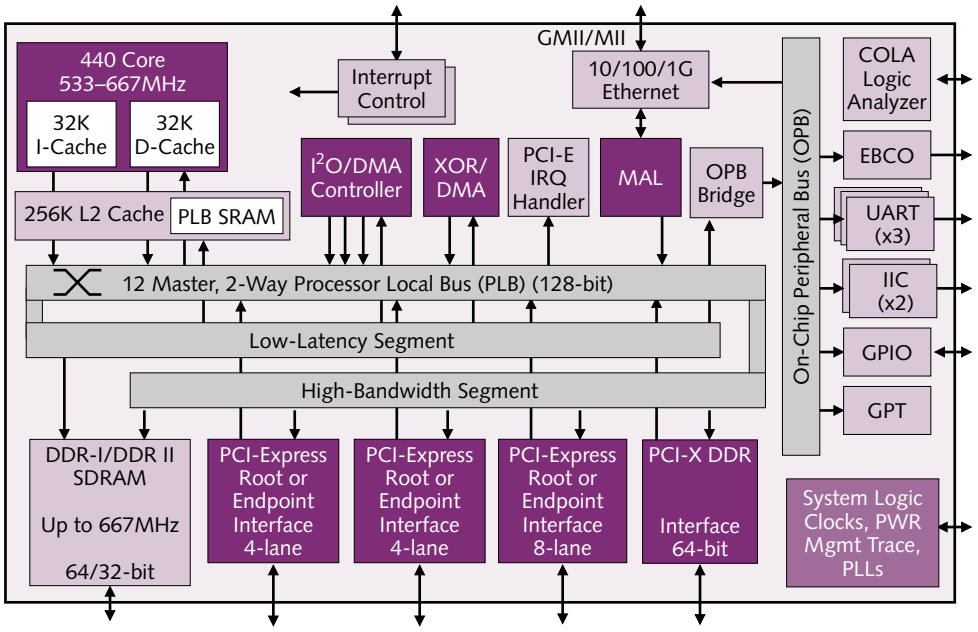

Figure 1. PowerPC 440SPe block diagram. With three independent PCI Express interfaces and 266MHz PCI-X, the 440SPe is well equipped for network-storage applications. Note the multi-tier implementation of IBM's CoreConnect bus, here with four 128-bit pathways, including a two-way crossbar supporting 12 masters. The fastest tier—the processor local bus—runs at one-fourth the speed of the PowerPC 440 core, up to 166MHz in the 667MHz version of the chip. Any two masters may access the bus simultaneously, and any master can access any slave.

chips have wider applications than Cavium's chips, which contain a great deal of hardware specific to data-plane packet processing. Nevertheless, it's obvious that the market for high-performance networking processors is heating up quickly.

In all, Broadcom announced four new SiByte processors at FPF: the dual-core BCM1255 and BCM1280 and the quad-core BCM1455 and BCM1480. They are direct descendants of Broadcom's successful dual-core BCM1250, introduced in 2001. All are based on a powerful MIPS64-compatible processor core custom-designed by SiByte, a startup company acquired by Broadcom in 2000. The SB-1 is the only MIPS64 core with four-way superscalar pipelines; other MIPS64 cores are two-way superscalar machines. (See *MPR 6/26/00-04*, "SiByte Reveals 64-Bit Core for NPUs.")

Feature	Broadcom SiByte BCM1255	Broadcom SiByte BCM1280	Broadcom SiByte BCM1455	Broadcom SiByte BCM1480
Architecture	MIPS64	MIPS64	MIPS64	MIPS64
CPU Core	SB-1	SB-1	SB-1	SB-1
CPUs	2	2	4	4
Core Frequency	800MHz–1.2GHz	800MHz–1.2GHz	800MHz–1.2GHz	800MHz–1.2GHz
Superscalar	4-way	4-way	4-way	4-way
ALU Instr/Sec	6.4–9.6 billion	6.4–9.6 billion	12.8–19.2 billion	12.8–19.2 billion
L2 Cache	512K	1MB	1MB	1MB
DRAM Controller	DDR1 400MHz DDR2 800MHz	DDR1 400MHz DDR2 800MHz	DDR1 400MHz DDR2 800MHz	DDR1 400MHz DDR2 800MHz
PCI-X	64-bit 133MHz	64-bit 133MHz	64-bit 133MHz	64-bit 133MHz
PCI Express	—	—	—	—
HyperTransport*	—	3	—	3
SPI-4.2*	—	3	—	3
Gigabit MII (GMII)	4	4	4	4
Ethernet MAC	Gigabit Ethernet (w/ FIFO mode)	Gigabit Ethernet (w/ FIFO mode)	Gigabit Ethernet (w/ FIFO mode)	Gigabit Ethernet (w/ FIFO mode)
Power (typical)	13W @ 1.0GHz	17W @ 1.0GHz	19W @ 1.0GHz	23W @ 1.0GHz
Price (10K)	$599	$699	$999	$1,199
Production	2Q05	2Q05	3Q05	3Q05

Table 2. Broadcom announced four new SiByte processors at FPF. This table highlights their distinguishing features (CPU cores, L2 cache, performance, and I/O) in purple text. *The HyperTransport and SPI-4.2 controllers share the same I/O ports and are independently software switchable. n/a = not available.

Except for the number of CPU cores (two or four), the only features that distinguish Broadcom's new chips from each other are their on-chip L2 caches and I/O. The BCM1255 has 512KB of L2 cache; the others have 1MB. The BCM1280 and BCM1480 have three software-configurable I/O ports that can independently operate as HyperTransport or SPI-4.2 interfaces; the BCM1255 and BCM1455 do not. Instead, the latter two chips must rely on their other I/O resources, which are common across the whole product line: four Gigabit Media-Independent Interfaces (GMII); four GbE MACs with configurable FIFO modes; and a 64-bit 133MHz PCI-X interface. All the chips also have miscellaneous other I/O, such as UARTs, PCMCIA, SMbus, and GPIO. Table 2 summarizes the important features of Broadcom's new SiByte processors.

Coherent Caches and Memory

Little about these processors has changed since *MPR* covered the BCM1400 in depth after MPF 2002. (See *MPR 10/16/02-01*, "Chip Combines Four 1GHz Cores.") Figure 2 shows a block diagram of the new chips. The GbE and PCI-X controllers are new additions—the BCM1250 has only a 32-bit 66MHz PCI interface. Moore's law rolls on, but Broadcom is aiming for a clock speed only 20% higher, which suggests the company is being more conservative this time or thinks 1.2GHz is fast enough for the target applications. It should be fast enough; only PMC-Sierra announced a faster processor in this session at FPF. Note that a 1.2GHz quad-core SiByte processor with its four-way superscalar pipelines can theoretically execute 19.2 billion integer instructions per second, which matches the maximum theoretical peak performance of a 16-core Cavium Octeon at 600MHz. And in applications requiring high floating-point performance, Broadcom's SiByte processors have a clear advantage over the FPU-less Octeons.

Some important features are carried over from the BCM1250. All the shared L2 caches in Broadcom's multicore processors support coherency in hardware, even across multiple chips. In addition, the CPUs can bypass the L2 cache

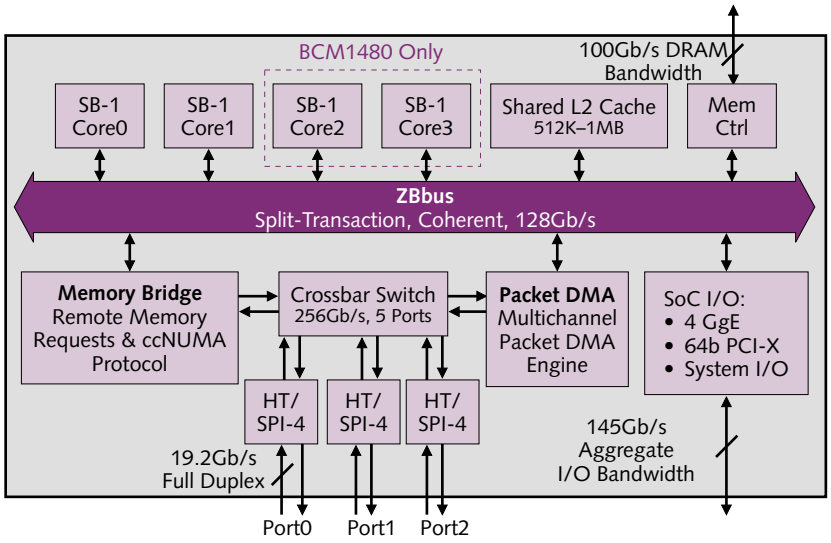

Figure 2. Broadcom BCM12xx/BCM14xx block diagram. The new BCM1255 and BCM1280 are dual-core processors, whereas the BCM1455 and BCM1480 are quad-core chips. The BCM1255 has 512KB of L2 cache; the others have 1MB. Other distinguishing features are I/O interfaces: the BCMxx80 chips have HyperTransport and SPI-4.2. All the chips have four GbE MACs, four Gigabit Media-Independent Interfaces (GMII), and PCI-X.

altogether when prefetching data, which avoids cluttering the cache with data that only one CPU needs for a short time. Headers can be stripped off packets and cached while payloads stream into main memory.

A programmable hash engine can look 256 bytes deep into a packet, distinguish IPv4 addresses from IPv6, and create a decision tree for routing packets to various buffers and queues. In addition, the protocol-agnostic hash engine can forward packets that won't be handled by this processor to another destination—all without loading the internal bus or involving the CPU cores. As was planned for the BCM1400, the BCM1280 and BCM1480 can use HyperTransport as a conduit for ccNUMA commands to shared memory. (For more details of the last feature, see our previously referenced article on the BCM1400.)

The 1.0GHz BCM1400 was supposed to cost $1,200, had it shipped in 2003. Now Broadcom says the 1.2GHz BCM1455 will cost $999, and the 1.2GHz BCM1480 will cost $1,199 when they ship in 2005. In view of Cavium's pricing—the 16-core 600MHz Octeon will cost only $750—the new Broadcom processors seem expensive. Cavium is taking advantage of lower manufacturing costs by using a 0.13-micron fabrication process, whereas Broadcom is using 90nm technology. As the cost curve of 90nm fabrication declines, Broadcom should be able to reduce prices. Note that the existing BCM1250 costs about 33% less than it did two years ago. PMC-Sierra hasn't yet announced pricing for its dual-core RM11200, so it's too early to say whether Broadcom's pricing for the BCM1255 and BCM1280 is out of line.

Broadcom says the dual-core chips are scheduled to sample this quarter and will be ready for production in 2Q05. Working samples of the quad-core chips are available now and production is scheduled to begin in 3Q05. Pushing these chips out the door will be a big step for Broadcom, which virtually inherited the BCM1250 ready-made from SiByte and hasn't delivered a significantly new version in four years. That's a long drought in a highly competitive market.

Faraday Attacks ASSPs

Faraday was the odd man out in the high-performance embedded-processors session at FPF. While the other five companies unveiled new ASSPs, Faraday introduced its NetComposer-II (NC-II), a structured ASIC. The NC-II is a basic L4–L7 networking processor with Faraday's metal programmable cell array (MPCA) in three optional metal layers, which the fab can add during chip manufacturing. The MPCA lets designers use one million to three million programmable gates and 512Kb to 1.5Mb of SRAM, depending on the design. This flexibility allows designers to customize the chip without suffering the higher costs and delays of a full-custom ASIC project. As with all structured ASICs, the NC-II offers a compromise of design cost, manufacturing cost, time to market, performance, and differentiation. Table 3 summarizes those trade-offs.

Programmable gates in a structured ASIC can switch faster than those in an FPGA: 166–333MHz in the NC-II. However, they're not quite as fast or as compact as standard cells in an equivalent fabrication process. Faraday estimates a 10% speed penalty and a 35% area penalty, compared with a standard-cell ASIC. Customers willing to make those trade-offs can save more than $1 million on a design project. Most of the savings are the result of the NC-II's requiring less engineering effort to customize and only three additional masks for the MPCA metal layers, instead of the full mask set of an ASIC. (Counting the via layers between the metal layers, the NC-II requires a maximum of five extra masks.)

Customers can use Faraday's programmable gates to add application-specific logic, such as cryptography engines, packet-processing accelerators, and I/O controllers. The I/O possibilities are especially intriguing. The NC-II has an eight-port programmable serializer/deserializer (SerDes) that designers can configure to support any combination of Fibre Channel, GbE, HyperTransport, PCI Express, RapidIO, Serial ATA, SPI-4.2, or the Xilinx Attachment Unit Interface (XAUI).

In addition, the SerDes has 62 multipurpose I/O (MPIO) pins that the third metal layer can program for PCI/PCI-X and interfaces to TTL/CMOS logic and stub-series terminated logic (SSTL-2). For example, by programming the MPCA and MPIO pins, designers can implement a Utopia interface. The NC-II's versatility with I/O is one of its most impressive features. It also distinguishes the NC-II from Faraday's previous NC-I, which doesn't have the SerDes.

Inside the fixed portion of the NC-II is Faraday's custom-designed FA626 ARMv4-compatible processor core. ARM cores are generally associated with low power, not high performance. Certainly, the 32-bit uniscalar FA626 is less powerful than the 64-bit superscalar MIPS cores from Broadcom, Cavium, and PMC-Sierra. Nor is it a match for the 32-bit superscalar PowerPC cores from AMCC and Freescale. Nevertheless, the FA626 is capable of powering a Layer 4+ packet processor for low-end to midrange applications, especially when it gets help from some custom logic in a structured ASIC like the NC-II.

The FA626 core has an eight-stage ALU pipeline, a multiply-accumulate unit, a barrel shifter, 32KB instruction/data caches, and a two-level translation

	ASSP	FPGA	ASIC	NetComposer
Customer Manufacturing Time	None	None	120 days	30 days
Customer NRE* Cost	None	None	>$1 million	$300,000
Performance	Good to very good	Poor to good	Very good	Very good
Customer Chip Cost	Low	High	High	Low
Design Differentiation	Low	High	Very high	High

Table 3. Faraday's NetComposer structured-ASIC technology attempts to strike the best balance between cost and performance. Cost estimates in this table assume a 0.13-micron CMOS fabrication process for all chips. *NRE: nonrecurring engineering costs.

lookaside buffer (TLB). In UMC's 0.13-micron process, it runs at 600MHz at a core voltage of 1.08V. The on-chip 128KB L2 cache runs at 200MHz. Overall, the NC-II's power consumption is remarkably low, an obvious benefit of using an ARM core: only 3W (typical, 125°C.). Figure 3 shows a block diagram of the NC-II, with the FA626 core at the upper right.

For on-chip interconnects, Faraday uses its proprietary M-Hub crossbar, a nonblocking switch fabric. The M-Hub has 128-bit datapaths and runs at 200MHz, providing 3.2GB/s of on-chip bandwidth. Among the bus masters attached to the crossbar are the FA626 processor core, L2 cache controller, 64-bit DDR memory controller, DMA controller, and programmable cell array. AMBA high-speed and peripheral buses provide bridges to the on-chip 32-bit SRAM/flash controller and miscellaneous I/O interfaces, such as UARTs, I^2C, and GPIO.

When an application requires more processing power or bandwidth than a single NC-II can provide, the programmable SerDes allows designers to attach multiple NC-II chips together. One method is to connect the chips directly, in serial or parallel, by programming the SerDes to implement a standard or proprietary I/O interface. Another method is to connect the chips to an external switch fabric, again by programming the SerDes to support the desired interface. This flexibility allows designers to adapt the NC-II to many different system architectures, almost without regard to the I/O requirements. It's an especially valuable feature at a time when I/O standards are in flux. PCI-X, HyperTransport, PCI Express, RapidIO, and other interfaces are all vying for the attention of designers.

Faraday says the NC-II will be available in 3Q05 for $48 when purchased in 1,000-unit quantities. That price is surprisingly competitive with those of ASSPs aiming at the same networking applications. Of course, almost all competing ASSPs have more-powerful CPU cores than the NC-II, and a growing number of ASSPs are sprouting multiple CPUs. However, the NC-II's programmable logic allows designers to create a chip that performs more like an ASIC than an ASSP with a general-purpose processor core.

In general, an ASSP makes more sense than a structured ASIC if it has the features needed for a particular system design, and if the system doesn't depend on the ASSP for product differentiation. The most grateful customers for the NC-II will be designers who have unusual problems to solve; who can't wait for an ASSP to support the I/O interfaces they need; or who need to differentiate their product at the chip level.

Freescale's First Dual-Core PowerPC

Three ways to describe Freescale's new MPC8641D: it's Freescale's first dual-core PowerPC processor; it's essentially a dual-core version of the equally new MPC7448; and it's a dual-core MPC7448 with the peripherals and interfaces of a PowerQuicc III MPC8548 networking processor. If you haven't heard of the MPC8548, that's because it's brand-new, too, and will be available as a separate part. The MPC8548 is essentially a PowerQuicc III MPC8540 with several enhancements and a process shrink to 90nm.

Freescale is aiming the dual-core MPC8641D at midrange-to-high-end line cards and network-services cards, which explains the merger between the MPC7448 and PowerQuicc III. Adding a second PowerPC processor core to the chip will further reduce the system's chip count and expand the range of possible applications. A design that requires two MPC7448 processors and a system controller can meet the same specifications with a single MPC8641D. For designs that don't need the processing power of the dual-core MPC8641D, Freescale is also offering a single-core version, the MPC8641. In all other respects, it's identical to the MPC8641D.

The processor core duplicated in the MPC8641D is the 32-bit PowerPC e600, a slightly enhanced version of Freescale's 32-bit G4+ core, which appears in most of the 0.13-micron MPC744x chips. Compared with the G4+, the e600 can execute Altivec instructions out of order; has twice as much on-chip L2 cache (1MB per core vs. 512KB); and adds

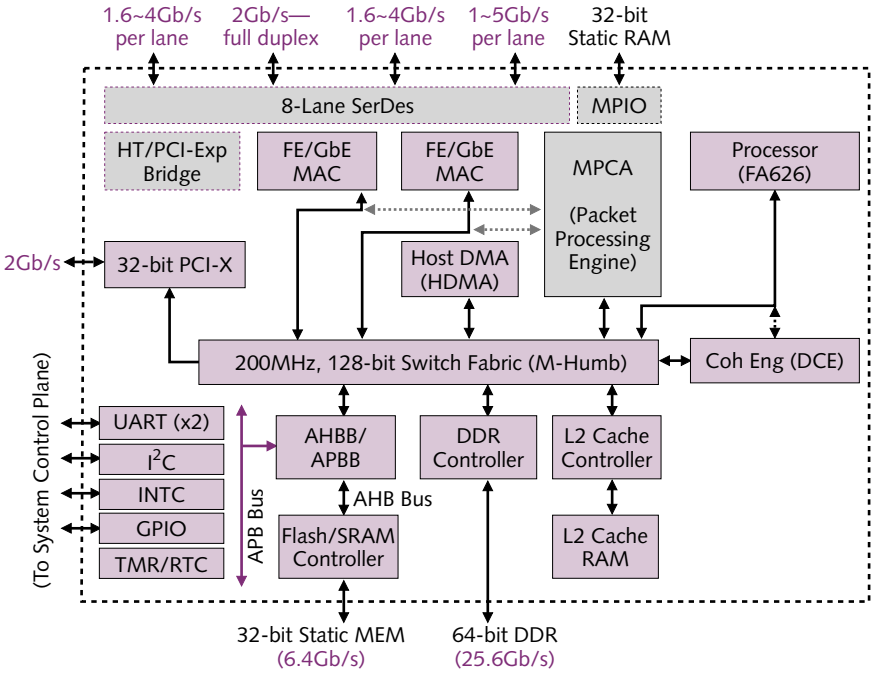

Figure 3. Faraday NetComposer-II block diagram. The FA626 processor is Faraday's own custom-designed ARMv4-compatible core. A nonblocking crossbar switch fabric connects all the bus masters. The programmable eight-lane SerDes can support several different types of I/O protocols and interfaces, including Fibre Channel, Gigabit Ethernet, HyperTransport, PCI Express, RapidIO, Serial ATA, SPI-4.2, Utopia, and XAUI.

ECC protection to the L2 cache. The e600 core will also appear in the new MPC7448. Some future chips from Freescale will have processor cores with more enhancements. (See the sidebar, "New PowerPC Cores Promise Higher Performance," in *MPR 5/10/04-02*, "Freescale Secures PowerQuicc.")

Freescale says the shrink to a 90nm copper silicon-on-insulator (SOI) process will boost the e600's clock frequency beyond 1.5GHz in the MPC8641D. Dual DDR2 memory controllers will operate at bus frequencies up to 667MHz. However, Freescale is coy about actual production clock speeds, perhaps because the chip is about a year away from sampling (2H05). Certainly, 1.5GHz is a conservative target for this design. The fastest e600-based processor now shipping is the 1.42MHz MPC7447A, which is fabricated at 0.13 micron.

Concerns about power consumption and thermal characteristics will probably influence Freescale to restrain the maximum clock speed of the 8641D, as long as it meets performance requirements. The process shrink will drop the core voltage to the 1.0V–1.1V range, maybe lower. Freescale estimates power consumption at 15W–25W (typical) for the dual-core 8641D, and 10W–15W for the single-core 8641. Because the junction temperatures are rated below 105°C, the company is touting 10-year reliability for these parts. Longevity is an important consideration for many embedded systems, whereas almost nobody worries about a desktop PC lasting that long.

Integrating the equivalent of a PowerQuicc III adds lots of features to the 8641D, as the block diagram in Figure 4 shows. In addition to the dual DDR2 memory controllers, there are four GbE controllers, a dedicated PCI Express controller (with one to eight lanes), and a configurable I/O controller that can support either PCI Express (one to eight lanes) or serial RapidIO (one or four lanes). The PCI Express and configurable controllers have their own DMA and link together over a nonblocking crossbar switch called the OCeaN, or On-Chip Network. (See *MPR 12/17/01-01*, "Motorola's MPC8540 Parts OCeaN.") There are also the usual UARTs, I^2C interfaces, timers, and so forth. This basic chip architecture is well proven in other members of the PowerQuicc family, and the 8641D is compatible with PowerQuicc III software.

The four identical GbE controllers have their own FIFO buffers, DMA, and hardware acceleration for L1–L4 packet processing. They offload TCP/IP and User Datagram Protocol (UDP) checksum calculations from the CPUs for both incoming and outgoing traffic, and they support IPv6 addresses in hardware. They can store up to 16 addresses internally for quick lookups, without referring to address tables elsewhere in memory. There are also 16 queues—eight each for incoming and outgoing packets—to support quality-of-service (QoS) filtering. Additional hardware helps accelerate other packet-processing functions, such as firewalling, classification, and tagging for virtual LANs. The FIFOs have their own 8/16-bit I/O interface running at 200MHz, good enough for OC-48 line rates (2.4Gb/s).

Overall, Freescale's first dual-core PowerPC is a logical and thoughtful design. It merges a proven CPU core with a proven networking processor to create a well-integrated chip that can replace two or three separate chips. For customers that can take advantage of its integration, it will cut costs, save power, reduce board space, and double the amount of processing capacity available to programmers. The single-core version offers the same function integration with even greater savings, and it provides a straightforward upgrade path—Freescale will offer both chips in 960-pin, 33×33mm ceramic packages. Moreover, these chips are software compatible with each other and with existing PowerPC MPC744x and PowerQuicc III processors.

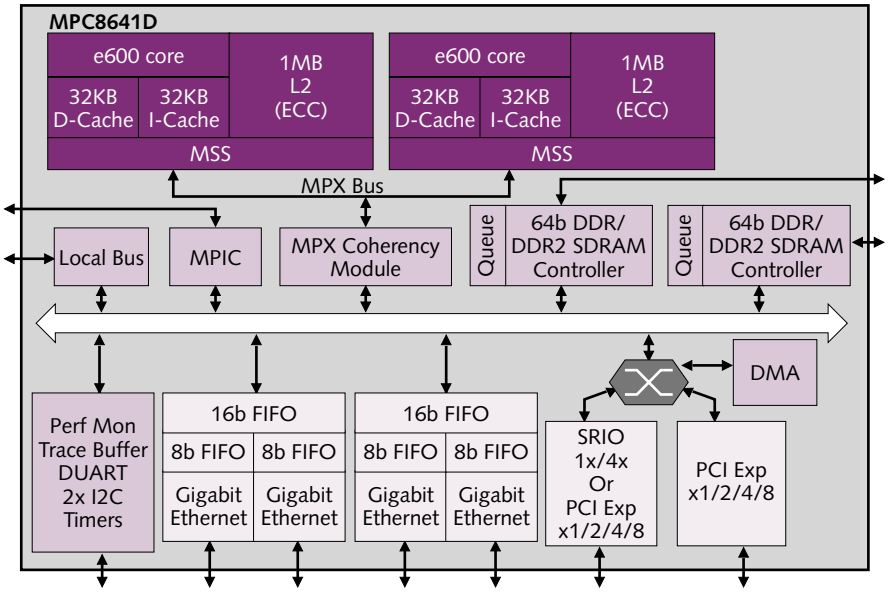

Figure 4. Freescale MPC8641D block diagram. This chip essentially merges two PowerPC MPC7448 processors with the peripherals and interfaces of a PowerQuicc III MPC8548 networking processor. Note the four GbE controllers, dedicated PCI Express controller, and configurable PCI Express/RapidIO controller. Dual DDR2 memory controllers have 64-bit interfaces and will run at speeds up to 667MHz. The PowerPC e600 processor cores are enhanced G4+ cores, with 1MB of L2 cache per CPU. Another newly announced chip, the MPC8641, is a single-core version of this design.

PMC-Sierra Cranks Up the Clock

In terms of clock speed, the fastest embedded processor announced at FPF is PMC-Sierra's 1.8GHz RM11200. It pairs two of PMC-Sierra's new E11K MIPS64-compatible CPU cores with new I/O interfaces and some other integrated features first seen in the RM9200 and RM9220. In addition, the RM11200 will be PMC-Sierra's first chip manufactured in a 90nm process.

By far the most important new features in the RM11200 are the dual E11K CPU cores, PCI Express interfaces, dual DDR2 memory controllers, and enhanced on-chip crossbar bus. The E11K core is a moderate improvement over the MIPS64-compatible E9K core, which extended the two-way superscalar instruction pipelines from five stages to seven stages to reach 1.0GHz. (See *MPR 9/15/03-01*, "PMC Enhances Portfolio.") Those pipelines remain unchanged in the E11K. Better circuit design (in particular, more dynamic logic) and the process shrink account for the RM11200's 80% increase in clock frequency. Instead of tinkering with the pipelines, PMC-Sierra has enlarged the caches. Each E11K core has a 64KB instruction cache (vs. 16KB in the E9K); a 32KB data cache with ECC protection (vs. the parity-protected 16KB data cache in the E9K); and a 512KB L2 cache with ECC (vs. 256KB in the E9K). In addition, all the caches are now eight-way set-associative (vs. four-way in the E9K).

I/O resources in the RM11200 are significantly improved. One new feature is a PCI Express controller, which can operate as a single eight-lane interface or dual four-lane interfaces. There are now four GbE controllers, compared with three GbE interfaces in previous chips. And there are two DDR2 SDRAM controllers on chip, each with 64-bit interfaces. They run at the double data rate of 800MHz (actual bus clock 400MHz), providing a total of 12.8GB/s of memory bandwidth. That sum matches the total memory bandwidth available in the fastest new processors from Broadcom and Cavium, and it exceeds the memory bandwidth provided by the other new embedded processors announced at FPF. In a dual-channel DDR configuration, the RM11200's CPU cores can access different banks of memory simultaneously. In a single-channel configuration, the RM11200 will arbitrate memory accesses between the CPUs.

We think the most fascinating feature of the RM11200 is its enhanced crossbar bus, which connects all the CPU cores and on-chip peripherals together asynchronously. The crossbar uses the Nexus asynchronous switch technology licensed by Fulcrum Microsystems. It mates CPUs and peripherals that run at greatly different clock speeds, and it does so without the multiple tiers of bus bridges found in other designs. The crossbar is a 16-port nonblocking bus that provides 192Gb/s per port, which totals 1.5Tb/s of aggregate bandwidth (16 ports × 192Gb/s per port = 3.072Tb/s, but the crossbar needs one port for input and another for output to make one channel). Latency from port to port is only 3ns and won't slow down under heavy traffic conditions. All transactions are pipelined. There's no global clock regulating this crossbar, so clock skew among the devices attached to the ports isn't a problem. Figure 5 shows a block diagram of the RM11200.

CPU Cores Get Priority

Notice that a processor switch connects the dual CPU cores to the crossbar. This switch enforces data coherency between the L1 and L2 caches of both CPUs and gives the CPUs exclusive access to the crossbar. Either CPU can snoop the other CPU's caches and find the data it needs over a dedicated coherency port. To avoid stalling the pipelines while snooping the L1 caches, each CPU maintains shadow tags for the main tags in its L1 data cache.

Each CPU has its own port into the crossbar switch that provides 200Gb/s of peak bandwidth. The crossbar treats all traffic the same. However, traffic between the CPUs and main memory gets precedence over requests from the I/O controllers, because the memory controller has two different queues. One queue gets higher priority, but the memory controller also guarantees bandwidth to the low-priority queue.

The RM11200 also has the "Direct Deposit" feature found in other multicore processors from PMC-Sierra. This feature allows the GbE, HyperTransport, and PCI Express interfaces to write data directly into either CPU's L2 cache without involving the CPUs. In other words, the CPUs need not explicitly prefetch data over those I/O interfaces; the data just seems to magically appear in the caches. Furthermore, Direct Deposit can automatically strip the packet header and stuff it into the L2 cache while shunting the packet payload in main memory. When the CPU is ready to process the packet, it finds the header and payload already separated, so it can get to work immediately.

Altogether, the RM11200 is a worthwhile improvement over the RM9220. Although it doesn't add any CPU cores, it

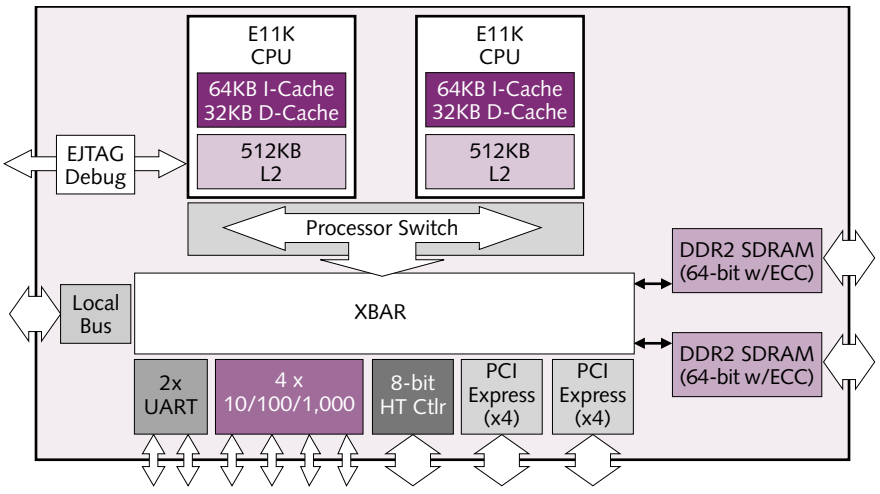

Figure 5. PMC-Sierra RM11200 block diagram. Dual MIPS64-compatible CPU cores with enlarged caches run at 1.8GHz. The PCI Express controller is new in this design, as are the dual-channel DDR2 memory controllers and a fourth GbE controller. Everything ties together over a unique 16-port crossbar that allows the CPU cores and on-chip peripherals to run asynchronously, without the regulation of a global clock signal.

Price & Availability

AMCC's PowerPC 440SPe is scheduled to sample in 1Q05, with production to begin in 3Q05; pricing is unannounced. Broadcom's dual-core SiByte BCM1255 and BCM1280 are scheduled to sample this quarter and ship in 2Q05 for $599 and $699, respectively. Broadcom's quad-core SiByte BCM1455 and BCM1480 are sampling now and scheduled to ship in 3Q05 for $999 and $1,199, respectively. Cavium's Octeon chips are scheduled to sample in 1Q05, with production to start in 2H05; pricing will range from $125 to $750 in 10,000-unit quantities. Faraday's NetComposer-II is scheduled for production in 3Q05 and will cost $48 in 1,000-unit quantities. Freescale's PowerPC MPC8641D and MPC8641 are scheduled to sample in 2H05 and enter production in 1H06; no pricing has been announced. PMC-Sierra's RM11200 is scheduled to sample in 2Q05, with production to begin in late 2005 or early 2006; no pricing has been announced. For more information, visit the vendors' websites:

- AMCC: *www.amcc.com/powerpc/*
- Broadcom: *www.broadcom.com/products/category.php?category_id=27*
- Cavium Networks: *www.cavium.com/octeon.html*
- Faraday Technology: *www.faraday-tech.com*
- Freescale: *www.freescale.com/webapp/sps/site/overview.jsp?code=DRPPCDUALCORE*
- PMC-Sierra: *www.pmc-sierra.com/processors/*

boosts performance by cranking up the clock speed and enhancing the cores. It's the first integrated processor from PMC-Sierra with PCI Express, yet it maintains HyperTransport and adds a fourth GbE controller. The dual-channel DDR2 controllers greatly improve memory bandwidth, and the asynchronous crossbar switch is a creative solution for a common problem in highly integrated processors.

Some Conclusions and Predictions

Without a doubt, Cavium has thrown down the gauntlet with its Octeon family of integrated networking processors. No other high-performance embedded processors can match Cavium's level of integration—not only in CPU cores, but also in hardware acceleration and security features. And Cavium's chips are compatible with MIPS64, probably the most popular CPU architecture for high-end networking and communications. Cavium has set a new bar that's hard to surmount.

That doesn't mean, however, that Cavium will corner the market. On the contrary, despite Cavium's experience with security processors and I/O bridge chips, it's still the newcomer when it comes to broader packet processing. Cavium must deliver what it promises and attract big customers. Over the past five years, we've seen many innovative network processors stymied by indifference in the marketplace. Risk-averse customers often prefer to stick with more-familiar solutions.

Broadcom has a chance to redeem itself by delivering its new dual- and quad-core processors on time and on spec. In raw processing power, Broadcom's quad-core BCM14xx chips are the only competition for Cavium's 16-core monster. Theoretically, both companies' processors can crunch 19.2 billion instructions per second. In real-world applications, Octeon will come closer to reaching its peak performance more often than the BCM14xx will, because 16 two-way superscalar cores offer more opportunities for parallelism than do a quartet of four-way superscalar cores. Nevertheless, Broadcom is flexing serious muscle—if it can deliver its chips on schedule. Any more slippage will turn customers toward the alternatives—and there are an increasing number of alternatives.

PMC-Sierra appears to be bringing up the rear of the MIPS64 parade with processors that have only two CPU cores. Until the most recent spate of announcements, a dual-core embedded processor was impressive; now it's starting to look commonplace. PMC-Sierra says dual-core processors are ideal for the sweet spot of the market and that chips with four or more cores are expensive and unproven in the marketplace. Nevertheless, we believe the company will need at least a quad-core design before too much time passes in order to compete against Broadcom and Cavium for higher-end networking applications. Meanwhile, the dual-core RM11200 is powerful enough to hold down the fort, especially with its fast 1.8GHz core, great memory bandwidth, rich array of I/O controllers, hardware support for cache coherency, and asynchronous crossbar switch. We suspect that, in real systems, the RM11200 will deliver more actual performance than a glance at the datasheet might imply.

With their latest introductions, AMCC and Freescale are making a good case that PowerPC is the rising CPU architecture for network processing. AMCC is the newest supplier for PowerPC chips and is off to a quick start with the 440SPe. However, the 440SPe is basically a tweak of IBM's existing 440SP and was created by a design team AMCC acquired from IBM last spring. AMCC's first post-acquisition original design, which will probably be announced a year from now, will give us a better indication of the company's future direction.

Freescale is another PowerPC vendor with something to prove after its recent spinoff from Motorola: namely, that it is viable as an independent company. If Freescale can survive the trauma of its corporate fission without losing momentum, it has a fleeting opportunity to gain market share for PowerPC against MIPS-based competitors struggling to meet delivery dates. We think the new dual-core MPC8641D is a good start. It's an intelligent amalgamation of two successful existing families of processors, the MPC744x and the PowerQuicc III. It's so logical, it's practically fail-safe. On the other hand, it's not quite edgy enough to capture the imagination. To keep from becoming an also-ran in this market, Freescale will have

to move beyond two CPU cores per chip and probably make the transition to 64 bits as well.

As mentioned before, Faraday is this group's odd man out. While the other vendors press forward with MIPS- and PowerPC-based ASSPs, Faraday is rolling the dice with an ARM-based structured ASIC. It offers an intriguing mix of hard-wired features and programmable logic. The NC-II is less powerful than the other embedded processors unveiled at FPF, but it's more versatile. With intelligent customizing, it can outperform an ASSP in some applications. In particular, we like the programmable SerDes that supports many popular I/O interfaces—a great way to hedge your bets while waiting for the market to settle on standards. If there's a penalty to pay in silicon area and performance, it's not huge, especially when balanced against the much higher costs of spinning a full-custom ASIC. The NC-II fills a growing gap among ASSPs, ASICs, and FPGAs. However, we think the NC-II would be more compelling if it had a more powerful MIPS64 or PowerPC core instead of an ARM core.

In summary, FPF was the showcase for an unusually strong slate of high-performance embedded processors. Rarely are so many directly competitive chips unveiled at a single event. On top of that, *MPR* is aware of additional competition coming in the near future. For customers, it's great to have so many choices. For vendors, it's a struggle for survival. We're looking into our crystal ball and we see emerging from the haze one word: ***attrition.***