The Linley Group

MICROPROCESSOR report

Insightful Analysis of Processor Technology

SOFT MACHINES TARGETS IPC BOTTLENECK

New CPU Approach Boosts Performance Using Virtual Cores

By Linley Gwennap (October 27, 2014)

Coming out of stealth mode at last week's Linley Processor Conference, Soft Machines disclosed a new CPU technology that greatly improves performance on single-threaded applications. The new VISC technology can convert a single software thread into multiple virtual threads, which it can then divide across multiple physical cores. This conversion happens inside the processor hardware and is thus invisible to the application and the software developer. Although this capability may seem impossible, Soft Machines has demonstrated its performance advantage using a test chip that implements a VISC design.

Without VISC, the only practical way to improve single-thread performance is to increase the parallelism (instructions per cycle, or IPC) of the CPU microarchitecture. Taken to the extreme, this approach results in massive designs such as Intel's Haswell and IBM's Power8 that deliver industry-leading performance but waste power and die area. But more-efficient designs, such as Cortex-A7, offer weak performance on single-threaded software, which still constitutes the majority of applications. VISC (which, according to the company, does not stand for virtual instruction set computing) delivers high performance on a single thread using a simpler, more efficient design.

As an added bonus, Soft Machines has created an intermediate software layer that can translate from any standard instruction set into threadable VISC instructions, as Figure 1 shows. The initial design uses this conversion layer to run ARM code, but the company claims it can create a conversion layer for x86 or other instruction sets as needed.

Soft Machines is no fly-by-night operation. The company has spent seven years and $125 million to develop and validate its technology. It currently has more than 250 employees, led by CEO Mahesh Lingareddy and president/CTO Mohammad Abdallah. Investors include AMD, GlobalFoundries, and Samsung as well as government investment funds from Abu Dhabi (Mubdala), Russia (Rusnano and RVC), and Saudi Arabia (KACST and Taqnia). Its board of directors is chaired by Global Foundries CEO Sanjay Jha and includes legendary entrepreneur Gordon Campbell.

Soft Machines hopes to license the VISC technology to other CPU-design companies, which could add it to their existing CPU cores. Because its fundamental benefit is better IPC, VISC could aid a range of applications from

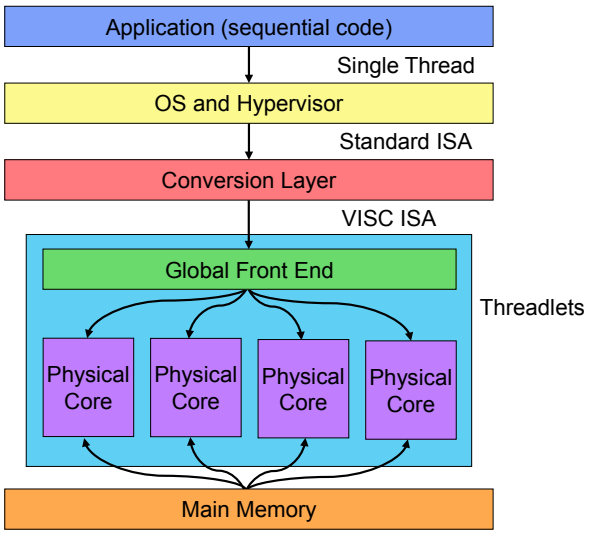

Figure 1. Soft Machines VISC technology. The conversion layer converts standard instructions (e.g., ARM) into VISC instructions. The hardware then converts a single instruction stream into multiple threadlets that can execute on multiple physical cores.

smartphones to servers. The company has applied for more than 80 patents on its technology.

Breaking the IPC Bottleneck

Since the first multiprocessor (SMP) computers appeared in the 1960s, programmers have realized that dividing their code to run on multiple CPUs improves throughput. In the past decade, multicore processor chips have driven SMP techniques into high-volume devices such as PCs, smartphones, and tablet computers.

Despite these advances, most programs continue to rely on a single instruction stream, or thread, to do most or all of the work. Breaking a program into multiple threads, ensuring that each can work independently with minimal inter-thread communication, is a challenging task. A few application types, such as graphics and packet processing, are simple to thread, but most others are not. For decades, development-tool vendors and other researchers have attempted to create a "magic compiler" that automatically creates a multithreaded program, but these tools either work only on a few easily parallelizable programs or require manual input from the programmer.

This failure puts the burden on CPU designers to improve single-thread performance. Once clock speeds hit a wall around 2005, performance per clock became the primary focus. Today's high-end CPUs, however, are well past the point of diminishing returns: adding another instruction-issue slot, another function unit, or a bigger reorder buffer burns more power while offering little IPC improvement. As a result, the IPC of high-end CPUs has increased at an annual rate of only 6% over the past decade, despite 40% annual growth in transistor count.

Thus, the implications of Soft Machines' VISC technology are tremendous. By breaking a single software thread into multiple hardware threads, the technology can combine multiple CPU cores into a single virtual core that delivers greater performance than any single CPU. This accomplishment could break the IPC bottleneck, accelerating performance scaling well beyond single-digit percentages to 2x or greater.

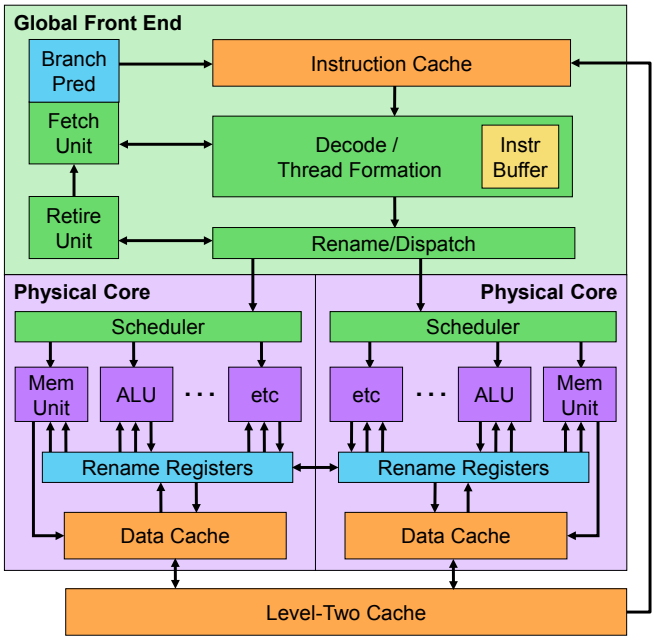

Figure 2. Conceptual diagram of a VISC processor. The global front end fetches a single instruction stream and divides it into threadlets that are dispatched to the multiple physical cores, which schedule and execute the instructions.

Pulling Apart the Threads

How can Soft Machines accomplish in hardware a task that software vendors have failed at? The hardware sees only a binary instruction stream, so it has less visibility than the compiler regarding data and code structures. The hardware has certain advantages, however, since it can access pointer values and other run-time information. At the conference, the company disclosed some description of its hardware design, although it did not provide complete details. The following provides a high-level view.

A VISC processor comprises a single instruction cache, a global front end, and two or more physical cores that each has its own instruction scheduling, register files, and data cache. The processor may also have a level-two (L2) cache that backs both the global instruction cache and the individual data caches. The initial test chip includes two physical cores and 1MB of L2 cache, as Figure 2 shows, but future designs are likely to include four or more cores.

The front end fetches instructions from the I-cache and places them into an instruction buffer. From this buffer, it attempts to form sequences of related instructions. For example, instructions with register dependencies can be grouped. These sequences are fairly short and are quite different from a conventional software thread; they can be better thought of as threadlets. The front end also performs global register renaming to avoid false dependencies. It can check pointers so as to group instructions that refer to the same memory location. The process of forming threadlets adds three cycles to the test chip's pipeline.

After creating the threadlets, the front end dispatches them to the physical cores. When it sends a threadlet to a core, it also allocates in that core the rename registers that are needed to execute the instructions. The goal is to minimize cases in which an instruction in one core requires a register allocated in another core, as this situation requires multiple cycles to resolve. Figure 2 shows a link between the two register files that handles this situation; the link becomes more complex in a processor with more than two cores.

Getting to the Core

Each core buffers the instructions received from the front end and can reorder them to avoid stalls. The buffer can

contain instructions from various threadlets, and these instructions can intermix freely. The core need not switch threads, because each threadlet has its own set of rename registers. When instructions are ready to execute, the core fetches data from its register file and completes the operations using its function units.

As in any out-of-order (OOO) design, speculative results are held pending until the completion of all previous instructions. In the VISC approach, the retirement unit must track instructions across all physical cores, since the original instruction stream could have been divided among multiple cores. Similarly, a branch misprediction detected in one core can affect instructions executing in other cores.

Thus, a VISC processor can be viewed as having global hardware for branch prediction as well as for fetching, grouping, tracking, and retiring instructions, but local hardware for scheduling and executing instructions and accessing memory. Compared with a high-IPC processor such as Haswell, the front end is of similar complexity, but the scheduler in each physical core is much simpler, as it manages only a few function units (versus eight in Haswell). The data cache also has fewer ports and can cycle faster. Because the execution resources of both cores can apply to a single thread, however, even a two-core VISC processor can deliver total ALU operations or memory operations per cycle that match or exceed those of Haswell.

VISC relies on a unique internal instruction set to do its magic, but true to its name, Soft Machines provides a conversion layer that converts from a standard instruction set to VISC. This approach is similar to how Transmeta processors executed x86 instructions (see *MPR 1/24/00,* "Transmeta Unveils Crusoe"), but VISC uses a completely different internal architecture. The company says the conversion overhead is less than 5%. In addition to its ARM-to-VISC translator, Soft Machines has prototyped an x86 translator and says it can develop other translators on customer request.

Better Load Balancing

The previous example shows how the global front end can break down a single thread, but it can also handle multiple software threads at the same time. In this way, a single VISC processor can emulate a traditional multicore design. But unlike traditional designs, it can more easily perform load balancing, matching processing power to the task at hand.

Figure 3 shows an example with two active threads: one heavy (high performance) and one light. In a processor with several identical cores, each thread would run on one core, wasting cycles for the light thread and limiting performance for the heavy thread. In a VISC design with two physical cores, the heavy thread is split into two hardware threads, one running on each core. Note that the second core automatically shares its resources between the heavy thread and the light thread, because a VISC core can mix instructions from multiple threadlets.

For More Information

A copy of the Soft Machines presentation from the Linley Processor Conference is available for free download at *www.linleygroup.com/events/proceedings.php?num=29* (registration required). For additional information on Soft Machines, access the company's web site at *www.softmachines.com*.

Soft Machines calls this approach "virtual cores." From a software viewpoint, the heavy thread runs on a single virtual core that comprises more than one physical core, while the light thread runs on a virtual core that uses only a portion of a physical core. This allocation of cores is invisible to software, except that the heavy thread runs faster than it would otherwise. According to the company, its front-end hardware will recognize which threads need more performance and allocate the virtual core resources appropriately. The operating system need not know the number of physical cores in the processor to assign threads or balance the load.

Demonstrated Performance

Soft Machines has designed and fabricated a test chip that implements its VISC architecture using two physical cores. It refused to discuss the number of function units or other basic microarchitecture capabilities of these cores but characterized them as A15-class CPUs. We interpret this statement to mean that each core can execute three to four operations per cycle with moderate instruction reordering.

The company also withheld the test chip's clock speed. Because it uses a pipeline with only 10 stages (including some extra stages for VISC scheduling), the CPU cannot match the high clock speeds of leading-edge x86 and ARM processors. We estimate the chip runs at several hundred megahertz. Even at this low speed, it completes some programs in less time than a low-end Haswell processor.

Despite its relatively simple design, the chip achieves spectacular performance. On the single-thread SPEC2006

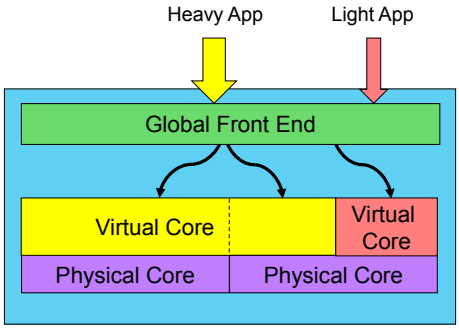

Figure 3. Soft Machines virtual cores. A single virtual core can comprise multiple physical cores or only a portion of a physical core. In this way, the front-end hardware aligns the performance of the virtual core with the needs of each thread.

test suite, the company reports an average IPC of 2.1, counting ARM instructions rather than VISC instructions. This IPC compares with 0.71 for Cortex-A15 and 1.39 for Haswell. (For consistency, Soft Machines measured the IPC on all three processors using GCC rather than Intel's favorite compiler, ICC.) Thus, the VISC chip achieved three times the IPC of ARM's highest-end CPU shipping today and 50% better IPC than Intel's fastest mainstream CPU.

Figure 4 shows more-detailed performance results comparing the VISC test chip against Cortex-A15 (measured in a Samsung Exynos processor). The figure shows the results of 62 different benchmarks, including components of SPEC2000, SPEC2006, EEMBC's digital entertainment benchmark (DenBench), and the Kraken JavaScript benchmark. While individual results range from 1.5x to 7x, the average gain across these tests ranges between 3x and 4x. Soft Machines points out that although the test chip runs Linux and other applications, the software contains workarounds for certain hardware bugs, so it will likely obtain better results with future tuning and bug fixes.

Although these results are impressive, they require some caveats to put them into perspective. A shorter CPU pipeline reduces branch penalties and other pipeline hazards, thereby improving IPC compared with a longer pipeline. In addition, a low CPU speed reduces the effective latency of caches and main memory (measured in CPU cycles), again improving IPC relative to a CPU with a faster clock. The latter effect might explain why the test chip appears to perform better on SPEC2006 than on SPEC2000, which has a smaller memory footprint.

Doubling Performance Is Realistic

The test results are difficult to interpret, owing to both the lack of information about the test chip's core CPU and the effects of the short pipeline. A better way to quantify the advantage of using VISC is the performance increase relative to a single core. By coupling the resources of two cores to execute a single thread, the maximum performance increase (compared with a single core of similar design) is 2x. In reality, the front end cannot fully utilize the resources of the second core, particularly for code that is difficult to break apart. For example, code that contains many branches or many register dependencies will be harder to thread and thus will use less of the second core.

Soft Machines has run many tests and simulations of its VISC technology. It estimates the second core improves performance by an average of 50–60% across a variety of benchmarks. This factor implies that the test chip would achieve about 1.3 IPC when using a single core—considerably higher than Cortex-A15. This higher IPC includes the effects of the lower clock speed and shorter pipeline.

Although the test chip has only two physical cores, Soft Machines has run simulations on a four-core design. As one might expect, the performance gains diminish for the additional cores: the third core adds 20–30% to single-thread performance, and the fourth adds only 10–20%. In total, the four-core design delivers about twice the performance of a single core. The unused resources in the extra cores, however, can be devoted to additional threads. For example, a four-core design could run two threads at close to their maximum performance.

Performance-critical applications will benefit from VISC, but the technology can also apply to low-power designs. As the test chip demonstrates, a VISC design can operate at a relatively low clock speed while achieving the same absolute performance as a traditional design operating at a higher clock speed. Thus, it should use less power, particularly if the voltage is reduced as well. Soft Machines, however, declined to reveal the power consumption of its test chip.

Other details also remain undisclosed, including die area. Details of how the processor handles privileged operations, inter-thread synchronization, traps, and interrupts could all affect how well it runs certain applications. Performance could vary widely across different workloads. The company provides additional information to potential customers under NDA and plans to make further disclosures over time.

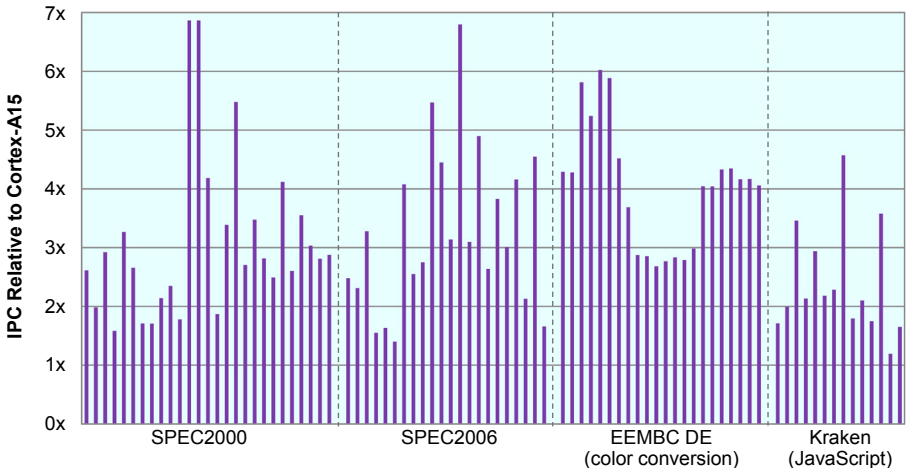

Figure 4. VISC performance results. This chart shows measured instructions per cycle (IPC) for Soft Machines' test chip normalized against Cortex-A15. The company did not disclose the clock speed of the test chip. (Source: Soft Machines)

A Rising Tide

Soft Machines has identified a critical (perhaps the most critical) problem in CPU design today: the minimal improvement in single-thread performance over the past decade. Despite the obvious need for multi-threaded software, most applications—even performance-intensive ones—continue to

rely on a single thread. The software tools for creating multithreaded applications remain primitive.

With the announcement and demonstration of its VISC technology, Soft Machines has taken a big step toward solving this problem. By shifting the burden to hardware, VISC aims to deliver the benefits of multithreading to all applications. The initial performance results are excellent: a 50–100% gain in single-thread IPC represents a decade of progress at the industry's current sluggish pace. As with any radical new technology, however, we remain skeptical, particularly given the startup's limited disclosure. Potential customers must fully assess the technology to determine how it will perform in an actual product across a range of workloads.

Assuming the technology works as advertised, it will change the way all processors are designed. CPU designers will stop trying to improve IPC by adding more hardware; in fact, complex high-IPC designs like Haswell could disappear in favor of smaller, simpler ones. Replacing large cores with clusters of simpler cores will improve performance and power efficiency, benefiting almost every type of processor. To deliver on this promise, Soft Machines must fully validate VISC and license it to leading processor vendors. ♦

The Linley Group

MICROPROCESSOR report

Insightful Analysis of Processor Technology

APPLIEDMICRO ARMS FOR EMBEDDED

New Helix Family Inherits DNA From X-Gene Server Processors

By Tom R. Halfhill (October 27, 2014)

AppliedMicro's X-Gene server processors are spawning a new family of ARM-compatible embedded processors intended mainly for communications. The first two members of the new Helix family use existing X-Gene die built in 40nm CMOS technology, but future products include new designs built in a 28nm high-*k* metal-gate (HKMG) process. All are compatible with the 64-bit ARMv8 architecture.

The Helix family will eventually supersede AppliedMicro's PacketPro APM86xxx embedded processors, which have 32-bit PowerPC 460 CPUs and are manufactured in 40nm technology. Those single- and dual-core chips are highly optimized for packet processing and communications. By contrast, Helix chips have two, four, or eight 64-bit CPUs, and we believe they have much of the same packet acceleration as the PacketPro Mamba and Diamondback processors (see *MPR 12/20/10,* "First PacketPro Chips Debut").

AppliedMicro pitches the four- and eight-core Helix processors for service-provider routers, networking line cards, enterprise Layer 2–4 switches, enterprise storage-area networks (SANs), prosumer network-attached storage (NAS), and networked industrial machinery. The dual-core Helix chips target enterprise multifunction printers, low-end service routers, and top-of-rack (ToR) switches.

As Figure 1 shows, Helix integrates AppliedMicro's custom ARM-compatible CPUs with the usual I/O interfaces found in communications processors. Even the dual-core models have at least one 10G Ethernet controller, which is uncommon. All Helix processors also implement third-generation PCI Express, Serial ATA (SATA), and USB. AppliedMicro says the 40nm chips will begin production in 1Q15, followed by the 28nm designs in 2H15.

Some Helix processors are X-Gene derivatives and others are new silicon. The Helix 1 series is based on the same die as X-Gene 1 (40nm G), and the larger members of the Helix 2 series are based on the same die as X-Gene 2 (28nm HPM). We expect the dual- and quad-core Helix 2 chips to debut on a smaller die before the eight-core Helix 2 chips appear.

In the embedded-processor market, Helix faces MIPS64-compatible and Power Architecture processors

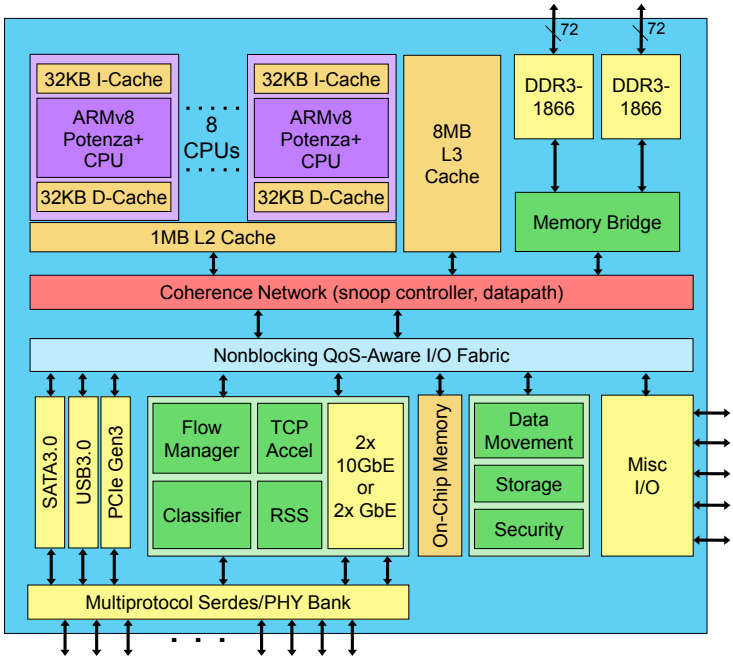

Figure 1. Block diagram of AppliedMicro's Helix 2 processor. The dual- and quad-core versions resemble this eight-core design. All have similar I/O interfaces but less memory bandwidth than the company's X-Gene server processors.

from Broadcom, Cavium, and Freescale, as well as ARMv8-compatible processors coming soon from those vendors. Additional competitors could include AMD, Intel, Qualcomm, and Texas Instruments, although some of those vendors have yet to announce their ARMv8 products.

Improved CPU Boosts Performance

To launch Helix, AppliedMicro is repackaging its X-Gene server chips (see *MPR 9/1/14,* "X-Gene 2 Aims Above Microservers"). If some Helix products can salvage silicon that isn't quite server grade, they can increase the company's effective yields and win designs in the embedded market while incurring only the incremental costs of packaging and marketing. Also, if Helix 1 debuts on schedule, it will reach the market ahead of rival ARMv8 products.

Table 1 compares all the initial Helix 1 and Helix 2 processors. Helix 1 uses the same "Potenza" CPU as X-Gene 1, whereas Helix 2 uses the same "Potenza+" CPU as X-Gene 2. Both CPUs are AppliedMicro's ARMv8-compatible designs created under an ARM architecture license. Like most ARMv8 vendors, AppliedMicro seeks to differentiate its processors by using a custom CPU instead of licensing an off-the-shelf core from ARM.

Both Potenza and Potenza+ are four-issue superscalar CPUs that implement a single thread and out-of-order execution (see *MPR 9/17/12,* "X-Gene CPU Design Strikes a Balance"). Both CPUs are more powerful than the 32-bit PowerPC 460 cores in PacketPro. In Helix, their maximum advertised clock frequency is 2.0GHz, whereas X-Gene server processors target 2.4GHz in 40nm technology and 2.8GHz in 28nm HPM. AppliedMicro is restraining the embedded processors to reduce power but will supply the higher clock speeds on customer demand.

Each CPU includes integer, floating-point, and 128-bit SIMD units, as well as 32KB L1 instruction and data caches. Building on experience with the original Potenza design, AppliedMicro enhanced the Potenza+ microarchitecture to increase the instructions per cycle (IPC) by 10–15%. (It withheld details of these improvements.) Pairs of CPUs share a rather small 256KB L2 cache, but all the CPUs share a large 8MB L3 cache.

Helix offers much less memory bandwidth than X-Gene. Whereas the latter has four 64-bit DDR3/4 controllers supporting effective transfer rates of up to 1,866MT/s, the former has only one or two 64-bit DDR3-1600 or DDR3-1866 controllers. Consequently, Helix provides no more than half of X-Gene's peak DRAM bandwidth. AppliedMicro says the bandwidth is sufficient for the target applications—a contention that competitors dispute. In any case, reducing the number of DRAM controllers saves hundreds of pins, which cuts packaging costs and perhaps will enable the company to repurpose some X-Gene die that emerge from fabrication with subpar controllers.

Ethernet connectivity is crucial for communications processors, so all Helix devices support both 10GbE and GbE. Those capabilities aren't universal in the dual- and quad-core members of competing product families, although newer chips are adding 10GbE. In the eight-core class, however, Helix suffers a comparative disadvantage. Rival products have more 10GbE and GbE ports, and Broadcom and Cavium are on the verge of shipping eight-core processors with 40GbE (see *MPR 3/31/14,* "Broadcom's XLP500 Raises Midrange").

The four- and eight-core Helix 2 products are similar to the Helix 1 versions. Using the enhanced Potenza+ CPU, they offer slightly more performance at the same clock speed. We estimate that moving to 28nm HKMG technology shaves 5W off the TDP of the quad-core version and 10W off the TDP of the octa-core version. An important addition to the Helix 2 series is a dual-core processor, enabling AppliedMicro to target lower-cost designs that PacketPro currently serves.

Black-Box Hardware Acceleration

Helix's other I/O interfaces are typical for embedded processors, and they support the third-generation implementations

	Helix 1	Helix 1	Helix 2	Helix 2	Helix 2	Helix 2
CPU Type	Potenza		Potenza+			
CPU Cores	4 CPUs	8 CPUs	2 CPUs	4 CPUs	4 CPUs	8 CPUs
CPU Freq	1.2–2.0GHz*		1.2–2.0GHz*			
L2 Cache	512KB L2	1MB L2	256KB L2	512KB L2	512KB L2	1MB L2
L3 Cache	8MB shared L3 cache		8MB shared L3 cache			
DRAM Type	2x 64-bit DDR3-1600 +ECC		1x 64-bit DDR3-1866 +ECC			
DRAM B/W	25.6GB/s		14.9GB/s	14.9GB/s	29.8GB/s	29.8GB/s
Ethernet MACs	2x 10GbE, 4x GbE		1x 10GbE, 4x GbE		2x 10GbE or 2x GbE	
PCI Express (Gen3)	2x8 PCIe	4x4 PCIe	1x4 + 2x1 PCIe	1x4 + 2x1 PCIe	1x8 PCIe	2x4 PCIe
Serial ATA	2x SATA3.0	2x SATA3.0	1x SATA3.0	1x SATA3.0	2x SATA3.0	2x SATA3.0
Power (TDP)	20W†	40W†	8W†	16W†	15W†	30W†
IC Process	TSMC 40nm G		TSMC 28nm HPM			
Production	1Q15 (est)		2H15 (est)			

Table 1. AppliedMicro's Helix embedded processors. The company has yet to announce model numbers, but the most important distinguishing features are the CPUs, DRAM interfaces, Ethernet controllers, TDPs, and production schedules. *AppliedMicro will supply Helix 1 chips at 2.4GHz and Helix 2 chips at 2.8GHz on customer demand, but our estimated TDPs are for 2.0GHz parts. (Source: AppliedMicro, except †The Linley Group estimate)

that will be common when these products ship: PCIe Gen3, SATA3.0, and SuperSpeed USB3.0. The SATA interfaces make the smaller chips suitable for SAN and NAS subsystems.

AppliedMicro is withholding many details about the hardware accelerators. Helix will definitely have cryptography acceleration and other engines that offload packet-processing tasks from the CPUs. In embedded processors of this class, a regular-expression (reg-ex) engine that accelerates deep packet inspection is desirable, as is an engine that accelerates data compression and decompression.

The server processors include hardware for TCP acceleration, Receive-Side Scaling (RSS), Data Center Bridging, and I/O Virtualization. RSS enables a network interface card (NIC) to steer TCP segments for a given packet flow to the same CPU, thereby creating an affinity between flows and CPUs. X-Gene also inherits the RAID acceleration and SlimPro engine from the company's PacketPro line.

SlimPro is a dedicated management processor that actually uses a 32-bit ARM Cortex-M3 core with dedicated on-chip RAM and ROM. It handles functions such as secure booting, power management, and run-time code authentication. SlimPro can also hold cryptographic keys for the security engine (see *MPR 6/20/11,* "ARM CPU Secures APM Processor").

We suspect that Helix (and X-Gene) processors have inherited additional PacketPro hardware. They likely have packet classifiers, an in-line crypto engine, a queue manager (QM), and a traffic manager (TM). Using this hardware, the PacketPro APM86692 can perform Layer 3 forwarding and packet classification at 12 million packets per second (Mpps). The QM/TM can distribute packets received from the Ethernet ports to the CPUs, performing load balancing or buffering packets until the appropriate CPU is available. Thus, it can ensure quality of service (QoS) for network traffic traversing the processor. If, as we surmise, Helix has these features, it should be able to perform the same functions.

To differentiate X-Gene server processors from their embedded cousins, AppliedMicro may enable these engines and other accelerators only for the Helix products. Doing so would not only optimize the embedded chips for their different markets but also save some slightly out-of-spec die that the company would otherwise discard.

Eight-ARM Octopus Chips Grapple for Sales

Table 2 compares the octa-core Helix 2 with two octa-core competitors: a still-unnamed member of Cavium's Octeon III CN72xx family and Freescale's QorIQ LS2085A. All three processors may reach production by the end of 2015.

One of Cavium's ARMv8-compatible ThunderX processors would be a better match than the MIPS64-compatible CN72xx, but too few details are currently available to draw a fair comparison. Ditto for Broadcom's ARMv8 Vulcan processors. AMD's octa-core Hierofalcon embedded processor is another potential challenger, but it's really a server design that is less optimized for communications (see *MPR 2/24/14,* "AMD to Sample First ARM Chip").

Helix 2 has the most powerful CPU—it's a four-issue superscalar ARMv8 core with instruction reordering. Freescale's LS2085A also supports ARMv8 but uses ARM's Cortex-A57 instead of a custom-designed core; it's a three-issue superscalar CPU with some reordering capability. Cavium's CPU is a simpler dual-issue core with rudimentary reordering, sacrificing single-thread performance to conserve power.

Those differences partly account for the preliminary power-consumption estimates. Freescale estimates its LS2085A will consume 30W TDP (thermal design power—a near-maximum rating). We estimate that Helix 2 will also consume about 30W TDP and that Cavium's CN72xx will consume 20W maximum (meaning its TDP will be around 18W).

In networking and communications, packet-acceleration hardware is often more important than the CPUs. If, as we suspect, Helix inherits PacketPro's DNA, the new processors will be suited to communications. But we believe Cavium and Freescale have faster accelerators that will outperform Helix on IPSec and packet forwarding. Also, the Cavium and Freescale processors have reg-ex engines for deep packet inspection and hardware accelerators for data compression and decompression—useful features that the Helix chips likely omit.

	AppliedMicro Helix 2	Cavium Octeon III CN72xx	Freescale QorIQ LS2085A
CPU Architecture	ARMv8	MIPS64 R5	ARMv8
CPU Cores	8x Potenza+	8x Octeon III	8x Cortex-A57
CPU Freq (max)	2.0GHz	2.5GHz	2.0GHz
L2 Cache	4x 256KB L2	2MB shared L2	4x 1MB L2
L3 Cache	8MB shared L3	None	1MB shared L3
IPSec Perf	8Gbps*	30Gbps	20Gbps
Compression Perf	Not disclosed	20Gbps	12Gbps
Reg-Ex Perf	Not disclosed	20Gbps	10Gbps
Layer 3 Fwding	24Gbps*	40Gbps	40Gbps
DRAM Interface (+ECC)	2x 64-bit DDR3-1866	2x 64-bit DDR4/3L-2400	2x 64-bit DDR4-2133
DRAM Bandwidth	29.8GB/s	38.4GB/s	34.1GB/s
Ethernet MACs	2x 10GbE, 2x GbE	40GbE, 10GbE, GbE	8x 10GbE + 8x GbE + 8-port switches
PCI Express (Gen3)	2x4 PCIe	4x PCIe	1x8, 4x4, 4x2, or 4x1 PCIe
Serial ATA	2x SATA3.0	2x SATA3.0	2x SATA3.0
Power	30W TDP*	20W (max)*	30W TDP
IC Process	28nm HPM	28nm HKMG	28nm HPM
Production	2H15 (est)	2Q15 (est)*	3Q15 (est)

Table 2. AppliedMicro's Helix 2 versus competitors. Helix has the most powerful CPU and the most cache, but the other processors have more memory bandwidth and Ethernet connectivity. (Source: vendors, except *The Linley Group estimate)

Price and Availability

AppliedMicro's Helix 1 quad- and octa-core embedded processors are sampling this quarter; the company expects production to start in 1Q15. Helix 2 is scheduled to begin production in 2H15. The company has yet to announce prices. For more information, access *www.apm.com/products/embedded/Helix-family*.

More Cache Eases Bottlenecks

The Cavium and Freescale products offer more memory bandwidth but less cache than Helix 2. Although all three have dual 64-bit DRAM controllers with optional ECC, the Helix 2 controllers are a little slower, providing only 29.8GB/s of peak bandwidth. Cavium's CN72xx delivers 38.4GB/s, and Freescale's LS2085A delivers 34.1GB/s. Helix 2 compensates by offering 9MB of total cache—much more than Freescale's 5MB and Cavium's 2MB.

Networking interfaces favor the Cavium and Freescale processors. Although Cavium hasn't disclosed details, the CN72xx will have 40GbE, 10GbE, and GbE ports; it's the only one in this comparison to support 40GbE. Freescale's LS2085A lacks that feature but offers eight 10GbE and eight GbE ports. Moreover, each of the LS2085A's Ethernet clusters integrates an eight-port wire-speed switch. Helix 2 has only two 10GbE and two GbE ports without a switch. Whether this connectivity is sufficient depends on the application and on Helix 2's undisclosed packet acceleration. If it has less acceleration than the competing processors, it can get by with less network bandwidth because it will serve in applications that handle lower packet rates.

All three processors support PCIe Gen3, SATA3.0, and SuperSpeed USB3.0. Cavium's CN72xx adds Interlaken and Interlaken-LA interfaces for attaching network search engines and other external devices. None of these processors has serial RapidIO (sRIO), which rarely matters for their target markets.

Overall, Helix 2 offers more CPU performance and cache in return for less memory bandwidth and network throughput. Its packet acceleration is uncertain, but it's probably less suitable for deep packet inspection than the other processors. The most likely applications for Helix are those requiring control-plane or mixed control/data-plane processing at packet rates for which one or two 10GbE ports will be sufficient.

More Opportunities, But More Competitors

Helix processors take advantage of the company's significant investments in ARMv8 CPUs and multicore designs for the server market. They also reuse some existing PacketPro technology created for the embedded market. By introducing larger embedded processors capable of tackling bigger applications, AppliedMicro can attract new customers and offer upgrades to PacketPro customers that need higher performance. Also, repackaging some off-spec X-Gene server silicon as Helix products will boost effective yields and cut costs.

Chips designed to bridge both the server and embedded markets will be less optimized than competing chips designed solely for one market, however. And unlike the server market, which is dominated by one bull elephant, the embedded domain is populated by a pack of hungry hyenas, so there's more competition.

Today, AppliedMicro's single- and dual-core PacketPro chips are designed solely for embedded applications, and they compete with similar products from major vendors like Broadcom, Cavium, and Freescale. By introducing processors with up to eight CPUs, AppliedMicro is challenging those companies' midrange embedded processors, which are unburdened by design decisions that partly aim to meet server requirements.

An AppliedMicro design optimized for servers won't have all the features that a midrange communications processor needs. (Witness AMD's Hierofalcon.) Conversely, a design optimized for communications will have excess baggage for servers—and the whole point of an ARM-based server is to achieve greater power efficiency. Multiple chip designs are the answer, but a smaller company like AppliedMicro will strain to offer a product line as diverse as its competitors'.

On the other hand, it's possible that Helix is actually the main strategy and X-Gene is the spawn. Although the company announced X-Gene first, it's in effect using Helix to expand its traditional market (embedded) while using the same design to enter a new growth market (power-efficient servers). Given the size and ferocity of the bull elephant ruling the server range, maybe it's wiser for AppliedMicro to optimize its processors for communications and welcome any new server business as a bonus.

AppliedMicro has an opportunity to seize a critical time-to-market advantage. Since Helix 1 is scheduled to ship in 1Q15, the company can win designs before competitors get their first embedded ARMv8 processors out the door. (Indeed, it already claims some design wins.) And if Helix 2 ships on schedule in 2H15, the 28nm products will hit the market at about the same time as competitors' first ARMv8 products at that node, thus keeping Helix in the race. AppliedMicro's early investment in ARM is now about to pay off. ♦

The Linley Group

MICROPROCESSOR report

Insightful Analysis of Processor Technology

FREESCALE'S LS1 GETS 64 BITS

New QorIQ LS1043A and LS1023A Processors Use Cortex-A53

By Tom R. Halfhill (October 27, 2014)

Once an exclusive feature of servers, workstations, and supercomputers, 64-bit CPUs are now spreading even to some low-end embedded processors. Freescale's new QorIQ LS1043A and LS1023A are the first 64-bit chips in the entry-level LS1 series. They integrate up to four ARM Cortex-A53 CPUs with a cryptography engine, packet acceleration, 10 Gigabit Ethernet, and DDR4 memory control. Despite their maximum target clock frequency of 1.5GHz, they consume only 8W or less—cool enough for fanless systems.

Freescale unveiled the new products at the recent Linley Processor Conference, pitching them for low-power customer premises equipment (CPE). Applications include integrated-services branch routers, security appliances, industrial controllers, and edge devices that implement software-defined networking (SDN) and network-function virtualization (NFV).

As Figure 1 shows, they have several features common to Freescale's other QorIQ processors, including the packet accelerators known as the Data Path Acceleration Architecture (DPAA). These programmable accelerators enable the new LS1 chips to parse, classify, queue, and forward network packets at wire speeds of up to 10Gbps. The SEC 5.4 crypto engine supports IPSec rates of up to 5Gbps. These accelerators offload the CPUs so completely that the throughput is the same for the dual- and quad-core models.

Freescale hopes to sample the new chips next quarter and begin production in 4Q15. They will join three other LS1-series processors that have been sampling since April and one that's already in production. The company hasn't announced pricing, but we estimate the LS1023A and LS1043A will cost $40 to $50 in volume—about the same as similar Power-based chips in the QorIQ T1 series.

Waiting for Cortex-A53 to Ripen

Whereas competitors like AppliedMicro, Broadcom, and Cavium have designed their own ARMv8-compatible CPUs under their ARM architecture licenses, Freescale has so far preferred to license ARM's off-the-shelf cores. The higher-end LS2-series uses the more-powerful Cortex-A57. Existing LS1-series chips have 32-bit ARM cores. For these

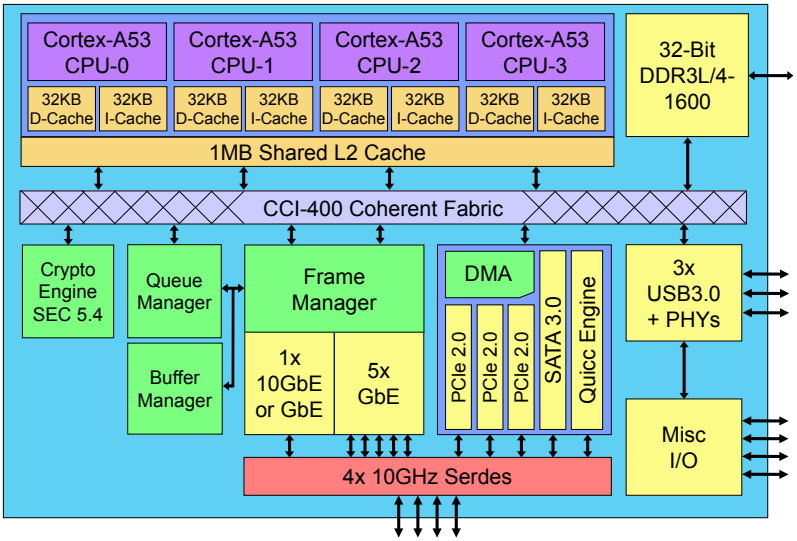

Figure 1. Block diagram of Freescale QorIQ LS1043A. The dual-core LS1023A is otherwise identical to this quad-core design. Freescale's ARM-based QorIQ processors have much in common with the Power Architecture QorIQs, but they substitute ARM's Coherent Cache Interconnect for the CoreNet fabric in the Power-based products.

initial 64-bit LS1 models, Freescale is licensing Cortex-A53, which is ARM's smallest 64-bit CPU (see *MPR 11/5/12,* "Cortex-A53 Is ARM's Next Little Thing").

Although Cortex-A53 smartphone processors have been appearing for a while now, Freescale waited for ARM to fix some errata before using the CPU in embedded processors that require 10-year availability and are intended for enterprise equipment. Despite its lower power, it still packs a punch. In simulations, Freescale showed the quad-core LS1043A exceeding 16,000 CoreMarks at 1.5GHz. That's about 12% more CoreMarks per megahertz than the 32-bit Cortex-A7 CPUs in other LS1-series chips.

To keep the CPUs fed, a 32-bit DDR4-1600 DRAM controller provides 6.4GB/s of peak memory bandwidth. DDR4 supports much faster memory, but Freescale is limiting the speed to save power. (The controller also supports low-power DDR3L DRAM and has optional ECC.) Although 6.4GB/s is no better than what most dual-core 32-bit LS1 processors deliver, Freescale says it's sufficient for the target applications.

A larger cache reduces the need to access external memory. Both the LS1023A and LS1043A have 1MB of L2 cache—twice as much as other LS1 chips. Unlike most other QorIQ processors, however, there's no third-level platform cache. In small processors like these, according to Freescale, it's more sensible to put a larger L2 cache closer to the CPUs than to add a higher-level cache.

For packet I/O, both processors have five Gigabit Ethernet (GbE) controllers and one 10GbE controller that can alternatively drive a sixth GbE port. These controllers support XFI, Quad SGMII, 2.5Gbps SGMII, RGMII, and IEEE 1588 time stamping.

	Freescale QorIQ LS1023A	Freescale QorIQ LS1043A
CPUs	2x ARM Cortex-A53	4x ARM Cortex-A53
CPU Freq (max)	1.5GHz	
L2 Cache (+ECC)	Shared 1MB L2 cache	
Crypto Engine	SEC 5.4 (XOR, CRC)	
Other Engines	Quicc Engine, DPAA packet acceleration	
Packet Rate (max)	10Gbps	
IPSec Perf (max)*	5Gbps	
DRAM Interface	32-bit DDR3L/4-1600 +ECC	
DRAM B/W	6.4GB/s	
PCI Express	PCIe 3x1 or 4x1 Gen2	
Ethernet MACs	1x 10GbE + 5x GbE or 6x GbE	
Serial ATA	1x SATA3.0	
Serdes Lanes	4 lanes	
USB	3x USB3.0 + PHY	
Power	6W TDP	8W TDP
IC Process	TSMC 28nm HPM	
List Price (1,000s)	$40†	$50†
Samples	1Q15 (est)	
Production	4Q15 (est)	

Table 1. Freescale's new QorIQ LS1023A and LS1043A processors. Both products are likely based on the same die, enabling the company to reject fewer parts during verification. *Maximum IPSec performance with IMIX packets. (Source: Freescale, except †The Linley Group estimate)

The interface flexibility is important because these products must share only four 10GHz serdes among the Ethernet, PCI Express (PCIe), and Serial ATA (SATA) controllers. Using Quad SGMII, four GbE controllers can share one serdes lane, freeing three lanes for other purposes. The lone 10GbE controller is a limitation, however—high-reliability systems require duplicate 10GbE ports for redundancy.

Slower PCI Express Not a Handicap

Table 1 compares the new LS1 processors. Unlike LS2 chips, they support PCIe Gen2 (500MB/s) instead of Gen3 (985MB/s). Freescale says the lower-end processors likely won't need the higher data rate because customers will probably use these interfaces to attach Wi-Fi transceivers, print-engine ASICs, or supplementary SATA controllers for RAID subsystems; for those purposes, Gen2 is sufficient. Both processors have three PCIe controllers that can operate in 3x1 or 1x4 modes.

The SATA3.0 controller supports external disk drives and is sufficient for the target applications. One controller isn't enough to build a RAID box, however, so system designers must attach an external SATA controller to a PCIe interface. To support legacy applications, the new processors include a compact version of Freescale's Quicc Engine, which has been a mainstay of the company's communications chips since the Motorola days. This programmable accelerator supports older networking and industrial protocols such as Profibus, HDLC, and TDM. Existing customers will have to port their application code from Power to ARM, but some of them are growing accustomed to the routine, having previously ported code from the Motorola 68000 to Power.

Freescale's SEC 5.4 security engine supports the usual cryptographic standards and protocols plus Secure Boot, the Trust Architecture, and ARM's TrustZone. Additional virtualization hardware enables a hypervisor to partition the processing resources, memory, and I/O interfaces to isolate multiple virtual machines. Three SuperSpeed USB3.0 controllers reduce contentions for those commonly used interfaces, and their integrated physical-layer hardware eliminates external PHY chips.

Thermal design power (TDP) is relatively low for 64-bit multicore processors that can hit 1.5GHz: only 8W for the quad-core LS1043A and 6W for the dual-core LS1023A. These are near-maximum ratings, including I/O power, at junction temperatures not exceeding 105°C. Thus, the new communications chips are suitable for fanless systems.

Most Powerful Low-Watt ARMv8

By expanding the LS1 series with two 64-bit processors—including the first quad-core chip in the series—Freescale is significantly strengthening its lowest-power, lowest-priced communications products. The LS1043A and LS1023A

join three dual-core processors that use ARM's 32-bit Cortex-A7 and one dual-core processor that uses ARM's 32-bit Cortex-A9. (The A9 product, now called the LS1024A, is actually the former Mindspeed Comcerto C2200 gateway processor that Freescale acquired from Macom; see *MPR 5/5/14,* "Freescale Buys Mindspeed's Comcerto.")

Competitors can't yet match the LS1043A's combination of four 64-bit ARM CPUs and 8W TDP. Broadcom has yet to announce specific details of its ARMv8-compatible Vulcan products; ditto for Cavium's ARMv8 ThunderX family. Broadcom's ARM-based StrataGX aims for similar applications but offers no more than two 32-bit CPUs.

AppliedMicro recently announced a quad-core ARMv8-compatible embedded processor that has dual 10GbE controllers and much more DRAM bandwidth, but it's a relabeled eight-core server processor that dissipates about 20W TDP (see *MPR 10/27/14,* "AppliedMicro ARMs for Embedded").

Consequently, Table 2 compares the new Freescale products with two MIPS64-compatible processors from Broadcom and Cavium plus one of Freescale's Power Architecture QorIQ chips. All are recent 64-bit designs. Freescale's T1042 is the only one without 10GbE—a characteristic omission of the T1 series that emphasizes the significance of introducing this capability in the low-end LS1 series. One distinguishing feature that the T1 series does offer, however, is an eight-port GbE switch in the T1040, which is otherwise almost identical to the T1042 (see *MPR 7/2/12,* "Freescale Drops Quad-Core Threshold").

Another feature the LS1043A did not inherit from the T1 series is the pattern-matching engine (PME), Freescale's name for its regular-expression (reg-ex) accelerator. This omission is a bigger disappointment. Most dual- and quad-core communications processors now include this feature, among them the rival Broadcom and Cavium chips in Table 2. Both of those products also include accelerators for compressing and decompressing data. The combination of a reg-ex engine and compression engine is valuable for network-security software that must rapidly decompress and scan email attachments for malware.

MIPS Processors Still Look Good

Broadcom's XLP208, a member of the XLP II family, is the only dual-core processor in this group, but it has the most powerful CPUs. The FC4400 is a quad-issue superscalar CPU with four-way multithreading, and its 2.0GHz maximum clock frequency is 11% to 33% faster than that of the other CPUs in this comparison. As a result, it excels at both single- and multithread performance and will likely outrun these quad-core products. It also has two 10GbE ports, suiting it to applications that need duplicate network connections for higher reliability (see *MPR 10/22/12,* "Broadcom Samples 28nm XLP II").

The other MIPS64-compatible processor in this group is Cavium's Octeon III CN7130. Like the LS1043A, its CPUs are dual-issue superscalar cores with in-order execution. At their maximum clock speed of 1.8GHz, however, they are 20% faster than the CPUs in the LS1043A. As noted above, the CN7130 also has advantages in reg-ex and compression acceleration (see *MPR 6/24/13,* "Octeon III Starts at Low End").

The T1042 and XLP208 offer twice the memory bandwidth of the LS1043A. The XLP208 is particularly generous because it also has much more cache than the other processors (3MB total). Whereas all these chips have 1MB of L2 cache, only the XLP208 and T1040 have a third-level cache, and the XLP208's is much bigger. It's

	Freescale QorIQ LS1043A	Freescale QorIQ T1042	Broadcom XLP II XLP208	Cavium Octeon III CN7130
CPUs	4x ARM Cortex-A53	4x Power e5500	2x MIPS64 FC4400	4x cnMIPS64-III
CPU Threads	4 threads	4 threads	8 threads	4 threads
CPU Freq (max)	1.5GHz	1.4GHz	2.0GHz	1.8GHz
L2 Cache	1MB shared L2	256KB L2 per CPU	512KB L2 per CPU	512KB shared L2
L3 Cache	None	256KB shared platform cache	2MB shared L3	None
Accel Engines	SEC 5.4, Quicc Engine, DPAA	SEC 5.4, Quicc Engine, DPAA, PME 2.1 reg-ex	Crypto, reg-ex, RSA, compression	Crypto, reg-ex, compression
L3 Forwarding	10Gbps	7Gbps†	7Gbps†	10Gbps†
IPSec Perf*	5Gbps	5Gbps	3.3Gbps	7Gbps
Reg-Ex Perf (max)	N/A	5Gbps	5Gbps	7Gbps
DRAM Interface	32-bit DDR4-1600	32/64-bit DDR4-1600	64-bit DDR3-1600	32-bit DDR3-1600
DRAM Bandwidth	6.4GB/s	12.8GB/s	12.8GB/s	6.4GB/s
PCI Express	PCIe 3x1 or 4x1 Gen2	4x PCIe Gen2	PCIe 4x1, 2x2, or 1x4 Gen3	3x PCIe Gen3
Ethernet MACs	1x 10GbE + 5x GbE or 6x GbE	5x GbE	2x 10GbE or 8x GbE	1x 10GbE or 8x GbE
Serial ATA	1x SATA III	2x SATA II	None	None
USB	3x USB3.0 + PHYs	2x USB2.0 + PHYs	3x USB	Not disclosed
Power (TDP)	8W	10W	10W†	8W†
IC Process	28nm HPM	28nm HPM	28nm HP	28nm HKMG
List Price (1,000s)	$50†	$50†	$100†	$80†
Production	4Q15 (est)	1Q15 (est)	2Q13	2Q14

Table 2. Freescale's QorIQ LS1043A and three similar 64-bit processors. The lack of ARMv8 competition at this power/performance level pits the LS1043A mainly against the MIPS-based rivals, but new products from AppliedMicro, Broadcom, and Cavium are coming soon. *Maximum IPSec performance with IMIX packets. (Source: vendors, except †The Linley Group estimate)

Price and Availability

Freescale's quad-core QorIQ LS1043A and dual-core LS1023A are scheduled to sample in 1Q15 and begin production in 4Q15. The company has yet to announce pricing, but we estimate they will cost about $50 and $40, respectively, in 1,000-unit volumes. For more information, access www.freescale.com/LS1043A.

likely to hit the cache more often and has plenty of bandwidth available when it does need to go off chip.

All four chips in Table 2 have about the same TDP, by our estimates. In the sibling rivalry between the LS1043A and T1042, we estimate the latter will dissipate more power while driving its additional acceleration hardware and wider DRAM interface.

The LS1043A has the fastest SATA interface and perhaps the fastest USB ports, although Broadcom and Cavium are withholding some details about their USB controllers. Our price estimates make the XLP208 look relatively costly, because its powerful CPUs and large L3 cache require more silicon and its performance is hard to beat. Even if our estimate is high, the XLP208 probably remains the most expensive chip in this group.

Freescale Offers Broadest Product Line

In summary, the LS1043A is a power- and cost-efficient processor that has the features required for its target applications while holding power consumption within an envelope suitable for quiet, reliable fanless systems. It brings the combination of 64-bit ARM muscle, low power consumption, and 10GbE to lower-end products than Freescale and competitors have offered before. By the time the LS1043A ships, however, some dual- and quad-core ARMv8 processors from AppliedMicro, Broadcom, and Cavium may appear.

The dual-core LS1023A offers the option of even lower power dissipation if the target application can get by with less CPU performance. Both new processors show that Freescale isn't merely replacing Power Architecture CPUs with ARM CPUs. The LS1 series still hasn't replicated all the features of the Power-based T1 series. The T1042, as noted above, has an eight-port Ethernet switch, and both the T1042 and T1040 have reg-ex engines. Instead, the LS1 products take advantage of ARM's lower power dissipation to target new applications.

Even so, the new ARMv8 products may lure existing QorIQ customers to switch to ARM. Customers using 64-bit Power CPUs probably don't want to port their code to 32-bit ARM but now can move to 64-bit ARM. Even those customers on older 32-bit Power CPUs may want to jump directly to ARMv8 rather than switching to ARMv7 and then later to ARMv8.

Freescale developed its own Power CPU cores, but it has licensed CPU cores for its initial ARM-compatible products. This approach saves a bundle on CPU development and architecture licensing, and it enables the company to quickly take advantage of ARM's newest designs. That company is already licensing its next-generation 64-bit cores, Artemis and Maya, and we expect Freescale to use these cores in future LS products. Instead of taking the one-size-fits-all approach, Freescale licenses "little" CPUs for its LS1 processors and "big" CPUs for the LS2 products, matching CPU performance and power efficiency to the needs of the target applications.

In the long run, Freescale may develop its own ARMv8 CPU. This approach could be necessary to match the high-end performance of competitors such as Broadcom's Vulcan and AppliedMicro's X-Gene. Whichever way Freescale goes, it already is committing most of its development resources to ARM. The LS1043A and LS1023A will help reinforce one of the company's traditional strengths—the breadth of its embedded product line. ♦

The Linley Group

MICROPROCESSOR report

Insightful Analysis of Processor Technology

THE EDITORIAL VIEW

LONG ROAD FOR ARM IN SERVERS

HP Moonshot Launches First ARM Cartridges

By Linley Gwennap (October 27, 2014)

Now that the first 64-bit ARM processors are shipping in servers, we can take stock of ARM's progress in this market. HP has begun shipping two new ARM-compatible cartridges for its ProLiant Moonshot server, one using AppliedMicro's X-Gene processor (see *MPR 9/1/14,* "X-Gene 2 Aims Above Microservers") and a second using Texas Instruments' KeyStone processors with integrated DSP cores (see *MPR 9/19/12,* "The New Look of DSPs"). HP's published specifications show that the ARM vendors have a long way to go to deliver on their promises of better power efficiency and lower cost.

Moonshot is a 4.3U chassis that can hold 45 cartridges along with power supplies, fans, and Ethernet switches. Each cartridge packs one or more processor chips plus memory, storage, and networking. In addition to the new ARM products, HP offers several cartridges based on x86 processors. All use the same form factor and the same cooling, and they provide the same basic functions, but HP positions each for different workloads, as the table shows. Moonshot offers customers a broad range of choices, and those customers will determine which cartridges are most popular.

HP has not published SPEC benchmark scores for the ARM processors; even if it did, the scores would not be directly comparable to Intel's SPEC ratings, which are inflated through the use of the highly tuned Intel compiler (ICC). When measured using GCC, we estimate the 2.4GHz X-Gene will outperform the 2.4GHz eight-core Atom (Avoton) by 10–20%. Yet the X-Gene cartridge is rated at 55W (typical), versus 27.5W for the Avoton card. The X-Gene power rating includes 32GB more DRAM and 10GbE networking, but even subtracting 10W for these factors, the ARM card uses 60% more power.

HP also offers the m710 cartridge, which uses Intel's Xeon processor. To fit the notoriously power-hungry Xeon into the dense form factor, HP chose a Xeon E3 (Haswell) processor clocked at just 1.8GHz. Even at this speed, Haswell will outperform X-Gene by about 40% on SPECint

Cartridge	m300	m350	m400	m700	m710	m800
Target Workload	Multi-tier web solutions	Web hosting	Memcached, NoSQL	Hosted desktop	Video transcode, big data	Signal processing, real-time analysis
Processor(s)	Intel Avoton	4x Intel Avoton	APM X-Gene	4x AMD X2150	Xeon E3-1284L v3	4x TI 66AK2H
Total CPU Cores	8x Atom v2	4x8 Atom v2	8x ARMv8	4x4 Jaguar	4x Haswell	4x4 Cortex-A15
GPU/DSP	None	None	None	Radeon-128	Iris Pro	4x8 C66 DSP
CPU Frequency	2.4GHz	1.7GHz	2.4GHz	1.5GHz	1.8GHz	1.0GHz
Max DRAM per Processor	32GB	16GB	64GB	8GB	32GB	8GB
Ethernet Ports	2xGbE	8xGbE	2x10GbE	2xGbE	2x10GbE	2xGbE
Relative SPECint*	1.0	2.9†	1.15†	0.9	1.6	0.7†
Total Gflops*	Not applicable	Not applicable	Not applicable	704	870	640
Card Power (typical)	28W	70W	55W	33W	65W	66W
List Price	$1,354	$3,199	$2,249	$1,878	$2,049	$2,899

HP Moonshot cartridges. The company offers several processor cartridges, each with a different application focus and different performance characteristics. The list price is for a single cartridge with maximum DRAM. (Source: HP, except *processor vendors and †The Linley Group estimate)

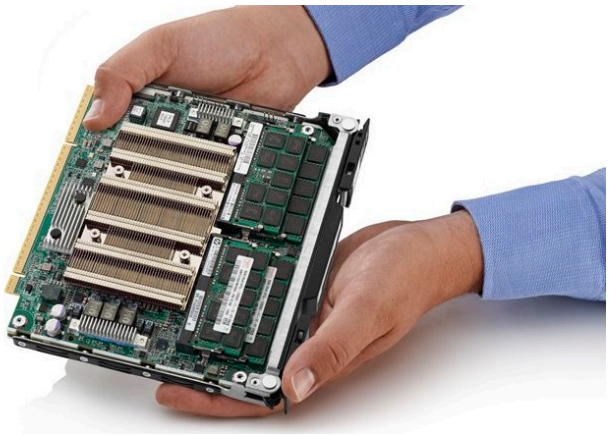

You're in good hands with X-Gene. The AppliedMicro processor fits on a Moonshot board that measures 16.7x18.4cm with a maximum height of 2cm.

while using about 30% more power (after accounting for the difference in DRAM).

Thus, on a general-purpose metric such as SPECint, the ARM cartridge doesn't deliver better performance per watt; in fact, it falls behind the Intel cartridges. X-Gene has an important advantage over either Intel processor, however, in its support for 64GB of DRAM. HP points out that this capability accelerates memory-intensive data-center applications such as web caching (e.g., Memcached) and NoSQL.

The X-Gene cartridge carries a list price of $2,249, according to HP, but if we assume $400 for the extra 32GB of DRAM, the comparable price would be $1,849. This makes the X-Gene card about 10% less expensive than the Haswell card and 35% more than Avoton, although the latter lacks 10G Ethernet. On the basis of these prices, the ARM processor offers no advantage in SPECint per dollar. For applications that benefit from the extra DRAM capacity, however, X-Gene's price/performance is attractive.

The KeyStone cartridge has a different value proposition. Its 1.0GHz ARM CPUs won't keep up with the other cards on standard server benchmarks, but it has 32 DSP cores optimized for signal processing and real-time data analysis. This design is well suited to speech recognition as well as network-virtualization (NFV) functions such as cloud RAN and voice processing. These segments are each small, however. PayPal is using the DSPs for real-time analysis of its transaction stream to help identify fraud. For scientific computing, the KeyStone cartridge delivers 640Gflops (single precision), but both the AMD and Xeon cartridges exceed this rating thanks to their powerful GPUs.

By offering the X-Gene product, HP has exploited a hole in Intel's server-processor lineup: Intel forces its customers to buy up to an expensive Xeon E5 in order to get more than 32GB of DRAM per processor. X-Gene supports twice the memory size of Atom or Xeon E3. It also offers 10G Ethernet, which Atom lacks. KeyStone uniquely integrates DSPs, which none of Intel's x86 processors include.

None of these advantages, however, comes from the ARM instruction set or ARM CPUs. Contrary to original promises, the ARM processors have failed to demonstrate any power-efficiency or cost benefit relative to x86. Their advantages stem from Intel's decision to not offer certain configurations, either to maintain higher margins or to avoid niche markets.

If it becomes concerned about the competition, Intel can close the holes in its lineup at any time. To move out of workload-specific niches, future ARM-compatible server processors—such as AMD's A1100, AppliedMicro's X-Gene 2, and Cavium's ThunderX—must deliver much better performance per watt and performance per dollar on general-purpose applications. Only in this way will they address the needs of the broader server market. ♦