

M I C R O P R O C E S S O R

www.MPRonline.com

◆ THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE ◆

NETLOGIC BROADENS XLP FAMILY

Multithreading and Four-Way Issue with One to Eight CPU Cores

By Tom R. Halfhill {7/26/10-01}

NetLogic is unleashing its first barrage of networking and communications processors since acquiring RMI last year. Nine new chips are scheduled to sample this fall, each with the four-way multithreading and four-issue super-scalar features of the previously announced eight-core XLP832. The new chips have one, two, four, or eight CPUs.

The single-core XLP104, XLP204, and XLP304 processors are designed for small-business networking equipment supporting packet-throughput rates of 100Mbps to 4Gbps. For enterprise equipment requiring packet rates of 2Gbps to 40Gbps, NetLogic announced the dual-core XLP208, XLP308, and XLP408, plus the quad-core XLP316 and XLP416. At the high end of the family, the previously announced eight-core XLP832 will be joined by another eight-core chip, the XLP432. These two chips, which are designed for network infrastructure, scale from 10Gbps to 160Gbps. To reach the upper end of that range, systems can link two, three, or four XLP832 processors in a coherent network over an integrated interchip interface.

XLP processors can serve both control- and data-plane applications and are a major expansion of the product line that NetLogic acquired with RMI. The existing XLS and XLR families are still competitive, but the market is moving fast. In recent months, we've seen important announcements and introductions from Cavium, Freescale, and Intel. Freescale and Cavium are NetLogic's most direct competitors, because their respective QorIQ and Octeon II chips have similar RISC architectures and integration. Intel's x86-architecture Xeon LC processors (code-named Jasper Forest) are less integrated, requiring two chips to offer features that the RISC vendors offer in one chip.

Fabrication technology is driving much of the action. Freescale is moving to 45nm production this year. Intel's Xeon LC processors are already in 45nm production and

will probably move to 32nm next year. NetLogic is moving from 90nm to 40nm for XLP production next year, skipping the 65nm node altogether. This transition will push the new chips to 2.0GHz, surpassing Cavium's 1.5GHz Octeon II chips and nearly matching Freescale's new 2.2GHz QorIQ processors.

With four-way multithreading and four-issue super-scalar pipelines, the XLP chips' CPUs have an instruction-throughput advantage over those of Cavium and Freescale at any clock speed. Per-clock performance nearly matches that of Intel's Nehalem CPU. Although some Intel processors surpass NetLogic's 2.0GHz maximum frequency, they also run much hotter.

Cavium, which recently announced an Octeon II with 32 CPUs, has the greatest number of CPUs. Intel's muscular Xeon chips are in production now, well ahead of the newly announced products from Cavium, Freescale, and NetLogic. Nevertheless, the new XLP chips and their improved CPUs will broaden and rejuvenate the product line that NetLogic acquired with RMI.

New CPU Accelerates Control Plane

All XLP processors use a new MIPS64-compatible CPU core that RMI began designing about four years ago, before NetLogic's acquisition announcement in June 2009. As part of the acquisition, the company inherited RMI's architectural license from MIPS, so NetLogic is free to create its own implementations of the 32- and 64-bit MIPS architectures. One result is the 64-bit EC4400 CPU, which implements the MIPS64 Release 2 instruction-set architecture (ISA) and is more powerful than any CPU licensed by MIPS Technologies.

The EC4400 design team, led by VP of Processor Architecture David Hass, set out to boost control-plane (or

single-thread) performance without compromising NetLogic's strong position in data-plane (or total instructions-per-watt) performance. Also with the RMI acquisition, NetLogic gained the XLR and XLS families of networking processors and their customers, which include Alcatel-Lucent, Dell EqualLogic, Fujitsu, Hewlett-Packard, Huawei, IBM, Juniper, McAfee, Motorola, NEC, and ZTE. Any loss of data-plane performance would jeopardize further business with those companies.

Notably missing from RMI's customer list is Cisco. Although Cisco buys network-search engines and physical-layer (PHY) chips from NetLogic, RMI was never able to win significant designs at Cisco with the XLR and XLS. Since the acquisition, however, NetLogic has won a place on Cisco's "preferred vendor" list, so perhaps the new XLP family will open that door.

RMI's previous CPU (also MIPS64 compatible) strongly emphasizes data-plane throughput. It's a single-issue in-order machine that can manage four hardware threads. To enable multithreading, it replicates four complete copies of the CPU's architectural state, including the registers, program counter, flags, interrupts, and other elements. The CPU can switch instruction threads on every clock cycle—there are no time-consuming context switches in which the CPUs must save their architectural states in memory (see [MPR 5/17/05-01](#), "A New MIPS Powerhouse Arrives").

NetLogic refers to each thread context as an nxCPU. Operating systems perceive each thread as a physical CPU. Intel has never implemented more than two threads per CPU using the similar technology it calls Hyper-Threading, although the unreleased Larrabee processor has four threads per CPU. MIPS Technologies began licensing the MIPS32 34K multithreaded CPU core in 2006, but NetLogic's implementation is wholly original (see [MPR 2/27/06-01](#), "MIPS Threads the Needle").

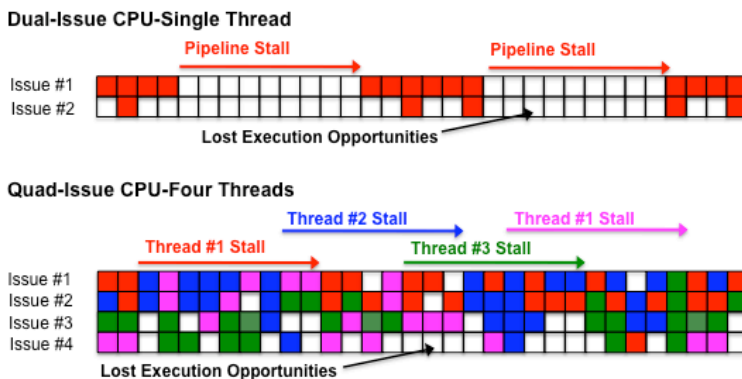


Figure 1. Multithreading and superscalar execution in NetLogic's EC4400. Using a combination of four-way threading and four-way superscalar issue, NetLogic's new CPU can fill many instruction-pipeline slots that would be wasted by data dependencies, cache misses, and other pipeline stalls. Each color represents a different instruction stream.

Boosting Performance through Multithreading

Multiprocessing is especially useful for data-plane processing and other highly parallelizable applications. Multithreading implements multiprocessing at the chip level. Using open-source software like Snort and Clam Antivirus with real packet traffic, NetLogic has measured performance improvements between 50% and 100% when stepping from one thread to four threads. The improvement tends to be greater if the processor's single-thread performance is poor to begin with.

Using only two threads on Intel processors, other benchmarking has shown a 50% throughput increase for IPSec, another packet-processing protocol. For code with less data parallelism, stepping from one thread to two generally gains 5% to 30%, depending on how well the software balances the performance across the two threads.

For single-thread processing, the single-issue pipeline of RMI's original XLR and XLS CPUs suffer a disadvantage compared with the superscalar CPUs from Cavium, Freescale, and Intel. Freescale and Intel also boost single-thread performance by reordering instructions. To address these shortcomings, the new EC4400 adds four-issue superscalar execution and instruction reordering while retaining four-way multithreading.

The ability to issue instructions from different threads during the same clock cycle is known as simultaneous multithreading, or SMT. All threads share four issue slots per cycle, so the CPU's peak throughput is still four instructions per cycle. Nevertheless, multithreading can improve the efficiency of superscalar execution by filling instruction slots and pipeline bubbles with instructions from threads that are ready to go. The processor wastes fewer cycles waiting for stalled instructions to clear.

Figure 1 shows how SMT allows a processor to hide many of the clock cycles lost to data dependencies and other pipeline hazards by using instructions from other threads. When running single-thread code, any thread can issue up to four instructions per cycle. When all four threads are active, the processor acts like four single-issue CPUs. The goal is to achieve the best of both worlds—four-issue superscalar execution maximizes single-thread throughput, and four-way multithreading maximizes total throughput. (For a fuller description of SMT, see [MPR 12/6/99-01](#), "Compaq Chooses SMT for Alpha.")

NetLogic's implementation of SMT allows threads to be turned on or off at runtime under program control. Moreover, it's controllable on a per-CPU basis: each CPU can run one, two, or four threads. This feature lets developers tune the number of threads to the application, even when a multicore chip is serving multiple purposes. For example, in small systems that piggyback control- and data-plane processing on one chip, the operating system can run on a CPU with single or dual

threading while packets flow through CPUs with four-way threading. This flexibility is also useful in applications that run multiple operating systems on the same chip.

Instructions Take Flight

To improve the odds of sustaining peak performance, the EC4400 has a new thread scheduler and five 16-slot instruction queues for the CPU's seven function units. (Some function units share queues.) The thread scheduler determines which four instructions to fetch from which thread, then feeds the instructions into the queues. It's configurable, so it can assign equal shares of clock cycles to each thread or assign different priorities.

The thread scheduler's default priority scheme is an I-count algorithm, which chooses the thread that's using the least number of slots in the queues. This scheme avoids blocking the queues with too many stalled instructions, which typically happens when a particular instruction stream is waiting for a memory access.

Each ALU, along with the FPU, has its own instruction queue. Two ALUs execute simple integer instructions and share their queues with the two load/store units. Another ALU executes simple integer instructions and branch instructions. The fourth ALU executes simple and complex integer instructions (complex operations include multiplies and divides). The FPU executes single- and double-precision floating-point instructions. Figure 2 is a block diagram of the EC4400 core.

To use its function units more efficiently, the EC4400 can issue instructions from the queues out of the original program order—another improvement over the previous CPU. Like some other out-of-order processors, the EC4400 uses register renaming to avoid bottlenecks in the archi-

tectural register file. Pending instructions store their operands in a larger pool of registers; when an instruction completes, the processor renames the associated registers to represent programmer-visible architectural registers.

The EC4400 has 256 registers in the integer pool and 192 registers in the floating-point pool. The processor dynamically maps these registers to the 32 integer and 32 floating-point architectural registers required for each thread. Register renaming leaves ample extra registers for those cases when a newer instruction reuses an architectural register that's still being used by an instruction in flight.

Of course, the processor always retires instructions in their original program order. Everything is put back in place by the reorder buffer. The EC4400's buffer can manage up to 100 instructions at a time, a juggling feat expected more of server processors than of embedded processors. (Intel's Nehalem CPU can manage 128 micro-ops in flight; x86 micro-ops roughly correspond to RISC instructions.)

Note the addition of an FPU to the EC4400—a feature missing from the existing XLR and XLS families. Floating-point math is usually unnecessary for packet processing, but the FPU is useful in applications outside of networking, such as aerospace and office equipment. Cavium's MIPS-compatible chips are intended primarily for data-plane processing and lack FPUs, but Freescale and Intel support floating point in their respective QorIQ and Xeon processors.

Deeper Pipeline, Better Prediction

The EC4400 has deeper 12-stage instruction pipelines, compared with 10 stages in the XLR and XLS processors.

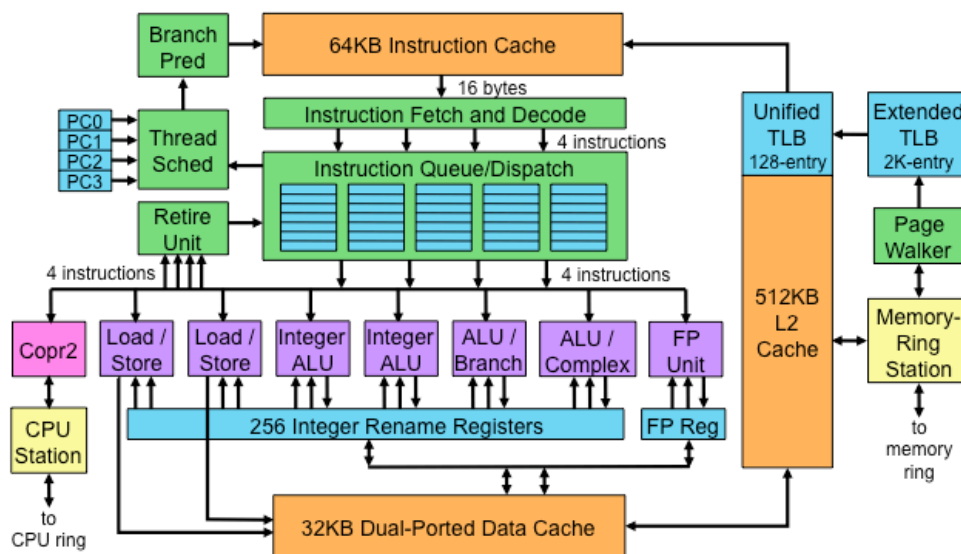


Figure 2. NetLogic EC4400 CPU block diagram. This sophisticated CPU implements simultaneous multithreading, register renaming, and instruction reordering. It can sustain four instructions per cycle, approaching the performance per cycle of Intel's high-end Nehalem CPU. XLP-family processors have one to eight EC4400 CPUs per chip. (Source: NetLogic)

Consequently, XLP-family processors can reach higher clock frequencies in a given fabrication process. Moving production from 90nm to 40nm sweetens the pot. Although TSMC has had trouble achieving good 40nm yields—always a risk with a new process—NetLogic is betting that yields will improve by the time XLP chips enter production. Other TSMC customers with shorter-lifecycle products (like AMD’s ATI graphics processors) will forge the way.

One tradeoff for XLP’s deeper pipelines is that mispredicted branches inflict a greater penalty in wasted clock cycles, because the processor must discard more partially completed instructions when flushing and repriming the pipelines and queues. To avoid this penalty, the EC4400 beefs up the resources devoted to branch prediction. Whereas the previous CPU had a G-share predictor with 4,096 entries, the EC4400 has both G-share and bimodal predictors, each with 8,192 entries. And it has a new 1,024-entry branch-target buffer.

NetLogic also doubled the size of the translation lookaside buffer (TLB) to 128 entries and added an extended TLB with 2,048 entries on the chip. Address lookups that miss the main TLB can skip to the extended TLB, so it’s like having a two-level cache for memory addresses. Lookups that miss the extended TLB can invoke a hardware page walker instead of calling a slower software interrupt handler. Although the main TLB supports both fixed- and variable-size pages, the extended TLB supports only fixed-size pages. Developers can disable the hardware walker and extended TLB under program control, if desired.

To conserve power, each CPU can independently and dynamically vary its clock frequency, and the CPUs can collectively and dynamically vary their core voltage. These adjustments require no software control, and the packet

I/O interfaces are independent of the CPU speed. CPU clock frequencies can drop to one-eighth of the maximum frequency (250MHz in these 2.0GHz implementations), and voltage can vary from 0.7V to 1.1V. The power-management logic complies with Intel’s voltage-regulator module (VRM) specification.

Three-Ring Circus for Packets

An improved on-chip ring network, first disclosed last year with the XLP832, links the EC4400 CPUs to each other and to the memory subsystem, I/O controllers, and acceleration engines. Actually, the chip has three rings, all bidirectional and all running at the core clock frequency. Rings can run at high clock speeds because they tend to have short on-chip trace lengths. One drawback is that data transfers usually need multiple hops to reach their destination node, but a bidirectional ring greatly shortens the worst-case path.

As the block diagram in Figure 3 shows, the main ring is dedicated to CPUs and transfers 64 bytes per CPU per clock cycle. (NetLogic refers to this ring as the Fast Messaging Network, or FMN.) The second ring is dedicated to I/O controllers and accelerators, and the third ring connects the L2 and L3 caches to DRAM. A nonblocking crossbar switch (not shown) connects the rings together.

Separating internal traffic on three rings avoids contention and the bus saturation that limits throughput when using traditional multidrop buses. The I/O ring segregates the traffic between I/O controllers and memory, so I/O reads and writes don’t interfere with inter-CPU communications on the CPU ring. Memory transfers between the L2 cache, L3 cache, and DRAM never traverse the crossbar switch, avoiding that bottleneck as well. NetLogic says the internal network provides a total of 40Tbps of bandwidth.

NetLogic has improved the XLP processors’ CPU load balancing. The XLR and XLS families distribute packet processing among the CPUs according to a scheme implemented in hard-wired logic, which isn’t very flexible. For the new chip, NetLogic designed its own MIPS32-compatible CPU core to accelerate common networking tasks, including load balancing. Each XLP chip has more than 40 of these small cores.

NetLogic says the hardware load balancers had to invoke numerous exceptions to parse and inspect the special packet headers that some customers use. The new balancers are programmable in C using standard development tools, so they can handle any header. NetLogic provides firmware for the common protocols and can modify the firmware if necessary. Alternatively, customers can

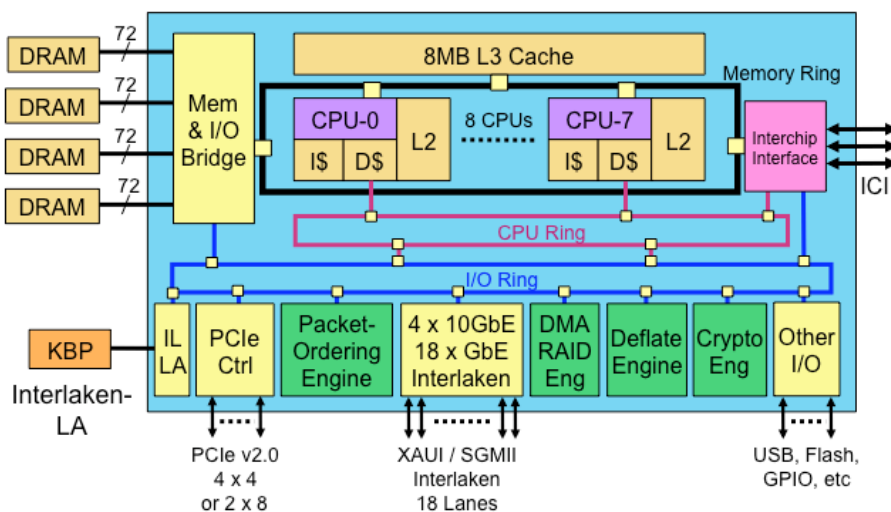


Figure 3. NetLogic XLP832 block diagram. This design has eight EC4400 CPU cores, but the chip architecture is essentially the same for the smaller designs with one, two, or four CPUs. Three internal rings connect CPUs to I/O controllers and memory. (Source: NetLogic)

do the work themselves, which is useful if they want to protect a proprietary header format.

The load balancers can assign newly arrived packets to CPUs on a round-robin basis or by routing traffic to the CPU with the lightest load. Packets can be resequenced, reassigned to different threads, bounced among different CPUs, forwarded to the security engine, and finally sent toward their destinations in their original order. Packet routing through the on-chip network requires no intervention by the CPUs, which are free to concentrate on processing.

Another custom-designed MIPS32-compatible processor core drives the cryptography engine, making it easier for programmers to support new encryption modes. The regular-expression (reg-ex) engine uses NetLogic's own proprietary NETL7 cores instead of MIPS cores, but they have a similar programming model. This engine has some local memory and can access external memory without requiring a dedicated reg-ex memory interface. (Note that the XLP4xx and XLP832 chips lack a reg-ex engine; they require a NETL7 external coprocessor.) The packet-ordering engines in the XLP chips continue to rely heavily on dedicated logic, not on the fully programmable processors of the load balancers, but they are configurable.

An Exceptional Embedded CPU

All together, the EC4400's improvements—particularly its four-way superscalar issue and out-of-order execution—make this a powerful CPU design. The EC4400 retains the four-way multithreading of its predecessor and enhances it with SMT, going well beyond the capabilities of other embedded processors.

Indeed, the EC4400 will challenge the instructions-per-cycle prowess of Intel's Nehalem CPU. Both have four-issue superscalar execution and instruction reordering, but the EC4400 has four-way multithreading, beating Nehalem's two-way Hyper-Threading. Both CPUs can juggle similar numbers of instructions in flight (100 RISC instructions versus 128 RISC-like micro-ops). Nehalem's 16-stage pipeline is deeper than the EC4400's 12-stage pipeline, but that's not necessarily an advantage. Intel's x86 processors need more stages to decode their complex instructions and break them up into micro-ops, whereas the MIPS-compatible EC4400 dines on simpler RISC instructions to start with. On paper, at least, the EC4400 is a good match for Intel's best CPU.

NetLogic estimates that Intel's Xeon still has a 5% to 15% advantage over the EC4400 in high-end control-plane applications but asserts that XLP chips will lead the field in data-plane processing. Intel has better branch prediction and faster load/store units—features that NetLogic chose not to match, instead favoring power conservation. NetLogic says it's willing to concede the extremes of the control-plane performance spectrum to address the middle 80% of the networking and communications market.

Compared with Freescale's new Power e5500 CPU, the EC4400 can execute two additional instructions per cycle and has greater reordering capabilities. The single-threaded Power e5500 will also suffer more pipeline bubbles, although its shorter seven-stage pipeline makes recovering from them less painful. Cavium's cnMIPS64 CPUs are easily outgunned by the EC4400. The Cavium CPUs are even more limited than the Power e5500 by their in-order processing. Clock for clock, we expect the XLP CPU to outperform the Power e5500 and cnMIP64 CPUs by 30% to 50% on single-threaded code.

Despite the new emphasis on control-plane processing, NetLogic was loath to surrender any ground in the data plane. By retaining four threads per CPU, the EC4400 preserves the programming model of its predecessor. In fact, the new CPU benefits more from multithreading than the old CPU did, because superscalar execution yields sharply diminishing returns beyond two-way issue. A four-issue CPU leaves more instruction slots unfilled most of the time. SMT (and instruction reordering) allows the EC4400 to usefully fill most of these unused slots. The CPU can devote up to four slots per cycle to a single thread or one slot per cycle to four simultaneous threads. Because neither the Freescale nor Cavium CPUs implement multithreading, the EC4400 can achieve two to three times the per-clock throughput of these competitors on highly parallel code.

Cavium favors physical CPUs over virtual CPUs. Although Cavium's CPUs are single threaded, the company's Octeon II processors will have as many as 32 CPUs per chip. NetLogic is fighting back with the eight-core XLP432 and XLP832, whose four threads per CPU add up to 32 nxCPUs. We think these chips can deliver similar data-plane performance, and a power/performance comparison tilts the balance toward NetLogic. Whereas Cavium says its CN6880 will consume 65W maximum at 1.5GHz, NetLogic says the 2.0GHz XLP832 won't exceed 50W (see [MPR 5/31/10-02](#), "Cavium Pushes Octeon to 32 CPUs").

Smaller Chips Expand XLP Family

To reach the broadest range of the networking market, NetLogic needs processors smaller than the eight-core XLP832 announced last year. The latest announcements significantly broaden the XLP family to include chips having one, two, or four CPUs, plus another eight-core chip to supplement the XLP832. In addition to varying the number of CPUs, these chips have different amounts of L2/L3 cache, different I/O capabilities, and different levels of reg-ex and cryptography acceleration.

Table 1 summarizes the 10-chip XLP family as it stands today. NetLogic's nomenclature is relatively simple: the first digit signifies relative features, power, and cost; the last two digits indicate the number of nxCPUs (threads). Dividing the last two digits by four yields the number of physical CPUs per chip.

NetLogic's overall strategy is to scale memory bandwidth, I/O, and acceleration throughput to match the number of CPUs while preserving pin compatibility within each series. This strategy avoids performance bottlenecks as the number of CPUs increases. It also lets customers choose a particular performance level for their first design while leaving an upgrade path for future designs. For instance, a developer could start by estimating the maximum number of CPUs the application will ever need, then choose a series with that number of CPUs at the high end. The developer's initial design could start with a smaller number of CPUs and later upgrade to higher-performance chips within that series, as needed. The additional re-

sources on those chips will provide balanced performance, while pin compatibility averts a costly board redesign.

Starting at the low end, the XLP1xx series has only one member, the XLP104. It's a single-core chip suitable for mixed control- and data-plane chores in small-business networking equipment. It can manage packet-throughput rates up to 2Gbps and doesn't skimp on reg-ex or crypto acceleration. (NetLogic says its data-throughput estimates are based on large packets, so the processor may not be able to sustain these rates for minimum or even average packet sizes.) With two serial ATA (SATA) interfaces, the XLP104 can serve as a controller for network-attached storage (NAS). Integrated PHY interfaces for the USB 2.0

controllers will save a little money and board space by replacing external PHY chips.

NetLogic hasn't announced pricing for any XLP chips, but we estimate that some versions of the XLP104 will cost about the same as the lowest-priced member of the XLS family, which has a list price of \$40–\$50. By moving from 90nm fabrication to 40nm, the cost of manufacturing a single-CPU chip like the XLP104 is actually lower than the manufacturing cost of a single-CPU chip in the XLS family.

XLP2xx/X3xx Offer More CPUs

The XLP2xx series includes the single-core XLP204 and dual-core XLP208. These chips bump the effective clock rate of the DDR3 DRAM interface to 1.6GHz (from 1.33GHz in the XLP104), enlarge the L3 cache to 2MB, and add 10G Ethernet, four serial RapidIO controllers, two more USB interfaces, two more SATA controllers, faster reg-ex processing, and faster cryptography acceleration. In total, the XLP2xx chips have 20 high-speed serial (serdes) lanes, versus 11 lanes for the XLP104.

The XLP2xx chips also have larger packages and consume more power—up to 50% more than the XLP104. (The additional I/O capabilities appear to consume about 20% more power, and adding a second CPU bumps power consumption another 30%.) To replace existing chips like the XLS408, we estimate that the XLP2xx chips must be priced at \$125 to \$175.

The XLP3xx series comprises three chips: the single-core XLP304, dual-core XLP308, and quad-core XLP316. All

	NetLogic XLP1xx	NetLogic XLP2xx	NetLogic XLP3xx	NetLogic XLP4xx	NetLogic XLP8xx
Processor Product Numbers	XLP104	XLP204 XLP208	XLP304 XLP308 XLP316	XLP408 XLP416 XLP432	XLP832
CPUs Per Chip	1 CPU	1 or 2 CPUs	1, 2, or 4 CPUs	2, 4, or 8 CPUs	8 CPUs
L2 Cache w/ ECC	512KB	512KB per CPU	512KB per CPU	512KB per CPU	512KB per CPU
L3 Cache w/ ECC	512KB	204: 1MB 208: 2MB	304: 1MB 308: 2MB 316: 4MB	408: 2MB 416: 4MB 432: 8MB	8MB
Memory Controller w/ ECC	DDR3 1 x 64-bit 1.33GHz	DDR3 1 x 64-bit 1.6GHz	DDR3 2 x 64-bit 1.6GHz	DDR3 4 x 64-bit 1.6GHz	DDR3 4 x 64-bit 1.6GHz
Memory B/W	10.6GB/s	12.8GB/s	25.6GB/s	51.2GB/s	51.2GB/s
PCI Express	4 x PCIe	4 x PCIe	8 x PCIe	8 x PCIe	8 x PCIe
Ethernet Controllers	5 x GbE	1 x 10GbE or 8x GbE	2 x 10GbE or 8 x GbE	4 x 10GbE or 18 x GbE	4 x 10GbE or 18 x GbE
Serdes	11 lanes	20 lanes	24 lanes	32 lanes	50 lanes
Ser RapidIO	—	4 x sRIO	4 x sRIO	—	—
USB 2.0	2 with PHYs	4 with PHYs	4 with PHYs	4 with PHYs	4 with PHYs
Serial ATA	2 x SATA	4 x SATA	4 x SATA	—	—
Chip to Chip Interface	—	—	—	—	3 x ICI coherent
Interlaken-LA	—	—	—	50Gbps	50Gbps
RAID Engine	RAID5/6	RAID5/6	RAID5/6	RAID5/6	RAID5/6
RegEx Engine	1.0Gbps	204: 1.25Gbps 208: 2.5Gbps	10Gbps	—	—
Crypto Accel	1.0Gbps	204: 1.5Gbps 208: 3.0Gbps	304: 2.5Gbps 308: 5.0Gbps 316: 10Gbps	408: 10Gbps 416: 20Gbps 432: 40Gbps	40Gbps
Power (max)	2.5W–10W (500MHz–2.0GHz)	204: 12W 208: 15W	304: 20W 308: 25W 316: 30W	408: 25W 416: 35W 432: 50W	50W
Package	BGA ~600 pins	BGA ~1,000 pins	BGA ~1,000 pins	BGA ~2,000 pins	BGA ~2,000 pins
Samples	4Q10 (est)	4Q10 (est)	4Q10 (est)	3Q10 (est)	3Q10 (est)
Production	3Q11 (est)	3Q11 (est)	3Q11 (est)	2Q11 (est)	2Q11 (est)

Table 1. Key parameters for NetLogic's XLP family. All announced members of this family share the same basic characteristics: the new MIPS64-compatible EC4400 CPU core, 12-stage instruction pipelines, four-way superscalar execution per CPU, out-of-order execution, four threads per CPU, a 64KB L1 instruction cache per CPU, a 32KB L1 data cache per CPU, 40nm G fabrication, 2.0GHz maximum clock speed, and a dynamically controlled voltage range of 0.7V to 1.1V (core). All of the chips also have an assortment of miscellaneous I/O interfaces: GPIO, flash memory, UARTs, I²C, SPI, and SDIO. NetLogic hasn't announced pricing. (Source: NetLogic)

three offer a big jump in external memory bandwidth by integrating two DDR3 DRAM controllers, each 64 bits wide, versus one such controller in the XLP1xx and XLP2xx series. In aggregate, that's 25.6GB/s of memory bandwidth, compared with 12.8GB/s for the XLP2xx series and 10.6GB/s for the XLP104. The additional memory bandwidth is necessary to avoid bottlenecks in the quad-core XLP316.

In addition, the XLP3xx chips have much more I/O and faster acceleration logic while retaining a relatively small package with about 1,000 pins—approximately the same number of pins as the XLP2xx series. The XLP3xx chips double the number of PCI Express (PCIe) controllers, offer the option of two 10G Ethernet interfaces, increase the throughput of the reg-ex engine, and improve the performance of the cryptography accelerator.

In other respects, the I/O capabilities of the XLP3xx series match those of the XLP2xx series. Depending on the number of cores, the XLP3xx chips are suitable for networking equipment in small businesses and some larger installations. Their serial RapidIO interfaces also make them suitable for some cellular base-station designs. To compete with Freescale's single-core QorIQ P5010 and dual-core P5020, we estimate that the XLP3xx chips will cost \$150 to \$300.

Up to Eight CPUs in XLP4xx and XLP8xx

The XLP4xx series includes three chips: the dual-core XLP408, quad-core XLP416, and eight-core XLP432. These are higher-end devices optimized almost exclusively for packet processing. Consequently, they're Ethernet-rich, offering the option of four 10G Ethernet interfaces or 18 Gigabit Ethernet interfaces. They have the same PCIe and USB controllers as the XLP2xx/3xx series and twice as many DDR3 controllers as the XLP3xx series, but they drop serial RapidIO, SATA, and the reg-ex engine—features that are more useful for cellular, storage, and security applications.

On the other hand, the XLP4xx chips add an Interlaken-LA interface and improve cryptography throughput even further (to as much as 40Gbps). Interlaken-LA is a chip-to-chip networking interface originally developed by Cisco and Cortina. It supports data rates of up to 100Gbps, but the XLP implementation peaks at 50Gbps (eight lanes running at 6.25Gbps). Maximum power consumption soars as high as 50W for the eight-core XLP432; the dual-core XLP408 has a cooler 25W maximum. We estimate pricing for the XLP4xx chips at \$300 to \$600, depending on the number of CPUs.

Ruling the XLP-family roost is the XLP832, the only member (so far) of the XLP8xx series. In every respect save one, it's identical to the eight-core XLP432, but this difference is crucial—it's the only XLP processor with NetLogic's proprietary interchip interface (ICI). Actually, it has three such interfaces, allowing developers to connect

two, three, or four of these chips together in a glueless cache-coherent network, as Figure 4 shows.

NetLogic's ICI is another answer to Cavium's superior multicore integration. NetLogic believes that one XLP832 with 32 nxCPUs (threads) will outperform a 32-CPU Octeon II CN6880 while burning significantly less power. If that assertion is optimistic, or if an application demands even greater performance, NetLogic's ICI can link two, three, or four XLP832 chips together. Multiple Octeon II chips can only be linked in a pipelined fashion, complicating the programming model. Applications needing this much performance include core routing, intrusion prevention, and 4G-cellular core infrastructure.

When announced last year, the XLP832 was scheduled to sample in 4Q09 and begin production in 4Q10. The complex design has taken longer than expected, however. Sampling is now expected in 3Q10 and production in 2Q11. We estimate that the XLP832 will cost about \$750, like top-of-the-line members of the XLR family.

All members of the XLP4xx series are on the same revised schedule, implying that (at least initially) they use the same die as the eight-core models. The XLP104, XLP2xx, and XLP3xx are one quarter behind this schedule: sampling is expected in 4Q10, and production is expected in 3Q11. If NetLogic can meet those deadlines, the new XLP chips will arrive at nearly the same time as Cavium's Octeon II and Freescale's new QorIQ P3 and P5 series (see [MPR 7/5/10-01](#), "Freescale P5 Raises QorIQ's I.Q.").

XLP308 is Brainy and Speedy

Most control-plane software is single threaded, so these applications typically use single- or dual-CPU processors. In these applications, NetLogic's XLP308 will compete with Freescale's new QorIQ P5020, Intel's new Xeon LC3528 (Jasper Forest), and—to a lesser extent—Cavium's new Octeon II CN6320. Table 2 compares these four processors.

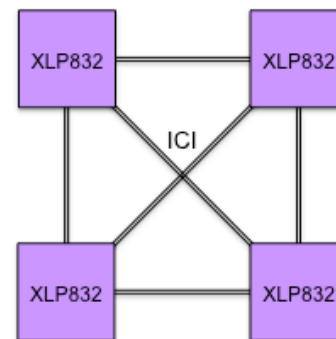


Figure 4. NetLogic's interchip interface (ICI). This proprietary interface gluelessly connects up to four XLP832 chips leaving only one hop between any two chips. It allows nonunified memory access (NUMA) and preserves coherence for memory and caches. Packets can be processed by CPUs in any chip. The chips' internal ring networks can pass messages and interrupts to local nodes or remote nodes.

Thanks to four-way superscalar issue, instruction re-ordering, and four-way multithreading, the dual-core XLP308 can issue up to eight instructions per clock cycle and can fill many otherwise unusable instruction slots. It comes close to the per-clock performance of Intel's LC3528 and easily beats Freescale's P5020 and Cavium's CN6320. Cavium's processor is further hampered by in-order execution—the only such machine in this group.

Throughput is the product of instructions per cycle and clock speed, and the XLP308 excels in the second parameter, too. In this group, it's practically a speed demon. At 2.0GHz, it's second only to the 2.2GHz P5020. Yet, the P5020 has a much shorter pipeline (7 stages versus 12

stages). Normally, deeper pipelines permit higher clock speeds. The P5020's swiftness is probably a benefit of superior SOI fabrication and Power e5500's simpler micro-architecture.

Intel's LC3528 finishes third in the clock-speed derby, managing a sustainable 1.73GHz. (It can briefly surge to 1.87GHz in Turbo Mode. Other Jasper Forest chips are faster but burn much more power.) Cavium's 1.5GHz CN6320 is hampered by 65nm fabrication, the oldest process technology in this group. Overall, the XLP and Xeon processors have similar single-threaded CPU performance and, by our estimate, a slight edge over the P5020. Cavium trails the pack—the CN6320 will deliver only about half the throughput of the others.

Things could change, however, by the time the XLP reaches market. Cavium, Freescale, and NetLogic are all on similar schedules, with their newest chips slated to debut in 2H11. Intel's LC3528 is available now and isn't manufactured in the company's latest 32nm process. By 2H11, Intel will probably have a 32nm product that's faster at the same power level compared with the current product. As AMD has learned in the PC market, it's hard to keep up with Intel's rapidly moving fabrication technology.

The Memory Bandwidth Champ

The XLP308 has more memory bandwidth than its competitors. With two 64-bit DDR3 interfaces running at an effective clock frequency of 1.6GHz, it has a total of 25.6GB/s. This bandwidth helps make up for a smaller cache subsystem compared with Intel (3MB versus 4.5MB), although it's larger than Cavium's 2MB.

For packet processing and other I/O-intensive tasks, the XLP has a clear advantage. Its primary network I/O interface can be configured as two 10G Ethernet ports or eight Gigabit Ethernet ports. Cavium and Freescale offer only one 10G interface, and Intel offers none. PCIe and serial RapidIO are a similar story—the XLP308 has twice as many such interfaces as the best of its rivals. Although Intel offers more USB ports (12)—a legacy of Xeon's PC heritage—they require a second chip, the BD3420 platform controller hub.

Reg-ex engines and cryptography accelerators are useful mainly for packet

	NetLogic XLP XLP308	Freescale QorIQ P5020	Intel Xeon LC3528	Cavium Oceon II CN6320
CPU Type	EC4400	Power e5500	x86 Core i7	cnMIPS64-R2
Arch. Width	64 bits	64 bits	64 bits	64 bits
Pipeline	12 stages	7 stages	16 stages	9 stages
Issue Rate	4 per cycle Out of order	2 per cycle Out of order	4 per cycle Out of order	2 per cycle In order
CPUs Per Chip	2 CPUs (8 threads)	2 CPUs (2 threads)	2 CPUs (4 threads)	2 CPUs (2 threads)
CPU Freq (max)	2.0GHz	2.2GHz	1.73GHz	1.5GHz
L1 Cache (I / D)	64KB / 32KB per CPU	32KB / 32KB per CPU	32KB / 32KB per CPU	37KB / 32KB per CPU
L2 Cache w/ ECC	512KB per CPU	512KB per CPU	256KB per CPU	2MB shared
L3 Cache w/ ECC	2MB shared	2MB shared	4MB shared	—
Memory Controller w/ ECC	DDR3 2 x 64 bits 1.6GHz	DDR3 / 3L 2 x 32 / 64 bits 1.3GHz	DDR3 2 x 64 bits 1.06GHz	DDR3 1 x 64 bits 1.6GHz
Memory B/W	25.6GB/s	20.8GB/s	16.9GB/s	12.8GB/s
PCI Express	8 x PCIe	4 x PCIe	4 x PCIe	2 x PCIe
Ethernet Controllers	2 x 10GbE or 8 x GbE	5 x GbE 1 x 10GbE	1 x GbE*	4 x GbE or 1 x 10GbE
Serdes	24 lanes	18 lanes	16 lanes	12 lanes
Serial RapidIO	4 x sRIO	2 x sRIO	—	1 x sRIO
USB 2.0	4 with PHYs	2 with PHYs	12*, no PHYs	2 with PHYs
Serial ATA	4 x SATA	2 x SATA	6 x SATA*	—
RAID5/6 Eng?	Yes	Yes	Yes	Yes
Crypto Accel?	Yes	Yes	—	Yes
RegEx Engine?	Yes	Yes	—	Yes
IC Process	40nm G	45nm SOI	45nm high-k	65nm G
Voltage (CPU)	Up to 1.1V	1.1V	Xeon: 1.99V† I/O hub: 1.05V	Not disclosed
Power	25W max	30W max	Xeon: 35W TDP I/O hub: 5W	17W max
Package(s)	BGA ~1,000 pins	FC-PBGA 1,023 pins 37.5mm	Xeon: FC-LGA 1,366 pins 42.5 x 45mm I/O hub: FC-BGA 951 pins, 27mm	FC-BGA 900 pins 31mm
Production	3Q11 (est)	2H11 (est)	2Q10	2H11 (est)
List Price (1,000s)	\$225†	\$225†	Xeon: \$302 I/O hub: \$31	\$59†

Table 2. Comparing NetLogic's XLP308 with competitors. All these dual-core processors are suitable for control-plane applications and systems that perform control- and data-plane processing on a single chip. *Requires Intel's BD3420 platform controller hub. (Source: vendors, except †The Linley Group estimate)

processing, and the XLP308 holds its own in this regard. Intel's LC3528 lacks these features, because it concentrates on control-plane processing. The x86 has instructions to assist string comparisons and is adding instructions to accelerate cryptography.

The usual price for higher performance and more I/O is higher power consumption. Yet, the XLP308 appears to dodge this bullet, if NetLogic's estimates are close to accurate. NetLogic specifies 25W maximum for the XLP308, which undercuts Freescale's estimate of 30W maximum for the P5020. Intel's thermal design power (TDP) for its two-chip solution is 40W. Cavium estimates only 17W for the CN6320, but at 1.5GHz, it's the lowest-clocked processor in this group. Our analysis concludes that the XLP308 has the best performance per watt.

Overall, the XLP308 looks like the best processor in this group. It has a small edge over Freescale's P5020, thanks to slightly higher throughput and slightly lower power estimates. The XLP308 scales to higher bandwidth for memory and Ethernet. It has clear advantages over Intel's LC3528 in chip count, power, and probably price, although Intel has a year to reduce those gaps by moving production to 32nm. Cavium's CN6320 is inexpensive but not well suited to high-performance control-plane tasks.

NetLogic Excels in Data Plane

Table 3 focuses on data-plane processing. At first glance, our comparison looks unfair, because it pits NetLogic's dual-core XLP208 against two quad-core processors from Freescale and a six-CPU chip from Cavium. NetLogic invites this comparison. Multithreading boosts the instructions-per-cycle efficiency of the XLP208, and it has a 33% clock-speed advantage to boot. Table 3 compares the XLP208 with Cavium's Octeon II CN6335 and Freescale's QorIQ P3041 and P4040. Six CPUs help Cavium compensate for narrower superscalar issue, in-order instruction execution, and single threading. (Alternatively, consider the quad-core CN6330, which is otherwise identical to the CN6335.) Freescale has no six-CPU processors—the next step up from four CPUs is the eight-core P4080.

Despite having half as many cores as the Freescale chips, the XLP208 matches their peak number of instructions per cycle (eight). CN6335's six CPUs can issue 12 instructions per

cycle, but peak throughput is rare for any processor; average throughput is more important. NetLogic estimates that four-way multithreading improves average instruction throughput by 50% to 100%, especially for data-plane code.

Clock speed is clearly in NetLogic's favor. NetLogic says the XLP208 will reach 2.0GHz—33% faster than the other chips in this group. Multiplying this advantage by the instructions-per-cycle advantage suggests that each XLP CPU has two to four times the performance of the competing CPUs. Unless NetLogic's thread utilization is poor, the XLP208 should easily beat Freescale's P3041 and P4040 and match or exceed the data-plane performance of Cavium's CN6335.

The CN6335 understandably lags at 1.5GHz, because Cavium is using 65nm technology. But Freescale's P3041 and P4040 hit the ceiling at the same clock frequency, despite 45nm SOI. As a rule of thumb, SOI is 10–20% more efficient than bulk CMOS. The Power e500mc CPU's

	Netlogic XLP XLP208	Cavium Octeon II CN6335	Freescale QorIQ P3041	Freescale QorIQ P4040
CPU Type	EC4400	cnMIPS64-R2	Power e500mc	Power e500mc
Arch. Width	64 bits	64 bits	32 bits	32 bits
Pipeline	12 stages	9 stages	7 stages	7 stages
Issue Rate	4 per cycle Out of order	2 per cycle In order	2 per cycle Out of order	2 per cycle Out of order
CPUs	2 CPUs	6 CPUs	4 CPUs	4 CPUs
Per Chip	(8 threads)	(6 threads)	(4 threads)	(4 threads)
CPU Freq (max)	2.0GHz	1.5GHz	1.5GHz	1.5GHz
L1 Cache (I / D)	64KB / 32KB per CPU	37KB / 32KB per CPU	32KB / 32KB per CPU	32KB / 32KB per CPU
L2 Cache w/ ECC	512KB per CPU	2MB shared	128KB per CPU	128KB per CPU
L3 Cache w/ ECC	2MB shared	—	1MB shared	2MB shared
Memory Controller w/ ECC	DDR3 1 x 64-bit 1.6GHz	DDR3 1 x 64 bits 1.6GHz	DDR3 / 3L 1 x 32 / 64 bits 1.3GHz	DDR3 / 3L 2 x 32 / 64 bits 1.3GHz
Memory B/W	12.8GB/s	12.8GB/s	10.4GB/s	20.8GB/s
PCI Express	4 x PCIe	2 x PCIe	4 x PCIe	3 x PCIe
Ethernet Controllers	1 x 10GbE or 8 x GbE	1 x 10GbE or 4 x GbE	1 x 10GbE or 5 x GbE	2 x 10GbE or 8 x GbE
Serdes	20 lanes	12 lanes	18 lanes	16 lanes
Serial RapidIO	4 x sRIO	1 x sRIO	2 x sRIO	2 x sRIO
USB 2.0	4 with PHYs	2 with PHYs	2 with PHYs	2, no PHYs
Serial ATA	4 x SATA	—	2 x SATA	—
RAID Engine	RAID5/6	RAID5/6	RAID5	RAID5
RegEx Engine?	Yes	Yes	Yes	Yes
Crypto Accel?	Yes	Yes	Yes	Yes
IC Process	40nm G	65nm G	45nm SOI	45nm SOI
Voltage (CPU)	Up to 1.1V	Not disclosed	1.0V	1.0–1.1V
Power (max)	15W	17W*	15W	21W
Package	BGA ~1,000 pins	FC-BGA 900 pins	FC-PBGA 1,295 pins	FC-BGA 1,295 pins
Samples	4Q10 (est)	2Q10	4Q10 (est)	4Q09
Production	3Q11 (est)	2H11 (est)	2H11 (est)	3Q10

Table 3. Comparing NetLogic's XLP208 with competitors. These processors have two, four, or six cores and are suitable for data-plane applications, although they can also shoulder control-plane duties in smaller systems. (Source: vendors, except *The Linley Group estimate)

Price and Availability

NetLogic's XLP family now includes 10 networking and communications processors, counting the previously announced XLP832. The new XLP104, XLP2xx, and XLP3xx processors are scheduled to sample in 4Q10 and begin production in 3Q11. The XLP4xx and XLP832 processors are scheduled to sample in 3Q10 and begin production in 2Q11. NetLogic hasn't announced pricing. The Linley Group estimates pricing at \$40 to \$50 for the XLP104, \$125 to \$175 for the XLP2xx series, \$150 to \$300 for the XLP3xx series, \$300 to \$600 for the XLP4xx series, and \$750 for the XLP832. More information is available at www.netlogicmicro.com/Products/MultiCore/XLP.htm.

shorter seven-stage pipeline is one limitation, although an equally short pipeline doesn't seem to restrain the Power e5500 CPU in Freescale's P5020. Perhaps Freescale is trading some SOI efficiency for lower power consumption instead of higher clock speed. The P3041 is rated at 15W (maximum), the same as the XLP208.

Remarkably, the XLP208 has the best cache subsystem in this group, despite having fewer CPUs to feed. It has a total of 3MB of cache, versus the CN6335's 2MB, the P3041's 1.5MB, and the P4040's 2.5MB. Keep in mind, however, that NetLogic's threads are virtual CPUs that must share the L1 and L2 caches attached to the physical CPUs, as well as the chip's global L3 cache. In the best case, each thread runs the same task on a different packet, so they all fetch the same batch of instructions, which avoids thrashing the cache. In the worst case, each thread runs a different task on a different packet and gets only one-fourth of the cache. The XLP's L1 instruction cache is twice as large as the others, which helps a little, but it may need to be four times larger, and the L1 data cache is still only 32KB. These shortcomings could erode the XLP208's performance in applications with large data sets.

With one 64-bit DDR3 controller running at an effective clock rate of 1.6GHz, the XLP208 has 12.8GB/s of aggregate bandwidth to main memory. That's as much as the CN6335, which has three times as many CPUs, but the XLP208 needs that much bandwidth to have any hope of sustaining its potential throughput. Freescale's P4040 is better provisioned in this respect, having two memory controllers and a total bandwidth of 20.8GB/s. Also, Freescale's controllers are designed to work with DDR3L DRAM in addition to standard DDR3 DRAM, allowing customers to choose lower power or higher performance.

XLP208 Has Superior I/O

All the chips in Table 3 have one 10G Ethernet controller, except for Freescale's P4040, which has two. (As mentioned above, the P4040 is basically a P4080 with half as

many CPUs.) The XLP208 matches the P4040's number of Gigabit Ethernet interfaces; the other chips have fewer. For most applications in this power range (21W or less), the P4040's extra packet interface is unnecessary.

For other I/O, however, NetLogic has the advantage. The XLP208 matches its best rival for PCIe and has twice as many serial RapidIO, SATA, and USB 2.0 interfaces as its next-best competitor. Cavium's CN6335 and Freescale's P4040 are the only chips in this group without SATA, so they're less suitable for applications requiring external storage and for security appliances with hard drives. (External SATA chips are cheap but use a PCIe port.) In addition, the XLP208 has hardware acceleration for RAID5 and RAID6; the only other chip in this group with RAID6 acceleration is the SATA-less CN6335. All these chips have hardware accelerators for reg-ex processing and cryptography, making them well suited to security applications.

Freescale's P4040 will enter production any day now, whereas the other chips are a year away from production. The only vendor to announce pricing is Cavium: the CN6335 will cost \$199 in 10,000-unit quantities. We estimate that the XLP208 will cost \$125 to \$175 and Freescale's P3041 will cost \$150. We estimate \$250 for P4040, unless Freescale makes it a loss leader for the P4080. Freescale's manufacturing costs for SOI tend to be higher, and the company's primary foundry partner is IBM, which is generally considered more expensive than TSMC, the foundry preferred by NetLogic and Cavium. (Freescale is shifting some manufacturing to Global Foundries, the independent foundry spun off from AMD in 2008.)

The XLP208 and P3041 have the same power rating, with the CN6335 slightly worse. Because no measured power data is available for any of these chips, we can consider all three essentially the same for now. The XLP uses fewer CPUs, but their more complex design requires more power than the simpler CPUs of its competitors. If the XLP208 delivers the performance advantage we estimate, it will have the best throughput per watt. In features and power/performance ratio, the XLP208 dominates this field.

If greater performance is needed, NetLogic's XLP316 matches up well against Freescale's top-of-the-line P4080, offering a similar advantage in absolute performance and in performance per watt. The P4080 is nearly in production, however, putting it a year ahead of the XLP308. NetLogic's high-end XLP432 offers performance similar to that of Cavium's 32-core Octeon II, but at a lower power level.

XLP Family Validates Acquisition

The expanded XLP family greatly strengthens NetLogic's competitive position against Cavium, Freescale, and Intel. Cavium's Octeon family enjoys a power advantage over NetLogic's XLR family, but the XLP reverses the situation. Although NetLogic doesn't offer as many physical CPUs as Cavium does, its wider superscalar pipelines, out-of-order

execution, and multithreading largely compensate by providing alternative resources for instruction- and data-level parallelism.

We believe the top-end members of the XLP and Octeon II families will offer similar performance, regardless of the difference in CPU count. Even if they don't, NetLogic's glueless interchip interface lets customers link two to four chips together, delivering far more aggregate performance than Cavium can muster with a similar programming model.

At the high ends of their respective product lines, Cavium and NetLogic both require a data-plane application to execute 32 threads to reach peak performance, so their programming models are essentially the same. The main difference appears when running single-threaded control-plane code: a single thread can use all the XLP CPU's quad-issue resources, whereas the same thread would use only one-fourth the resources of a quad-CPU Octeon II chip.

The XLP's beefy caches and memory bandwidth show that NetLogic isn't ignoring the problem of keeping voracious CPUs well fed. The XLP is generously provisioned

with packet interfaces and other I/O capabilities, too. The on-chip ring network segregates contentious traffic and avoids the bottlenecks of conventional buses. The integrity of the whole design suggests that the XLP's architects did their data-flow homework before slapping down more CPU cores.

At 2.0GHz, the XLP chips will offer slightly better single-thread performance than Freescale's new QorIQ P5 chips, despite the P5's advantage in clock speed and its far better single-thread performance compared with Cavium's Octeon II. The XLP's single-thread performance is similar to that of slower Intel Xeon processors, and the XLP uses much less power than faster Xeon processors.

NetLogic's single-, dual-, and quad-core XLP chips do a good job of addressing the networking market's mainstream segment, and the eight-core chips address the high end. If the XLP family has any gap, it's the lack of \$10-\$20 chips for the low-end market that Cavium is pursuing. NetLogic says it perceives a greater opportunity in the control-plane market. If the company can deliver XLP on time and meet its promised performance targets, it should find plenty of new customers in that market. ♦

To subscribe to *Microprocessor Report*, access www.MPRonline.com or phone us at 408-945-3943.