The Linley Group
MICROPROCESSOR REPORT
www.MPRonline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

MIPS APTIV CORES HIT THE MARK

New Family Shows Highest CoreMark/MHz for Licensable CPUs

By J. Scott Gardner (May 28, 2012)

The CPU core wars have heated up again, and microarchitecture aficionados will watch this new conflict with keen interest. MIPS is trying to jump back into the fray with three new MIPS32 cores that may hinder ARM's race for dominance. MIPS has improved the firepower across its entire range of CPUs, from low-cost microcontrollers to high-performance superscalar designs with multiple cores. These new CPU cores will help the company defend its market position in digital home and data networking while also strengthening its appeal to customers designing SoCs for mobile applications. MIPS faces a number of business challenges before it can seize the high ground in mobile markets, but the technical features of its Aptiv product family should make the market for licensable CPU cores much more competitive.

In subtle homage to the ARM marketing team, which segments its Cortex product families into -A, -R and -M designations, the Aptiv family trifurcates into "**m**icro," "**i**nter," and "**p**ro" to address every MIPS32 market. These new MIPS CPU cores extend the existing 14K, 24K/34K/1004K, and 74K/1074K families into the respective elements of the unified Aptiv brand. The new naming scheme may help customers and press more easily distinguish between the three different microarchitectures.

The microAptiv CPU core uses a simple in-order design to target low-cost embedded applications in much the same way that ARM uses its Cortex-M family. MIPS applies the interAptiv moniker to describe its newest CPU core that supports simultaneous multithreading (SMT), a feature not offered by ARM (Intel's Atom processor, however, has comparable SMT capability). The proAptiv CPU is the crown jewel in MIPS's three-pronged product launch, as the company believes its new out-of-order design will reach performance levels that match or exceed the estimated performance of ARM's Cortex-A15. The microAptiv CPU is available now, whereas production RTL for interAptiv and proAptiv are both expected by 3Q12.

Moving Beyond Dhrystone

Although no single benchmark can provide all of the information needed to evaluate a CPU, we believe that EEMBC *CoreMark* offers a better representation of CPU core performance than the nearly 30-year-old Dhrystone benchmark. MIPS published a certified EEMBC score for its high-end proAptiv CPU and is reporting noncertified scores for microAptiv and interAptiv. When scaled for frequency, the top-end proAptiv sets a new EEMBC single-core record for licensable CPUs, scoring 4.42 CoreMarks/MHz on preproduction RTL implemented in an FPGA. The only vendor with a higher published score is Intel, which scored roughly 5.1 CoreMarks/MHz (per core)

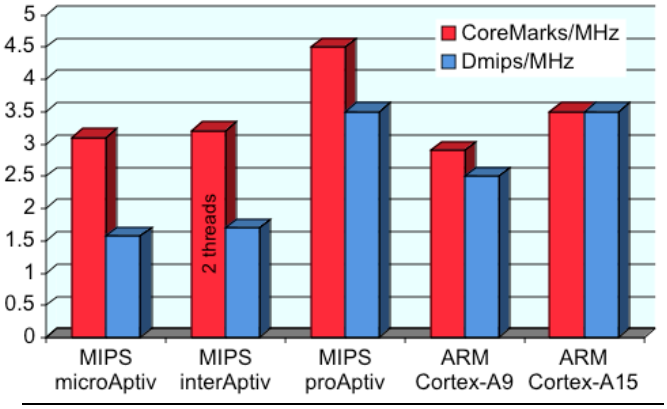

Figure 1. Performance comparison for MIPS's Aptiv family. The Aptiv CPUs outperform competing ARM CPUs on CoreMark but not Dhrystone. (Source: vendors, except CoreMark scores for ARM are Linley Group estimates)

for PC processors such as the Core i5 and Core i7. MIPS recently pushed its proAptiv CoreMark/MHz rating to 4.5 and plans to have EEMBC certify the new score.

Figure 1 shows CoreMarks/MHz and Dmips/MHz for all three Aptiv products and for ARM's top-end cores. ARM reported CoreMark scores of 2.08 for Cortex-A8 and 2.88 (per core) for Cortex-A9, but it has not submitted any EEMBC scores since 2010 and prefers to highlight Dhrystone for new CPUs such as Cortex-A7 and Cortex-A15. We estimate that Cortex-A7 should score 2.5 CoreMarks/MHz, and Cortex-A15 will reach 3.5 CoreMarks/MHz.

We include Dmips for comparison while noting that a CoreMark test executes over 2,000 times as many instructions as the venerable Dhrystone. Like Dhrystone, however, the CoreMark test is intentionally designed to fit into most L1 caches. CoreMark is heavily influenced by the performance of the branch predictor in microarchitectures with long pipelines; MIPS references branch-prediction performance as one of the microarchitectural features that contributes to proAptiv's excellent CoreMark scores. The CoreMark test can dispatch data-independent threads and scales linearly with the number of cores. In the case of interAptiv, the benchmark uses two threads and multithreading, giving the CPU a higher score.

Tiny MicroAptiv Targets Microcontrollers

The microAptiv CPU is a tiny, low-power MIPS32 core targeting high-volume embedded processors. We expect this market for licensed embedded cores to exceed eight billion chips per year by 2016. Applications for microcontrollers (MCUs) are driving this growth, most notably in products such as flash memory and smartcards. Many of these devices will require multiple CPU cores in each chip.

For a microcontroller, almost any 32-bit CPU core will meet the performance requirements. Die size and power consumption are much more important factors, which is why MIPS based microAptiv on the simple five-stage microarchitecture of the 14K family. The instruction set supports microMIPS, which implements instruction-compression technology (similar to ARM's Thumb) and allows mixing of 32-bit and 16-bit opcodes (see *MPR 11/16/09,* "MicroMIPS Crams Code").

Compared with the 14K, microAptiv's primary design enhancement relates to the DSP application-specific extension (ASE), which is already available on the high-end 74K/1074K and adds a dedicated DSP pipeline with support for dual 16-bit and quad 8-bit SIMD instructions. These instructions include single-cycle MAC operations (dual 8x8, dual 16x16, dual 16x8, and 32x32). The extra DSP hardware allows better performance on digital filters and other math-intensive applications. Compared with an M14K, most DSP operations require half as many cycles; a 16-bit FIR filter (k=32) completes in 75% fewer cycles.

The CoreMark benchmark includes matrix math and other SIMD operations that exploit the MAC hardware, but MIPS notes that the base MIPS ISA includes MAC instructions. Thus, the company attributes the relatively high CoreMark/MHz score to the five-stage pipeline and to the MIPS instruction set's ability to use branch delay slots to hide latency. Although the five-stage pipeline limits the maximum speed, scaling the CoreMark score to frequency emphasizes pipeline efficiency compared with deeply pipelined multi-issue CPU designs. These factors should also improve Dhrystone performance, however, and we are concerned about the large gap between microAptiv's Dhrystone and CoreMark scores. It remains to be seen if real application performance will match up to the high CoreMark score.

For applications that don't need deterministic performance, the microAptiv system architecture can be configured with a full MMU and caches. Alternatively, the CPU can be configured to directly address SRAMs as scratchpad memory while using a simple memory-protection unit to support a real-time operating system. Using a speed-optimized 65nm G process, microAptiv occupies less than 0.33mm^2 (excluding memory) while consuming less than 0.10mW/MHz of active power. In this trailing-edge process, the CPU reaches 500MHz and achieves 3.1 CoreMarks/MHz. Many applications will be able to take advantage of the extra instructions per cycle (IPC) to operate at a lower frequency while saving power.

	1004K	InterAptiv	1074K	ProAptiv
Max Clock Rate*	1.1GHz	1.1GHz	1.25GHz	1.2GHz
Max CPUs/Cluster	1x–4x SMP (34K)	1x–4x SMP	1x–4x SMP (74K)	1x–6x SMP
Max Instr Dispatch	1 per cycle†	1 per cycle†	2 per cycle†	3 per cycle†
Pipeline Depth	9 stages	9 stages	17 stages	16 stages
L1 Cache	0–64KB L1†	0–64KB L1†&	0–64KB L1†	32KB–64KB L1†
TCM	0–1MB TCM†	0–1MB TCM†	0–1MB TCM†	0–1MB TCM†
Extensions	DSP (2x 16-bit MACs/cycle)	DSP (2x 16-bit MACs/cycle)	DSP (2x 16-bit MACs/cycle)	DSP (2x 16-bit MACs/cycle)
CoreMarks/MHz	2.97 CM/MHz† (2 threads)	3.2 CM/MHz† (2 threads)	2.55 CM/MHz†	4.5 CM/MHz†
CPU Cluster System Bus‡	64-bit or 256-bit system OCP interface out of L2$ to system; 64-bit coherent I/O OCP interface			
RTL Released	2Q08	2Q12	3Q10	2Q12
First SoC Samples	2010	2013#	2011#	2013#

Table 1. MIPS CPU core comparison. The interAptiv design provides minor improvements over the older 1004K, whereas proAptiv is a major upgrade to the 1074K, adding a third instruction slot. *40nm G process, worst-case timing, 12-track SVT library, including 32KB L1 caches and TLB/MPU; †per CPU. (Source: MIPS Technologies, except ‡The Linley Group estimate)

Extends Multithreading Lead

Table 1 compares the interAptiv and proAptiv cores to current-generation multicore MIPS products. The interAptiv device is based on the 1004K (see

MPR 02/27/06, "MIPS Threads the Needle"). Surprisingly, few general-purpose embedded processors support multithreading, even though the design technique can improve latency tolerance in a power-efficient manner. Multithreading requires less energy than approaches that use larger, faster caches and speculative memory fetches to reduce latency. When an execution thread encounters a long-latency operation, the CPU can context switch to a different execution thread, allowing the pipeline to keep working. The tradeoff is an increase in die area (compared with a single-thread design) and greater design complexity.

Multithreading is a good design tradeoff to maintain performance when using a simpler, lower-power microarchitecture and memory system. For example, Intel's Atom capitalizes on two-way multithreading (Hyper-Threading) to ping-pong between instruction thread queues whenever the in-order pipeline stalls. To deliver these advantages, multithreading requires multithreaded software, or at least multiple simultaneous applications.

Using the MIPS multithreading ASE from the 34K/1004K family, interAptiv supports more-sophisticated multithreading than Intel's Atom. Figure 2 shows a block diagram of a single core in a design based on this processor. Each interAptiv core supports two virtual processor elements (VPEs), and each VPE replicates all of the privileged-mode state in hardware, allowing an operating system to treat an interAptiv core as if it were a pair of MIPS32 CPUs. The two VPEs can be assigned up to nine thread contexts (TCs), which operate as multiple threads in a single user-mode process. The interAptiv architecture supports up to four cores in a coherent SMP configuration.

MIPS provides flexibility to allow system-software developers to manage threads in the most efficient manner. A system can incorporate fine-grained multithreading in which hardware switches between threads every CPU cycle using a round-robin or priority-driven approach. Alternatively, a CPU can switch threads only when the current thread stalls. The system also supports quality of service (QoS) hardware to guarantee that a real-time thread gets enough CPU cycles. The CPU must contain enough registers to hold the state for all of the thread contexts, so the number of VPEs and TCs can be configured to allow the designer to trade off performance and die area.

The interAptiv design uses a multiplexer to select the thread context from which to fetch instructions. A fetched instruction queues up in a buffer corresponding to its TC, and an additional fetch stage dispatches instructions from these buffers, rotating among the TCs each cycle. If a thread is stalled, its TC is skipped until the stall is resolved. Multithreading creates other complications in the CPU implementation that are not necessarily visible to the programmer. If one TC stalls, interAptiv can roll back partially completed dependent instructions and send instructions from a different TC down the pipeline. The interAptiv design can be configured with a full IEEE 754 floating-point unit that can also be multithreaded.

Price and Availability

MIPS plans to deliver proAptiv and interAptiv RTL in June 2012. The microAptiv cores are available now. The company does not publicly disclose licensing fees or royalties. For more information about the Aptiv family, access *www.mips.com/aptiv.*

Although we must continue waiting for a licensable MIPS64 core, MIPS has expanded the MIPS32's user address space by adding segmentation to the TLB structure. The interAptiv design supports enhanced virtual addressing (EVA) to increase this address space to greater than 3GB. The data TLB always has eight entries, but the SoC designer can configure the instruction-TLB size (4–12 entries) and joint-TLB size (16–64 dual entries) to support the required number of threads. The SoC designer can also allocate separate scratchpad RAMs (up to 1MB each) for instructions or data store. All data memories and caches have optional ECC protection.

We estimate that multithreading adds about 10% to the area of a MIPS core (including 32KB L1 caches and two VPEs), but it provides a 20–40% performance gain on multithreaded applications (depending on how well the number of VCs matches the number of execution threads). Single-threaded applications won't receive any speedup, but interAptiv is an architecture targeting embedded applications that need efficient processing of multiple threads. The multithreaded performance gain is reflected in the reported EEMBC score of 3.2 CoreMarks/MHz for two threads running on a single core.

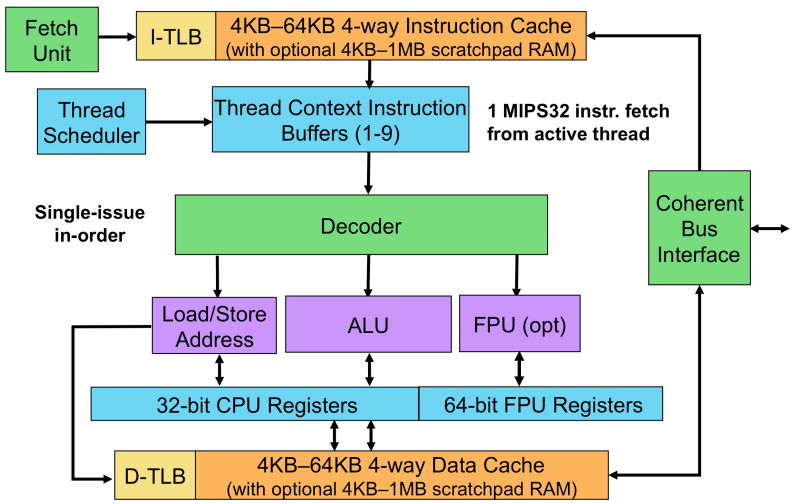

Figure 2. MIPS interAptiv block diagram. The multithreaded microarchitecture supports up to two virtual processing elements and up to nine thread contexts.

MIPS expects the interAptiv core to reach speeds of 1.1GHz in TSMC's 40nm G process (worst-case timing). The design should be power and area efficient, although MIPS has not completed its power measurements. The company estimates that a power-optimized triple-core interAptiv will fit into the same 3.6mm^2 used by a similarly configured dual-core Cortex-A9 (excluding Neon).

MIPS Takes a ProAptiv Approach

Instead of ceding the high-performance markets in networking, digital home, and mobile to competitors such as ARM, MIPS is hoping that proAptiv will get the company back in the performance race. The proAptiv family builds on its architectural heritage from the 74K/1074K's out-of-order superscalar design (see *MPR 10/11/10,* "MIPS Boosts Multiprocessing"). MIPS has focused on delivering leading performance while avoiding an excessive increase in area or power. On the basis of its CoreMarks/MHz score, proAptiv is 75% faster than the 1074K. The proAptiv design team achieved this big performance jump without resorting to brute-force approaches that increase performance by simply expanding the width of the machine.

Figure 3 shows the proAptiv design, which, like the 1074K, is an out-of-order machine that uses a deep pipeline and multiple execution units to achieve high clock speed and parallelism. The pipeline is one stage shorter than the 1074K's.The 1074K can dispatch up to two MIPS32 instructions per cycle, but proAptiv uses an instruction-bonding technique to dispatch up to three MIPS32 instructions per cycle. After decoding instructions, the CPU can bond, or fuse, certain pairs of adjacent instructions that do not contain interdependencies. These bonded pairs can be scheduled and dispatched as a unit, reducing the complexity of the scheduler and enabling proAptiv to use essentially the same scheduling algorithms as the two-issue 1074K.

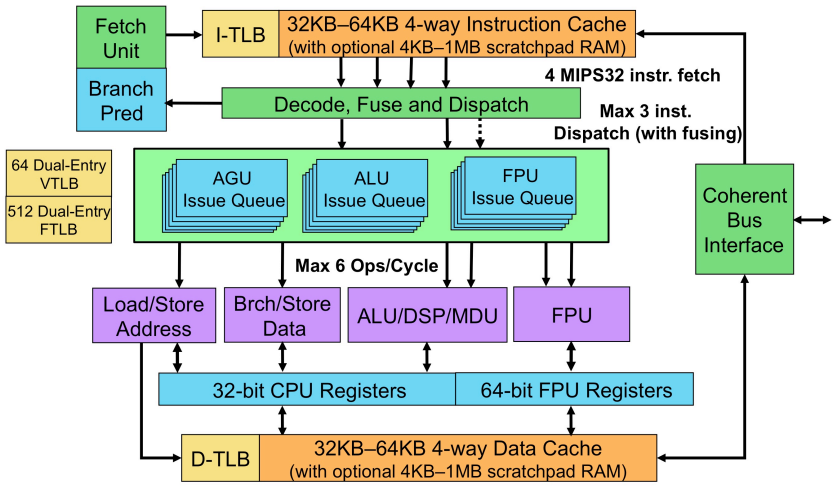

Figure 3. MIPS proAptiv block diagram. The out-of-order microarchitecture can issue up to four integer operations and two floating-point operations per cycle.

The MIPS architects believe that program dependencies and issue restrictions yield diminishing returns on architectures with wider dispatch. In an out-of-order architecture, the decoded instructions (with renamed registers) fill up scheduler queues that must stay filled with useful work as the execution units become available. If issue limitations keep the scheduler queues from emptying, there is no need to decode instructions any faster.

MIPS believes that its branch predictors, quad-instruction fetch, and bonded dispatch ensure that instruction supply is not a bottleneck. This design approach allows proAptiv to save power and area while delivering the performance of a wider machine. The design's high CoreMark scores seem to validate this design tradeoff.

Improved Branch Prediction

The MIPS designers are especially happy with proAptiv's branch-prediction performance, which the company believes is critical to the kind of real workloads reflected in CoreMark; this benchmark uses modern code with realistic branch behavior. Owing to such a long pipeline, any mispredicted branches incur a high performance penalty. More importantly, fetching unnecessary instructions (four at a time) wastes energy.

MIPS didn't disclose many details about its branch predictor, though proAptiv uses a branch history table (BHT) that is more than eight times larger than that of the 1074K. The proAptiv design also implements a multilevel branch target buffer (BTB) to store addresses of frequently taken branches. The BTB works in conjunction with the BHT to reduce the misprediction penalty—specifically, the time taken to redirect the fetch unit after a taken branch is encountered. Minimizing the misfetch penalty is particularly important for indirect jumps, since the branch target is unknown at fetch time. Indirect jumps frequently appear in applications that use object-oriented programming languages (C++, Java, and so on).

By ensuring a steady stream of decoded (and renamed) instructions for the out-of-order scheduler queues, proAptiv's execution units can take better advantage of the parallelism that allows a peak of four integer and two floating-point operations per cycle. The processor can simultaneously issue four integer operations every cycle, including a load/store, a branch, and two ALU operations. Multiplies, divides, and DSP operations are accelerated by new (undisclosed) algorithms. For example, the hardware integer divider runs twice as fast as the divider in the 1074K. These complex instructions share issue ports with the two ALU pipes, saving power but avoiding performance loss because multiply/

divide operations do not block the pipe: the simple ALU operations can issue underneath the complex operations. Most ALU operations execute in a single cycle.

The proAptiv CPU implements many of the same architectural elements used for interAptiv, such as the new coherence manger, power management, advanced trace interface, four-way caches (up to 64KB I/D), and support for scratchpad RAM (up to 1MB). Both architectures introduce a host of features to reduce power consumption. Some of these power-saving options include automatic shutdown of core clocks during bus requests, intelligent way selection in the four-way instruction cache, and support for lower-power 32-bit accesses to the data cache.

MIPS developed a second-generation coherence manager that works with both interAptiv and proAptiv. The new design enables a faster coherent system bus running at a one-to-one ratio with the CPU pipeline speed while supporting a larger L2 unified cache (256KB to 8MB). The new cache controller also runs at the CPU pipeline speed, but it adds configurable wait states for L2-cache access. MIPS estimates its new memory and bus architecture will reduce average system latency from approximately 24 cycles to approximately 11 cycles, compared with the previous system design for the 1004K/1074K.

The company predicts that proAptiv will operate at 1.1GHz in TSMC's 40nm G with worst-case parameters. This speed assumes a design that targets low-power embedded applications and omits LVT cells, voltage overdrive, and a high-performance standard cell library. For applications in digital home and networking, MIPS expects proAptiv will exceed 2.0GHz in 40nm G using voltage overdrive and other techniques to gain speed at the expense of power and area. The company has withheld detailed area and power parameters, but MIPS expects the CPU to be approximately half the size of a Cortex-A15 implemented in the same process node and with similar parameters.

Muscling Up Against ARM

Although many of the features carry over from previous generations, the Aptiv family is a complete overhaul of MIPS's entire product portfolio. By launching three product families at once, the company makes a strong statement about continuing to advance its product lines (despite recent rumors about it being for sale). The unified product-naming scheme will also simplify the message and enable customers to more easily draw direct comparisons to ARM CPUs.

The microAptiv family has some technical advantages over Cortex-M, especially when executing DSP-oriented applications. In many microcontrollers, memory and I/O dominate the device, so switching suppliers to get a new CPU core may be easier. Although MIPS has a long-standing relationship with Microchip as a semiconductor supplier of microcontrollers, ARM's Cortex-M is more broadly adopted by companies such as Texas Instruments, NXP, and many others. The microAptiv core can also compete for deeply embedded applications that include the huge smartcard and flash-memory markets, where the instruction set is less important.

The interAptiv family represents an area where MIPS has a unique offering with its multithreaded core. Owing to the rapid proliferation of multithreaded applications, MIPS can use interAptiv to reintroduce itself to customers who may have overlooked the 34K/1004K families. Intel's multithreaded Atom family has helped validate the importance of multithreading for embedded processors, including designs that require multitasking. MIPS may be able to create a similar technical value proposition as it reintroduces multithreading.

The proAptiv family is the most important product in the new launch, since MIPS needs to reestablish itself as a performance leader in the market for general-purpose RISC cores. For most of its existence, the company has been the leading supplier of high-performance CPU designs, whereas ARM has supplied simple designs that consume little power. Now, ARM is designing big CPU cores (Cortex-A15 and Atlas) to take the performance lead. But the laws of physics haven't changed, and these larger CPUs will incur higher silicon and power costs.

Table 2 compares MIPS's proAptiv and ARM's Cortex-A15, including our estimates for the CPU parameters that ARM has yet to disclose. Several semiconductor vendors have announced a target frequency of 2GHz for processors based on Cortex-A15, but these products use 28nm technology. We believe that both cores would achieve similar clock speeds as long as the chip designers are equally talented and use an identical library and iden-

	MIPS ProAptiv	ARM Cortex-A15
Instruction Set	MIPS32 R3	ARMv7
Max Coherent SMP	6 CPUs	4 CPUs
Max CPU Speed	2.0GHz*	2.0GHz*
Max Instr Dispatch	3 per cycle	3 per cycle
Reordering?	Yes	Yes
Pipeline Depth	16 stages	18 stages
L1 Caches	32KB or 64KB	32KB
L2 Cache	Integrated, up to 8MB	Integrated, up to 4MB
TCM	Up to 1MB	None
Special Features	DSP, opt. FPU, SMP, EVA, CorExtend, PDtrace	DSP, opt. FPU, Neon, SMP, TrustZone, LPAE, PTM
MACs/Cycle	2x16-bit	2x16-bit
CoreMarks/MHz	4.5 CM/MHz	3.5 CM/MHz
DMIPS/MHz	3.5 DM/MHz	3.5 DM/MHz
Interfaces	1x 256-bit OCP	1x 128-bit Amba4
Die Area (base CPU)	1.0mm^2*	2.0mm^2*
RTL Release	2Q12	4Q10

Table 2. Comparison of fastest announced cores from MIPS and ARM. *Soft cores synthesized in a standard-voltage 28nm LP process using a 9-track SVT library. (Source: vendors, except *The Linley Group estimate)

tical process parameters. But Cortex-A15 will benefit from the collective expertise of and competition between SoC designers from well-resourced companies such as Apple, Nvidia, Samsung, and TI. MIPS will likely be working with customers that won't make the same investment to push the architecture to its limits.

We've included an estimate for Dhrystone mips, since ARM assures us that many customers ask for this performance metric. We believe, however, that the CoreMark scores are more meaningful, and by this metric, proAptiv should be 33% faster than Cortex-A15. Neither benchmark, however, reflects the performance for ARM's Neon SIMD unit, which ARM expects to yield significant improvements in multimedia applications.

MIPS expects the proAptiv CPU to be half the size of Cortex-A15 in the same process, on the basis of ARM's statements about the A15 being twice the size of the A9. The bottom row of Table 2 contains a more important metric: proAptiv will be released to customers a year and a half after ARM ships Cortex-A15. ARM is already well along in the design process of its next-generation high-performance CPU, code-named Atlas, so any performance advantage that MIPS achieves may be temporary.

For now, the MIPS design team seems to have taken the performance lead away from ARM, and it deserves credit for this accomplishment. If MIPS can execute on its product strategy, it will have a strong trio of CPU cores that will make the CPU-IP markets more competitive and provide some answers to the vexing questions about its business strategy and long-term viability. ♦