Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

The Rise of Licensable SMP

New PowerPC 476FP Processor Core Challenges ARM and MIPS

By Tom R. Halfhill {2/16/10-01}

Four wild horses are a small herd, but four horses harnessed are a team. Getting them to work together requires training, coordination, and a skilled driver. The same principles apply to multicore SoCs. Two or more cores may share the same silicon, but symmetric multiprocessing (SMP) is the harness that helps them share a workload.

As embedded applications demand more performance, we're seeing more interest in licensable microprocessor cores specifically designed for SMP. Of course, chip designers can use any processor cores for this purpose, but only a few cores have the built-in features, coherency control, and coherent debugging that make SMP easier to implement. ARM introduced the ARM11 MPCore in 2004, followed by the Cortex-A9 MPCore in 2008 and Cortex-A5 MPCore last year. MIPS Technologies introduced the MIPS32 1004K Coherent Processing System in 2008. All these cores are licensable 32-bit embedded processors supporting two-, three-, or four-way SMP with coherent memory systems.

Now IBM is joining the race with the new PowerPC 476FP. It, too, is a licensable 32-bit embedded-processor core. It's a wide superscalar design intended for high-performance embedded applications. Top speed exceeds 2.0GHz, or 1.6GHz under worst-case conditions. It has an FPU, and it supports coherent SMP systems with up to eight cores—twice as many cores as ARM or MIPS. Developers can link two eight-core clusters to build a 16-core design, although memory coherency is limited to the cores in each cluster. Unlike the ARM and MIPS soft cores, the PowerPC 476FP will be licensable as a hard macro or as a synthesizable model.

IBM designed the PowerPC 476FP in collaboration with LSI Corp. (formerly LSI Logic), which is also the first licensee. LSI plans to use the new processor in next-generation multicore chips for networking, communications, and storage applications. LSI's contribution to the project was a configurable L2 cache tightly coupled to the processor. IBM designed an improved CoreConnect bus controller that can snoop the L1 and L2 caches to maintain coherency.

Both the hard and soft versions of the PowerPC 476FP core are scheduled for general availability by the end of this year. The hard core is engineered for fabrication in IBM's 45nm silicon-on-insulator (SOI) process—which, of course, limits manufacturing to IBM's fab. The soft core will be portable to other CMOS processes and foundries.

Wide Superscalar Execution

The PowerPC 476FP processor conforms to version 2.05 of the Power instruction-set architecture (ISA), with Book III-E embedded extensions. Power ISA v2.05 dates to 2007 and was superseded by Power ISA v2.06 a year ago. Among the features missing from v2.05 are new vector floating-point instructions. However, Book III-E does add support for hypervisors and virtualization in single-core and multicore systems. The PowerPC 476FP's lack of support for the latest Power ISA v2.06 suggests that IBM locked down the core's design in 2008 or early 2009, before v2.06 was adopted by Power.org, the Power Architecture governing body.

Although the original "PowerPC" name for this CPU architecture has been largely abandoned in favor of "Power Architecture," IBM continues to use "PowerPC" as a sub-brand for

embedded-processor cores. (See *MPR 8/21/06-01*, "The New Power Architecture.")

At first glance, the PowerPC 476FP is such a muscular processor that one wonders why any embedded application would need more than one of them. The PowerPC 476FP can issue four instructions at once through its nine-stage integer pipelines, using dynamic branch prediction, out-of-order execution, and speculative execution. As many as 32 instructions can be "in flight," passing through various pipeline stages on their way toward completion. Speculative prefetching helps keep this voracious processor fed with instructions.

Separate pipelines handle simple ALU instructions, complex ALU instructions, multiply/divide operations, branches, and load/store instructions. The multiply/divide pipeline accelerates some multiply and multiply-accumulate (MAC) operations for signal-processing tasks. (However, the PowerPC 476FP's DSP capabilities are rather limited.) Thanks to its unusually wide superscalar microarchitecture, the PowerPC 476FP delivers 2.5 Dhrystone mips per megahertz—the same as ARM's Cortex-A9 MPCore, and significantly more than the MIPS 1004Kf (1.56Dmips/MHz).

IBM's FPU supports single- and double-precision floating-point math to IEEE 754-1985 standards and Power ISA v2.05 specifications. It has separate pipelines for arithmetic and load/store operations, avoiding stalls while waiting for data. (The FPU shares the load/store pipeline with the ALUs.) The floating-point arithmetic pipeline is nominally six stages long, though some complex operations (division, denormal) have longer latencies.

The L1 instruction and data caches are each 32KB and four-way set-associative, with parity protection and 32-byte lines. Programmers can lock individual cache lines to keep the cache controller from evicting critical instructions or data. The write-through data cache is nonblocking. Cache accesses require two clock cycles and are pipelined.

Large, Fast L2 Cache

As mentioned above, LSI designed a configurable L2 cache controller for the PowerPC 476FP. It supports caches of 256KB, 512KB, or 1MB, with four-way set associativity. IBM hasn't disclosed the size of the L2 cache in the first chip design. Note that an L2 cache is required to use the newly enhanced version of IBM's CoreConnect on-chip bus, PLB6 (Processor Local Bus, version 6). PLB6 supports SMP coherency.

Two interfaces connect the L2 cache to the processor core. There's a 256-bit read interface shared by the L1 instruction and data caches, plus a 128-bit write interface from the L1 data cache. The read interface can fill an L1 cache line in a single cycle while predecoding the instructions. Each interface can run at core-to-bus clock ratios of 2:1, 3:1, and 4:1. Therefore, the maximum theoretical I/O bandwidth is 48MB/s, assuming a 2.0GHz core and 1.0GHz bus clock. All L2 cache transactions are parity protected with error-correction codes (ECC).

Figure 1 is a block diagram of the PowerPC 476FP. Clearly, this is a well-endowed embedded processor. In addition to the features already described, it has a memory-management unit (MMU) with translation-lookaside buffers (TLB), additional buffers to aid branch prediction, and a set of timers. The unified TLB has 1,024 entries, and the MMU can read four entries at a time. Page sizes can range from 4KB to 1GB. Separate eight-entry TLBs for instructions and data supplement the unified TLB.

In addition, the PowerPC 476FP has separate user and supervisor modes. With all these features, it's ready to run sophisticated embedded operating systems that manage large amounts of memory. The 42-bit address space for real memory supports up to four terabytes (4TB), and the virtual-memory address space is 49 bits (512TB). Even without its SMP capabilities, the PowerPC 476FP would seem to have more in common with server processors than with embedded processors. Its memory-management features will be welcome for high-end networking and storage applications.

Of course, the trade-offs for all these features are size and power. The PowerPC 476FP core is rather large: 3.6mm^2, not including an L2 cache. That is IBM's estimate, assuming fabrication in Big Blue's 45nm SOI process with eight layers of metal. Most other 32-bit embedded-processor cores would measure 1.0mm^2 or less in a similar process. And 3.6mm^2 is for the hard macro; the synthesizable version will almost certainly be larger.

Thanks to extensive clock gating and other optimizations, the hard core dissipates only 1.6W at 1.6GHz—a mere 1.0mW per megahertz. (IBM specifies 1.6GHz as the low-end process-voltage-temperature point. The operating range is 0.9V to 1.1V, –40° to +125°. Speed-sorted samples can exceed 2.0GHz.) One milliwatt per megahertz is remarkably low power for a PowerPC processor. However, as we'll see in a moment, it's more power hungry than competing processors from ARM and MIPS.

SMP Is the Defining Feature

As powerful as the PowerPC 476FP is, customers will probably want this processor for its SMP capabilities. Those features add some logic that cannot be removed from the core if the target application needs only one processor. Consequently, we expect the PowerPC 476FP to compete directly with the ARM Cortex-A9 MPCore and MIPS 1004K Coherent Processing System. To a lesser degree, it also competes with the new ARM Cortex-A5, a less powerful SMP core. Like the PowerPC 476FP, the ARM and MIPS cores are licensable 32-bit processors expressly built for multicore SMP. And they are synthesizable, portable to any fabrication process at any foundry.

The Cortex-A9 MPCore is the most complex processor core ARM has yet designed. It is ARM's only superscalar processor and is the only one with out-of-order execution and speculative execution. It can decode two instructions

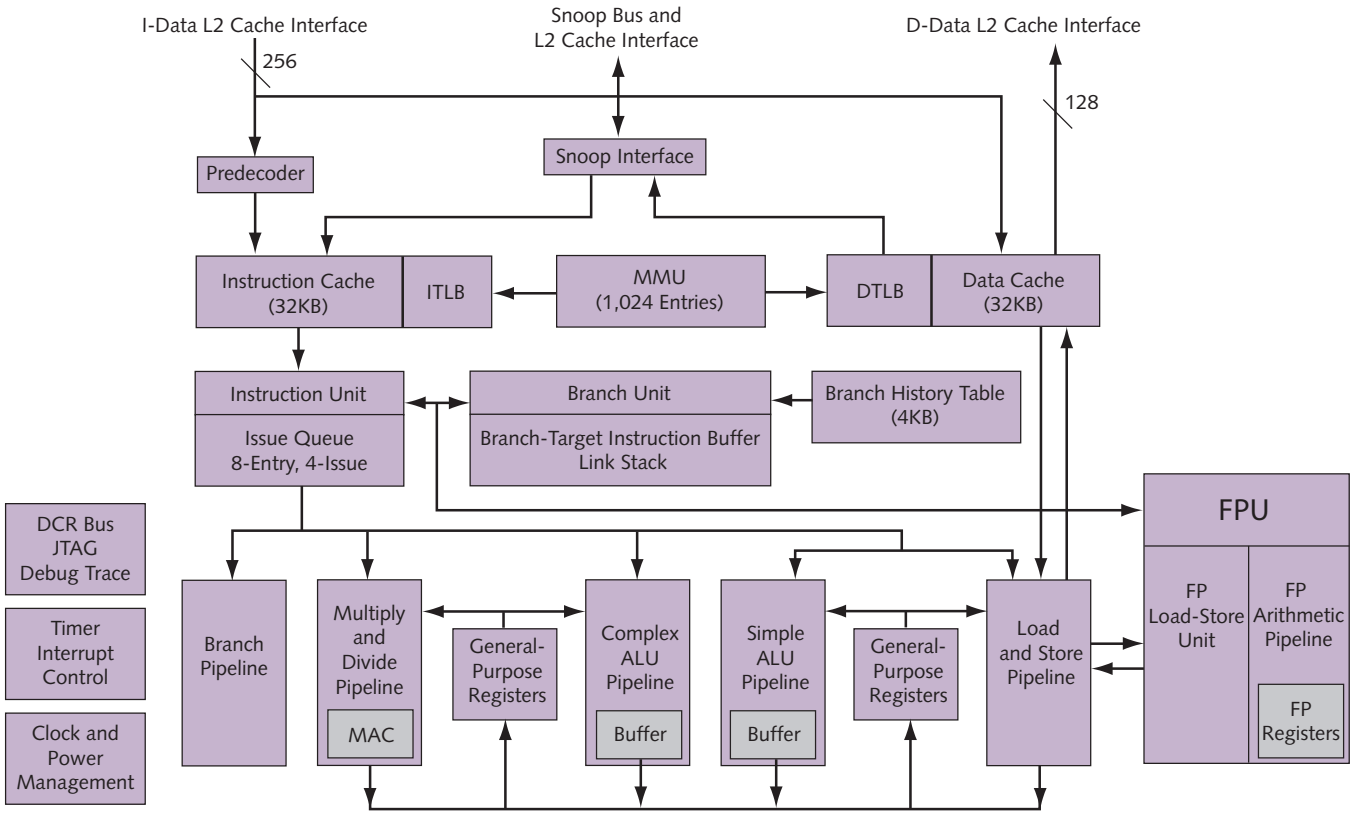

Figure 1. IBM PowerPC 476FP block diagram. This is one of the most complex 32-bit embedded-processor cores yet seen. It has four-issue superscalar execution, dynamic branch prediction, out-of-order execution, and speculative execution. The double-precision FPU has separate pipelines for arithmetic and load/store operations. Note the 256-bit and 128-bit interfaces to the configurable L2 cache, which has ECC and can be as large as 1MB. There are four built-in timers: a time base, a decrementer, a fixed-interval timer, and a watchdog. Big-endian memory addressing is standard, but the processor also supports little-endian addressing.

per clock cycle and issue four instructions per cycle. Integer pipelines are eight stages deep, although instructions need 9 to 11 clock cycles to traverse them. Branch prediction is dynamic.

Like other Cortex A-series processors, the Cortex-A9 has an MMU, so it can run sophisticated virtual-memory embedded operating systems, such as Linux, Windows CE, Symbian, and Android. For security, the Cortex-A9 has ARM's TrustZone protected-execution mode. To accelerate programs written in Java and Microsoft's .NET languages, it has ARM's Jazelle RCT extensions and optional Jazelle DBX extensions. (See *MPR 7/11/05-01*, "ARM Strengthens Java Compilers.")

For multimedia applications and light-duty signal processing, the Cortex-A9 has ARMv6 SIMD instructions. ARM's Neon extensions are optional, as is an L2 cache and a 32/64-bit FPU. Like other processors in the Cortex-A family, the Cortex-A9 adheres to the ARMv7-A ISA. ARMv7-A includes the latest 16-bit Thumb-2 instructions for better code density, plus the Thumb-2EE instructions associated with Jazelle. For backward compatibility, ARMv7-A supports the original 16-bit Thumb instructions, too. (See *MPR 6/17/03-02*, "ARM Grows More Thumbs.")

The Cortex-A9 supports two-, three-, or four-way coherent SMP. As mentioned above, the PowerPC 476FP supports coherent SMP with up to eight cores, or two coherent clusters of eight cores. IBM has a major advantage in SMP scalability, for embedded systems that need that much horsepower.

SMP Features: IBM's Advantage

ARM's coherent on-chip bus is its own 64-bit AMBA-3 AXI, configurable on the Cortex-A9 with one or two 64-bit interfaces. IBM's coherent on-chip interconnect is the aforementioned CoreConnect PLB6. As implemented here, PLB6 has separate 128-bit read and write channels running at half the core frequency—800MHz in a 1.6GHz hard core. The maximum point-to-point bandwidth would be 25.6GB/s at that bus frequency. The bandwidth of the Cortex-A9's on-chip bus likewise depends on the clock frequency, but the AXI bus must run faster to make up for its narrower channels, relative to IBM's PLB6.

In addition to scalability, the PowerPC 476FP has another SMP advantage over the Cortex-A9. Each IBM core can have its own snooped L2 cache, whereas the Cortex-A9 shares a single unsnooped L2 cache among all cores. An

SMP system built with the PowerPC 476FP can have much more L2 cache.

Otherwise, their SMP capabilities are similar. Both processors have a global interrupt controller that routes interrupts to the appropriate cores, and each core can respond individually, without interrupting other cores. Both processors have coherent I/O control. Both allow developers to debug an SMP design with memory-coherent debuggers and tools.

If the target application requires some signal processing or media processing, the Cortex-A9 is better equipped. Whereas the PowerPC 476FP has a few MAC instructions, the Cortex-A9 has ARMv6 SIMD instructions and the optional Neon coprocessor. Neon extends the ARMv7 ISA with more than one hundred 16- and 32-bit instructions, plus additional registers and pipelines. (For more about Neon, see *MPR 11/14/05-01*, "Cortex-A8: High Speed, Low Power.")

Our tentative conclusion is that ARM's Cortex-A9 will consume less power than IBM's PowerPC 476FP when both processors are fabricated at similar geometries and are cruising at similar clock speeds.

Comparing the performance metrics of the PowerPC 476FP and Cortex-A9 is difficult. Although ARM is usually generous with core-area and clock-speed estimates, the Cortex-A9 is a conspicuous exception. ARM has never released those numbers. *Microprocessor Report* has estimated the Cortex-A9's performance by gathering information from other sources and from clues disclosed at technical conferences. We compared the Cortex-A9 with the Cortex-A8 in a recent article about the Texas Instruments OMAP4. (See Table 1 in *MPR 11/9/09-01*, "More Applications for OMAP4.")

Estimating Cortex-A9 Power

Our estimates for the Cortex-A9 soft core assume fabrication in TSMC's 45nm-GP bulk-CMOS process. That process is probably inferior to the 45nm-SOI process in which IBM fabricates the PowerPC 476FP hard core. (IBM is manufacturing its mighty POWER7 server processor in the same process.) All other things being equal, SOI is generally better.

Note that TSMC's 40nm process actually uses 45nm design rules with an optical shrink. Engineers familiar with these processes say that TSMC's 40nm process slightly reduces the core area but has little effect on power consumption and clock frequency. Keep this in mind when comparing metrics for processor cores fabricated in these processes.

MPR estimates that a speed-optimized Cortex-A9 core (without Neon) will consume 0.48mW per megahertz and easily exceed 1.0GHz in TSMC's 45n-GP. That's about half the power of the PowerPC 476FP in 45nm-SOI. Add Neon to the Cortex-A9, and we estimate that power will rise to 0.8mW per megahertz, still well below the PowerPC 476FP's 1.0mW per megahertz.

We cannot estimate static leakage. TSMC's 45nm-GP is almost certainly leakier than IBM's 45nm-SOI. Silicon-on-insulator transistors leak less current than bulk-CMOS transistors do, unless the design trades off that advantage by driving the transistors to higher speeds. If the Cortex-A9 is fabricated in TSMC's lower-leakage 45nm-LP process, the chip will bleed less power when the logic is idle, but we estimate that dynamic power will rise slightly to 0.53mW per megahertz (or 0.89mW per megahertz with Neon). We can't accurately estimate leakage for either processor without better information about their gate counts and transistors.

Our tentative conclusion is that ARM's Cortex-A9 will consume less power overall than IBM's PowerPC 476FP when both processors are fabricated at similar geometries and are cruising at similar clock speeds. The PowerPC processor probably has a little more clock-frequency headroom. ARM says that selective samples of a Cortex-A9 hard macro fabricated in TSMC's 40nm-G process have reached 2.0GHz. IBM says the PowerPC 476FP hard macro can exceed 2.0GHz in 45nm-SOI under similar conditions. Under worst-case conditions, ARM backs off to "over 1.0GHz" and IBM specifies 1.6GHz.

The synthesizable version of the PowerPC 476FP will likely be slower and less power efficient than the hard core, so the Cortex-A9 probably wins the power-consumption contest. Remember, however, that these conclusions assume our power estimates for the Cortex-A9 are close to accurate. Without better data from ARM, our estimates are merely educated guesses.

MIPS Does Multithreading

Like the PowerPC 476FP, the MIPS32 1004Kf Coherent Processing System aims for high-performance embedded systems that need SMP. (The 1004K is available with or without an FPU; we are comparing the FPU version, the 1004Kf, with the similarly equipped 476FP.) Whereas MIPS tends to focus on consumer electronics, IBM leans toward networking and communications.

With the 1004Kf, MIPS is taking a wholly different approach to throughput performance. Nearly two years after its introduction, the 1004Kf is still the only licensable processor core to combine coherent multiprocessing with hardware multithreading. (See *MPR 4/28/08-01*, "Multicore Multithreading With MIPS.")

Each MIPS 1004Kf core can manage two simultaneous threads. In effect, a quad-core design works like eight cores—virtually matching the eight-way SMP scalability of the PowerPC 476FP. Instructions from two or more software processes can share the same pipeline at the same time, switching contexts in one clock cycle. Each context may be a lightweight program thread or a heavyweight task, such as an operating system or application program. (Intel refers to this technology as Hyper-Threading.) The 1004Kf inherits multithreading from its parent, the MIPS32 34K processor. (See *MPR 2/27/06-01*, "MIPS Threads the Needle.")

On the other hand, the MIPS 1004Kf is handicapped by a simple uniscalar instruction pipeline and in-order execution. Per thread, it's no match for the PowerPC 476FP's four-way superscalar pipelines, out-of-order execution, and speculation. Still, hardware multithreading is potentially a higher degree of instruction-level parallelism than superscalar execution. Judging which processor delivers better real-world performance will require careful benchmarking.

Benchmarks Reveal Little

Unfortunately, the only across-the-board benchmarks we have at this time are the usual lame Dhrystone numbers. Dhrystone is an ancient single-threaded workload that hardly taxes the abilities of these two instruction-juggling processors. The PowerPC 476FP achieves 2.5Dmips per megahertz versus the MIPS 1004f's 1.56Dmips per megahertz.

In other words, Dhrystone says the PowerPC 476FP can run a single thread on a single processor faster than the MIPS 1004Kf can. We already guessed that by comparing their microarchitectures. What about multiple threads on multiple processors? On that question, Dhrystone says nothing.

Our advice to IBM: Get thee to EEMBC. EEMBC's new CoreMark benchmark is a better thread-performance measure than Dhrystone. (See *MPR 6/8/09-01*, "EEMBC's Dhrystone Killer.") And EEMBC's MultiBench suite is expressly designed for testing multiprocessor systems. (See *MPR 7/28/08-01*, "EEMBC's MultiBench Arrives.")

We do have CoreMark scores for the ARM Cortex-A9 MPCore and a MIPS processor closely related to the 1004Kf core. The ARM Cortex-A9 delivers 11.522 CoreMarks per megahertz when running four threads on four cores. That's 2.88 CoreMarks per thread per megahertz. ARM's report is available on EEMBC's CoreMark website (*www.coremark.org*).

During the technical review process for this article, *MPR* urged MIPS to release CoreMark scores, too. MIPS was already using CoreMark internally and posted its first score on February 10. That score is for the 34K processor, not the 1004K. The two cores are nearly identical, except the 1004K supports the MIPS Coherent Processing System. Their thread performance should be nearly identical, too. MIPS tested a 34K core running two hardware threads and measured 2.919 CoreMarks per megahertz—essentially matching the Cortex-A9 per core.

MIPS used CoreMark to test a MIPS32 34K processor running two hardware threads. It essentially matched the Cortex-A9's score.

MIPS says that running two instances of CoreMark in two hardware threads on the 34K processor is 25% faster than running CoreMark twice on the same processor configured for single threading. That's an interesting result. It would be more interesting if IBM published CoreMark scores for a dual-core implementation of the PowerPC 476FP. Comparing those scores wouldn't settle the argument of multithreading versus multiprocessing, but it would be a useful data point. Although IBM uses CoreMark internally, we have yet to see any published scores.

Comparing Performance and Power

For SMP, both the PowerPC 476FP and MIPS 1004Kf processors seem equally matched—if you accept the proposition that running eight threads on four cores can match the performance of eight threads on eight cores. (It depends on the application; more on this below.) Both processors have equal facilities for intercore coherency, global interrupt control, coherent I/O control, and coherent multicore debugging.

MIPS appears to have an advantage in I/O bandwidth to the L2 cache. Whereas the PowerPC 476FP has a 256-bit read interface and a 128-bit write interface, the MIPS 1004Kf has separate 256-bit read and write interfaces. On-chip bus bandwidth is a toss-up. Whereas IBM's CoreConnect PLB6 has separate 128-bit read and write channels running at half the core speed, MIPS uses an Open Core Protocol (OCP) bus with separate 64-bit read and write channels running at full core speed.

Comparing performance metrics for the MIPS and IBM cores is problematic. MIPS says a speed-optimized implementation of the 1004Kf soft core can exceed 1.3GHz in TSMC's 40nm-G bulk-CMOS process. That's a worst-case maximum-frequency estimate, within production margins. Under similar conditions, IBM says the PowerPC 476FP reaches 1.6GHz.

Don't forget that the PowerPC 476FP is a hard core fabricated in a different process (IBM's 45nm-SOI). If the future soft-core version of the PowerPC 476FP were to be fabricated in TSMC's 40nm-G, the speed losses from synthesis and bulk CMOS would probably lower its maximum worst-case clock frequency to approximately the same range as the MIPS 1004Kf processor.

MIPS says that speed-sorted silicon of the 1004Kf can reach 2.1GHz in TSMC's 40nm-G. IBM says that similar samples of the PowerPC 476FP can exceed 2.0GHz in 45nm-SOI. Roughly speaking, the IBM and MIPS cores deliver similar clock-frequency headroom.

The PowerPC 476FP lags in power consumption. IBM quotes 1.0mW per megahertz per core, whereas MIPS quotes 0.17mW per megahertz for a speed-optimized 1004Kf core in 40nm-G. Both estimates exclude the L2 cache. The main reason for the power disparity, we suspect, is the PowerPC 476FP's complex microarchitecture. There's a power penalty for a four-issue superscalar design with out-of-order execution, speculation, and eight-way SMP capability. The MIPS 1004Kf delivers 3.6 times as many Dmips per watt.

The MIPS core is much smaller than the IBM core. IBM says the PowerPC 476FP hard core is 3.6mm^2 in 45nm-SOI, excluding the L2 cache. MIPS says the speed-optimized 1004Kf soft core is 1.0mm^2 in 40nm-G, also excluding the

Feature	IBM PowerPC 476FP	MIPS MIPS32 1004Kf	ARM Cortex-A9 MPCore	ARM Cortex-A5
CPU Architecture	Power ISA v2.05	MIPS32 R2	ARMv7 Cortex-A	ARMv7 Cortex-A
Architecture Width	32 bits	32 bits	32 bits	32 bits
Symmetric Multiprocessing	2–16 cores (2–8 coherent)	2–4 cores (all coherent)	2–4 cores (all coherent)	2–4 cores (all coherent)
Coherent On-Chip Bus	2 x 128-bit CoreConnect PLB6	2 x 64-bit OCP	2 or 4 x 64-bit AMBA-3 AXI	2 or 4 x 64-bit AMBA-3 AXI
Intercore Coherency	Yes	Yes	Yes	Yes
Coherent Cache Snoop	L1 and L2	L1 and (optional) L2	L1 and (optional) L2	L1 and (optional) L2
Global Interrupt Control	Yes	Yes, optional use	Yes	Yes
Coherent I/O Control	Yes	Yes, optional use	Yes	Yes
Coherent Debug Unit	Yes	Yes, optional use (MIPS PDtrace)	Yes (ARM CoreSight)	Yes (ARM CoreSight)
Threads Per Core	1	1 or 2	1	1
Integer Pipeline	9 stages	9 stages	8 stages (9–11 clocks)	8 stages
Superscalar Execution	4-issue ALUs, 2-issue FPU	—	2-issue decode, 4-issue dispatch	Limited (Branch + ALU)
Out-of-Order Execution	Yes, with speculation	—	Yes, with speculation	—
Branch Prediction	Dynamic	Dynamic	Dynamic	Dynamic
L1 Cache (I / D)	32KB I / D per core	8KB–64KB I / D per core	16–64KB I / D per core	4K–64KB I / D per core
L2 Cache	Configurable 256KB–1MB	Optional 128KB–1MB	Optional Up to 2MB	Optional Up to 2MB
MMU	Yes	Yes	Yes	Yes
FPU	SP / DP (Power ISA 2.05)	SP / DP (MIPS32 1004Kf)	Optional, SP / DP (New Cortex-A9 FPU)	Optional, SP / DP (ARM VFPv3)
16-Bit Instruction Subset	—	Yes (MIPS16e)	Yes (Thumb-2)	Yes (Thumb-2)
DSP / SIMD Extensions	Minimal (MAC instructions)	Yes (MIPS DSP ASE)	Optional (ARM SIMD, Neon)	Optional (ARM SIMD, Neon)
Java Extensions	—	—	Jazelle RCT, optional Jazelle DBX	Jazelle DBX + RCT
Custom Extensions	—	Yes (MIPS CorExtend)	—	—
Execution Modes	User, supervisor	User, supervisor, kernel	User, supervisor, system, TrustZone	User, supervisor, system, TrustZone
Core Frequency (Worst-Case F_{MAX})	1.6GHz IBM 45nm-SOI (hard core)	>1.3GHz 40nm-G (speed optimized)	>1.0GHz 45nm-GP (speed optimized)	480MHz–1.0GHz 40nm-LP, 40nm-G (speed optimized)
Dhrystone 2.1 (Per Core)	2.5Dmips / MHz	1.56Dmips / MHz	2.5Dmips / MHz	1.57Dmips / MHz
Core Area	3.6mm^2 IBM 45nm-SOI (hard core)	1.0mm^2 40nm-G (32KB caches + FPU)	~3.1mm^2 40nm-G (MPR estimate)	0.53mm^2 40nm-LP (16KB caches, no FPU)
Power (Typical) Per Core	1.0mW / MHz IBM 45nm-SOI (hard core)	0.17mW / MHz 40nm-G (speed optimized)	0.48mW / MHz* 45nm-GP (speed optimized)	0.12mW / MHz 40nm-LP (speed optimized)
Availability	December 2010	June 2008	April 2008	1Q 2010

Table 1. Feature comparison of the IBM PowerPC 476FP, ARM Cortex-A9 MPCore, MIPS 1004Kf, and ARM Cortex-A5 MPCore. All these 32-bit licensable embedded-processor cores are designed for coherent symmetric multiprocessing. The IBM processor supports larger-scale SMP, with up to eight coherent cores per cluster. The other processors are limited to quad-core coherent clusters, but the MIPS 1004Kf uses hardware multithreading to double the number of threads per core. Initially, IBM is delivering the PowerPC 476FP as a hard macro, manufactured in IBM's own 45nm silicon-on-insulator (SOI) process. A synthesizable version portable to other fabrication processes will follow. Although a prehardened core and SOI are advantages in IBM's favor, the other cores still manage to be smaller and use less power. **MPR* estimate, without Neon; with Neon, add 0.32mW per megahertz. (n/a: data not available.)

L2 cache. Both processor configurations have 32KB L1 caches and FPUs. Even after accounting for the difference in process geometries—and remember, TSMC's 40nm process uses 45nm design rules with an optical shrink—the MIPS 1004Kf has a clear size advantage.

Multithreading vs. Multiprocessing

Overall, the PowerPC 476FP and MIPS 1004Kf are approximately equal competitors, albeit with many variables clouding the comparison. The most important variable is hardware multithreading. If running two threads per core can really match the performance of running two threads on two cores, the 1004Kf has the edge. It could deliver nearly the same throughput with half as many cores. But if the 1004Kf stumbles over thread management or memory I/O, two cores are better than one.

Interestingly, this question echoes the debate between IBM and Sun Microsystems in the world of high-performance server processors. IBM's POWER chips use stratospheric clock speeds (up to 5.0GHz for POWER6+) and wide-issue superscalar execution to hammer big jobs into submission with brute force. Sun's UltraSPARC T2 processors use lots of hardware threads (eight per core) and lots of cores (eight per chip) to divide big jobs into smaller tasks. Which approach is better depends greatly on the nature of the application. We expect the same will be true in the world of embedded SoCs when the PowerPC 476FP is compared with the MIPS 1004Kf. (See *MPR 2/1/10-01*, "Server Processors: Chapter 2009 [Part 2].")

Table 1 compares features of the IBM PowerPC 476FP, ARM Cortex-A9 MPCore, MIPS 1004Kf, and ARM Cortex-A5 MPCore. All are licensable 32-bit embedded-processor cores with integral support for coherent SMP. All are intended for high-performance embedded systems, and all but the PowerPC 476FP are available now as synthesizable cores. ARM designed the Cortex-A5 for midrange applications, so it's the slowest processor in this fast group, but we include it because it supersedes the six-year-old ARM11 MPCore. (See *MPR 10/26/09-01*, "ARM's Midsize Multiprocessor.")

When comparing the speed, power, and core-area metrics in the table, keep in mind that the fabrication processes are roughly comparable but not identical. Whereas IBM is using its own 45nm-SOI process for the PowerPC 476FP hard core, metrics for the other cores assume fabrication in TSMC's 40nm bulk CMOS. Our estimates for the ARM Cortex-A9 mix 40nm-G and 45nm-GP, but the design rules for both of those TSMC processes are essentially the same.

Tough Competition for IBM

One of ARM's selling points is popularity. The ARM architecture is the most popular 32-bit microprocessor architecture in the world. That leadership position translates into an extensive ecosystem of partners, tools, software, and developer familiarity. ARM's position in SMP is less formidable, however. There's still room for competitors to make it a horse race.

Historically, ARM tends to favor lower power, whereas MIPS tends to favor higher performance. The Power Architecture is better known for higher performance, too. Recently, ARM has been reaching toward higher performance, because more embedded systems demand it. At the same time, other CPU architectures are striving for lower power, because mobility is becoming universal, so battery life is crucial. The result of this competition is a flowering of new processor cores that are pushing performance to new heights while limiting power consumption to manageable levels.

IBM's PowerPC 476FP brings a fresh horse to the race. Developers who have acquired the habit of rarely looking much further than ARM for licensable processors now have another option to consider. The PowerPC 476FP is a viable option, especially if the target application needs more than four-way coherent SMP. For smaller designs—especially when low power is vital—the 476FP will have trouble competing with ARM and MIPS. Still, more choices usually enable better choices.

Price & Availability

IBM is initially delivering the PowerPC 476FP processor as a hard macro, to be followed by a synthesizable version of the core. Both versions are scheduled for general availability by the end of this year. The first licensee is LSI Corp. (formerly LSI Logic), which designed the configurable L2 cache controller for the processor. Chips using the hard macro must be manufactured by IBM, using IBM's 45nm silicon-on-insulator (SOI) process. The soft core will be portable to other fabrication processes and foundries. IBM hasn't publicly disclosed licensing fees for either version of the PowerPC 476FP.

More information about IBM's PowerPC 476FP:
www-01.ibm.com/chips/techlib/techlib.nsf/products/PowerPC_476FP_Embedded_Core

More information about IBM's CoreConnect PLB6 bus:
www-01.ibm.com/chips/techlib/techlib.nsf/products/CoreConnect_PLB6_Bus_Cores