Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

CORTEX-R4X: EXTREME MAKEOVER

Intrinsity's Fast14 Technology Accelerates ARM's Processor Core

By Tom R. Halfhill {9/24/07-01}

Imagine the impact on professional sports if anabolic steroids were safe, legal, and effective. Quarterbacks might be throwing 80-yard bombs. Sluggers might be hitting 600-foot homers. Tiger Woods wouldn't need his woods. Cricket might be watchable.

Intrinsity's rejuvenated Fast14 technology is having a similar effect on microprocessors. In July, Microprocessor Report described a new Power Architecture processor core that Intrinsity designed for AMCC using Fast14 dynamic logic. In that collaboration, Intrinsity played the role of a design house as well as an intellectual-property (IP) provider by designing a new Power-compatible microarchitecture to AMCC's specifications. (See *MPR 7/23/07-01*, "AMCC's Titan Core.") Now, Intrinsity is playing a different role for ARM. Starting with an existing microarchitecture—ARM's Cortex-R4 embedded-processor core—Intrinsity is using Fast14 to transform the synthesizable model into a hard macrocell.

The result is the ARM Cortex-R4X, the extreme-makeover edition of the Cortex-R4. The Cortex-R4X can reach 600MHz (worst case) in a low-leakage 65nm CMOS process (TSMC 65LP). In contrast, a speed-synthesized version of the soft core reaches only 380–400MHz in the same process. If the soft core were pushed to match the higher clock frequency of the Cortex-R4X by using a faster 65nm process, it would suffer much worse current leakage and be unsuitable for many battery-powered systems. At 600MHz, the hard core's throughput is about 960 Dhrystone mips, and dynamic power consumption is about 198mW (typical, including memories). ARM says the Cortex-R4X will be available in 1Q08.

The Cortex-R4X won't replace the Cortex-R4. It's simply another option for ARM's customers, recalling the days when ARM routinely offered hard versions of all its processor cores. Most customers now want synthesizable cores, so ARM accommodates them. But some developers need better performance than a soft core can deliver, or they prefer to drop a prehardened macro into a chip design without the fuss of synthesizing and laying out a soft core themselves. The Cortex-R4X will cost a little more to license, but it can save months of development time while preserving a customer's investments in software. It is 100% software compatible with the synthesizable version, and ARM will license and support it like any other ARM processor.

For Intrinsity, the Cortex-R4X opens the door to a potentially lucrative line of business. Intrinsity calls the Cortex-R4X an "RTL FastCore"—an existing processor core whose register-transfer-level (RTL) logic is made faster by applying Fast14 technology. In contrast, Intrinsity refers to the AMCC Titan processor as an "ISA FastCore"—a wholly new microarchitecture that Intrinsity designed for an existing instruction-set architecture (ISA). Intrinsity is ready and willing to make RTL FastCores or ISA FastCores for anyone. Meanwhile, Intrinsity is working on additional RTL FastCores for ARM.

Another way of expressing this concept is that Titan is a custom design, whereas the Cortex-R4X is a semicustom mod. Either way, the distinguishing factor is Intrinsity's patented Fast14 technology, which has significant performance advantages now recognized by AMCC and ARM, two important embedded-processor companies. In addition, Intrinsity is working on a project for Agere Systems, which

was acquired last year by LSI Logic. It appears that Intrinsity is successfully making the difficult transition from a fabless semiconductor company to an IP provider and design shop. (See this month's editorial, *MPR 9/24/07-02*, "Intrinsity Turns a Corner.")

Souping Up the Cortex-R4

MPR has previously described the microarchitecture of the Cortex-R4, so we won't dwell on those details here. (See *MPR 5/16/06-01*, "ARM Reveals Cortex-R4.") However, several aspects of the Cortex-R4X deserve attention. To begin with, it's a richly configured version of the Cortex-R4, not a stripped-down speedster. ARM says the initial version will probably have 32KB instruction and data caches, three 32KB tightly coupled memories (TCM), error-correction codes (ECC) on all caches and TCMs, a memory-protection unit (MPU) capable of managing 12 memory regions, 64-bit AXI master and slave ports with parity, and a full CoreSight debug unit supporting eight watchpoints and eight breakpoints.

Although the Cortex-R4X is a hard macro, Intrinsity says there will be some configurability for individual customers. The instruction and data caches are expandable to 64KB, and each TCM is expandable to 64K, offering a total of 192K of scratchpad memory. In place of the 12-region MPU, Intrinsity can substitute a smaller eight-region MPU. Usually, hard cores are available only in a take-it-or-leave-it configuration.

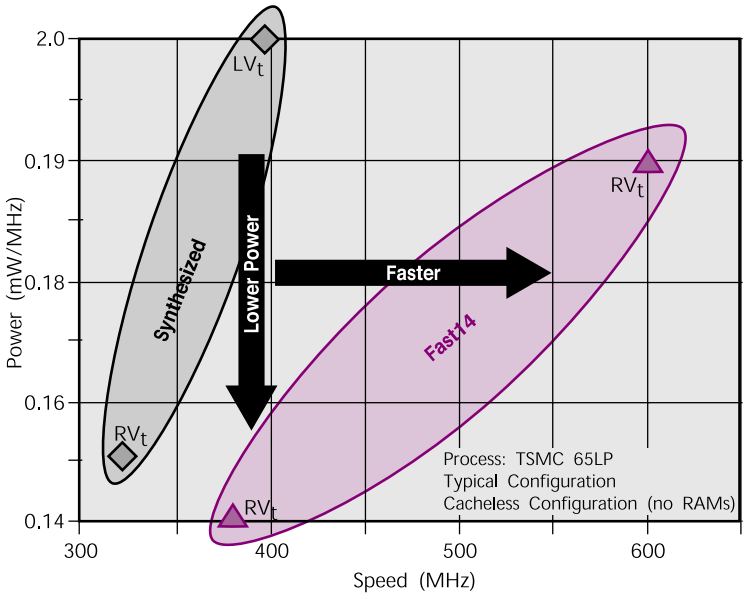

Figure 1. Intrinsity's Fast14 technology gives the Cortex-R4X hard core a dramatic speed advantage over the Cortex-R4 soft core in a low-leakage 65nm CMOS process. This chart compares the power/performance envelopes of both processors when they are fabricated in TSMC's 65LP. Note that the soft core must use speed-optimized low-threshold-voltage transistors (LVt) to reach 332MHz, whereas the hard core hits 600MHz with transistors operating at the regular threshold voltage (RVt). Without LVt transistors, the speed-optimized soft core leaks about 27mW; with LVt transistors, it leaks about 148mW.

One feature is missing from the Cortex-R4X: the FPU that ARM added later to the Cortex-R4F, primarily for the automotive and printer markets. (See *MPR 10/30/06-01*, "ARM Thumbs a Ride.") Although the Cortex-R4X is suitable for some automotive applications, floating-point arithmetic is useful mainly for engine controllers, where ARM believes the high performance of the Cortex-R4X is overkill.

Nevertheless, ARM is willing to license different configurations of the Cortex-R4X to customers that find the initial configuration (or TSMC's 65LP process) unacceptable for some reason. In those cases, Intrinsity must create a different hard macro to the customer's specifications. Intrinsity says its Fast14 design tools are highly automated and can rapidly generate minor variations of the macro, once the initial hardening is done. The initial configuration of the Cortex-R4X will occupy about 1.103mm^2 of silicon, including memories (0.855mm^2, core only). That's slightly larger than a similarly configured area-optimized Cortex-R4, but it is about the same size as a Cortex-R4 optimized with a high-speed synthesis library. ARM estimates that a smaller, slower configuration of the Cortex-R4X could shrink to about 0.6mm^2.

As Figure 1 shows, the Cortex-R4X is much faster than a Cortex-R4 synthesized in conventional static logic for TSMC's 65LP process. The Cortex-R4X can reach 600MHz at 10% below the core's nominal voltage of 1.2V, and the absolute maximum power is 340mW at 10% overvolt (1.32V). The temperature range for those specifications is –40°C to +125°C.

Typical power at 600MHz under nominal conditions—a figure more widely quoted in the industry—is 114mW (core only) to 198mW (with memories). A key point is that the Cortex-R4X can reach its maximum clock speed in this low-leakage process by using transistors with a regular threshold voltage (RVt), whereas the Cortex-R4 requires speed-optimized transistors with a lower threshold voltage (LVt) to reach its slower top speed. Those speed-optimized LVt transistors would leak much more power.

Remarkably, the Cortex-R4X will be the fastest-clocked implementation of an ARM processor core when it debuts early next year. Although Texas Instruments has demonstrated an ARM Cortex-A8 processor running at 1.0GHz, the initial implementation will run at 550MHz in TI's OMAP3 cellphone processor. By nearly halving the clock frequency, TI should be able to reduce power consumption by a like amount, and the superscalar Cortex-A8 would still be capable of good throughput at that speed. (See *MPR 7/24/06-01*, "The F1: TI's 65nm Cortex-A8.") Although a Fast14 implementation of the Cortex-A8 may seem logical, ARM says it's unlikely, because the F1 is already a semicustom implementation and is fast enough for the target application.

The Cortex-R4X is intended for different applications. Figure 2 shows ARM's sales estimates for various types of embedded systems in 2008. Although

the Cortex-R4X isn't the best choice for all these product categories, it's suitable for some higher-performance subsets—especially those requiring low overall power consumption and high throughput.

Fast14 Makes the Difference

All the performance improvements in the Cortex-R4X are thanks to Intrinsity's Fast14 technology, which MPR has described in past articles. (See *MPR 8/13/01-02*, "Intrinsity's Dynamic Designs.") Instead of traditional static logic, Fast14 uses 1-of-N domino logic (NDL) with four-phase overlapped clocking and a preponderance of NMOS transistors. Intrinsity's proprietary design tools automatically optimize the size of each transistor for its source/base load parasitics, required noise immunity, and required switching speed. The overlapped clock phases simplify timing closure. Fast14 circuits need no latches or P-channel transistors in the dynamic-logic paths, and they have only four levels of logic in places where a traditional static circuit would have 10 to 20 levels.

Figure 3 shows an example of the difference that Fast14's NDL makes. The figure illustrates two versions of a halt-propagate-generate cell used in adders and other logic circuits. The NDL version requires fewer than one-third as many transistors and is more than twice as fast.

Intrinsity uses a combination of industry-standard and proprietary hardware-design tools to transform a conventional processor core into a FastCore. In the case of the Cortex-R4X, Intrinsity started with ARM's Verilog model of the Cortex-R4. Intrinsity's design tools reconfigure the logic as NDL and customize the output of standard Cadence routing tools. For instance, Intrinsity uses techniques called "Twizzling" (a trademarked term) and "fat wires" to reduce noise in signal and clock paths. Wherever possible, Intrinsity's tools automatically insert redundant vias between metal layers for greater reliability.

Making a FastCore isn't an automatic push-button procedure—Intrinsity does plenty of low-level design work, too. However, Intrinsity says it has been steadily improving its tools toward the goal of greater automation. Although the initial version of the Cortex-R4X will be prehardened to a specific processor configuration, foundry, and IC process, Intrinsity is confident that its design tools can generate other variations more quickly. Creating a custom configuration of the hard core shouldn't take more time than a customer would spend hardening the soft core.

In any event, customers needn't wait for delivery of the hard core before developing their SoCs or writing software. The soft Cortex-R4 is available now, and the two cores are 100% cycle-compatible with each other. Developers can use processor- and system-level simulators to start their projects before the hard macro ships next year.

Competing Cores Still Impress

Despite the wonders of Fast14 technology, the Cortex-R4X still finds itself in a tight horse race with competing 32-bit processors—even though the others are fully synthesizable cores, not hard macros. ARC International, MIPS Technologies, and Tensilica have licensable soft cores that give the Cortex-R4X a run for its money.

Consider our following analysis merely a first-order approximation. It's virtually impossible to make true apples-to-apples comparisons among these processors, because there are so many variables. Those variables include the design, synthesis, and layout tools; optimizations applied during synthesis; physical libraries for logic and

Price & Availability

ARM says the Cortex-R4X hard macrocell will be available in late 1Q08 for TSMC's low-leakage 65LP CMOS process. Until then, customers can use the synthesizable Cortex-R4 for early development. ARM and Intrinsity may produce alternative configurations of the Cortex-R4X on request. ARM doesn't publicly disclose licensing fees for the Cortex-R4X.

For more information, visit

- *www.arm.com/products/CPUs/cortex-r4x.html*
- *www.intrinsity.com*

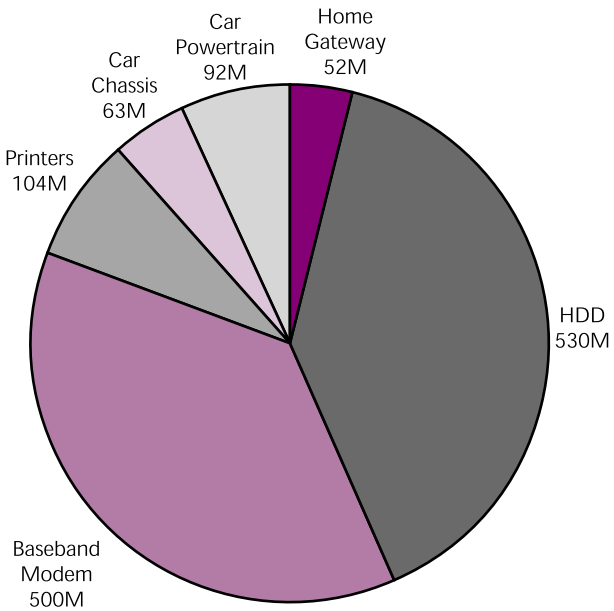

Figure 2. ARM is marketing the Cortex-R4X for embedded systems requiring a combination of high throughput and low power consumption. Some potential applications are battery powered, but others are tethered, so the Cortex-R4X is a crossover embedded processor—not quite as power-stingy as the smallest 32-bit processor cores, and not quite as fast as larger 32- and 64-bit processor cores. Even in tethered systems where power is plentiful, low heat dissipation is an advantage. (Data source: ARM.)

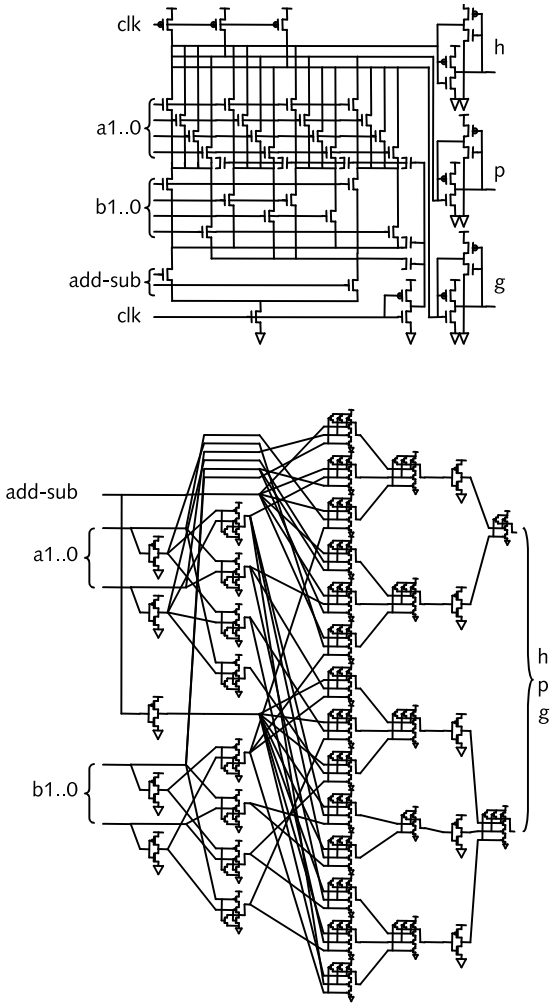

Figure 3. Using Intrinsity's Fast14 1-of-N domino logic (NDL), a halt-propagate-generate cell requires only 50 transistors (top), compared with 168 transistors for the same structure implemented in conventional static logic (bottom). The NDL version is about 2.5 times faster.

memories; the varieties of fabrication processes, even at the same geometry; the worst-case corner specifications; and other factors. Tensilica's Steve Leibson, a former *MPR* analyst, has written a cautionary white paper on this subject (*www.tensilica.com/products/WP_power_specs.htm*).

Featurewise, the ARC 750D is a close match for the Cortex-R4X. The ARC 750D is a preconfigured version of the user-configurable ARC 700. (See *MPR 3/14/05-02*, "ARC's Preconfigured Cores.") Like the Cortex-R4X, the ARC 750D has dynamic branch prediction, DSP instructions, a memory-management unit (MMU), configurable memories, non-maskable interrupts, and versatile I/O interfaces. ARC was unable to provide specifications for an ARC 750D fabricated in a leading-edge 65nm process, so we used data for a 750D fabricated in two variations of a common 90nm process.

When optimized for low power (TSMC 90LP), the ARC 750D consumes only about 0.12mW per megahertz—about one-third as much dynamic power as the Cortex-R4X consumes. Core area (excluding memories) is about 0.53mm^2, compared with 0.855mm^2 for the Cortex-R4X in a next-generation 65nm process. The catch is that the clock speed of a power-optimized ARC 750D peaks at a mere 200MHz, whereas the Cortex-R4X reaches 600MHz. Because both processors deliver approximately the same throughput in Dhrystone mips per megahertz, the ARM processor has a big advantage in raw performance. Implementing the ARC 750D in a speed-optimized 90nm process (TSMC 90GT) can boost the maximum clock rate to 700MHz while enlarging the core to 0.93mm^2, only a little bigger than the Cortex-R4X in TSMC 65LP. However, TSMC's 90GT is an expensive overdrive process that worsens both dynamic power consumption and current leakage. Notably, ARC does not publish a power-consumption estimate for a 700MHz ARC 750D in this process.

MIPS has a few processor cores competing with the Cortex-R4X, depending on whether the primary goal is low power or high throughput. For this analysis, the MIPS32 24KEc is a near match, especially in the power arena. Another core, MIPS32 74K, is a higher-throughput alternative. MIPS has performance estimates for both processors in TSMC's 65nm GP process, which is less power-stingy but faster than the 65LP process to which Intrinsity is porting the Cortex-R4X.

In 65GP, a speed-optimized MIPS 24KEc can reach 770MHz (worst case) while holding the core area to 0.7mm^2 and typical power consumption to 0.43mW per megahertz (at 1.0V). Those estimates make the MIPS 24KEc slightly smaller but slightly more power-hungry than a Cortex-R4X fabricated in 65LP. Fabricating the MIPS 24KEc in the same low-leakage 65LP process as the Cortex-R4X would almost certainly reduce the 24KEc's clock speed and worsen dynamic power consumption.

MPR is including the MIPS32 74K processor core in these comparisons because it's the latest MIPS core and is unabashedly designed for high throughput. The 74K has an uncommonly deep 17-stage pipeline with dual-issue superscalar execution and dynamic instruction reordering. (See *MPR 5/29/07-01*, "MIPS 74K Goes Superscalar.") In those respects, the 74K is more like ARM's powerful Cortex-A8 than the simpler Cortex-R4X. Nevertheless, the MIPS 74K is not wholly out of step in this comparison. In the 65GP process, using TSMC standard cells and Dolphin Technology memories, a speed-optimized 74K can break 1.0GHz. The post-layout core measures 1.7mm^2 and consumes 0.76mW per megahertz (including caches). Yes, the 74K is larger and hotter than the Cortex-R4X. But throttling back the clock frequency and targeting the 65LP process should bring the 74K more into line with the ARM core, so it's a viable alternative where high throughput is demanded.

ARM vs. Tensilica and FLIX

Tensilica's closest contender in this race is the Diamond 570T processor core, a member of the preconfigured Diamond

Feature	ARM Cortex-R4X	ARC ARC 750D	MIPS MIPS32 24KEc	MIPS MIPS32 74K	Tensilica Diamond 570T
Architecture	ARMv7	ARCompact	MIPS32	MIPS32	Xtensa
Max Core Freq	600MHz	200MHz	770MHz	1.04GHz	50–400MHz
(IC Process)	(65nm LP)	(90nm LP)	(65nm GP)	(65nm GP)	(65nm LP)
Pipeline	8 stages Uniscalar	7 stages Uniscalar	8 stages Uniscalar	17 stages Superscalar	5 stages Uniscalar + FLIX
Branch Predict	Dynamic	Dynamic	Dynamic	Dynamic	—
Instr Length	32 bits	32 bits	32 bits	32 bits	24 / 64 bits
Short Instructions	16 bits Thumb-2	16 bits	16 bits	16 bits	16 bits
DSP Instructions	Yes	Yes	MIPS DSP ASE-1	MIPS DSP ASE-2	Yes
Java Extensions	—	—	—	—	—
FPU	—	Optional 32 or 64 bits	Optional 64 bits	Optional 64 bits	—
MMU / MPU	MPU	MMU + TLB	Optional	Optional	MPU
Caches (I & D)	32K–64K each	0–64K each	0–64K each	8–64K each	16K each
Scratchpad RAM*	3 x 32K or 3 x 64K	0–512K	0–1MB	0–1MB	0–128K
NMI**	Yes	Yes	Yes	Yes	Yes
Bus Interface	AMBA-3 AXI 2 x 64 bits	BVCI, AHB, AXI 32–64 bits	OCP 2.1 64 bits	OCP 2.1 64 bits	AHB-Lite, PIF, XLMI, GPIO
Configurability	—	High	High	High	Low
Core Size (Base)	0.855mm^2	0.53mm^2	0.7mm^2	1.7mm^2	0.385–0.459mm ^{2}
(IC Process) Power / MHz	(65nm LP) 0.33mW†	(90nm LP) 0.12mW	(65nm GP) 0.43mW	(65nm GP) 0.76mW†	(65nm LP) 0.091mW–0.112mW
(IC Process) Dhrystone 2.1	(65nm LP) 1.6 Dmips/MHz	(90nm LP) 1.5 Dmips/MHz	(65nm GP) 1.44 Dmips/MHz	(65nm GP) 1.8 Dmips/MHz	(65nm LP) 1.52 Dmips/MHz
Introduction	1Q08	2004	2003	May-07	2006

Table 1. Feature comparison of the ARM Cortex-R4X, ARC 750D, MIPS 24KEc, MIPS 74K, and Tensilica Diamond 570T processors. All are 32-bit licensable embedded-processor cores. Only the Cortex-R4X is a prehardened macrocell; the others are fully synthesizable. Thanks to Intrinsity's Fast14 technology, the Cortex-R4X is a relative speed demon, even when fabricated in a low-leakage process. However, rival processors deliver similar throughput per clock cycle and may exceed ARM's performance in some applications. As a hard macro, the Cortex-R4X should save several months of development time over the soft cores. *ARM refers to scratchpads as tightly coupled memory (TCM). **Supports nonmaskable interrupts. †Power estimate includes caches. (Data sources: vendors.)

series introduced last year. (See *MPR 3/20/06-01*, "Tensilica's Preconfigured Cores.") Tensilica derived the 570T from the Xtensa LX core, and the 570T takes advantage of a unique Xtensa LX feature—Flexible-Length Instruction Xtensions (FLIX). Using FLIX, Tensilica defined some long instruction words that pack three operations into a 64-bit bundle. These FLIX instructions are in addition to the usual 16- and 24-bit instructions at the foundation of Tensilica's architecture. In effect, FLIX brings some superscalar powers to a simple uniscalar pipeline.

When synthesized for speed in the same TSMC 65LP process as ARM's Cortex-R4X, the Diamond 570T can reach 400MHz while consuming 0.112mW per megahertz. Excluding memories, the post-layout core is only 0.459mm^2. (When synthesized for area for TSMC's 65LP, the Diamond 570T shrinks to 0.385mm^2, and power consumption drops to 0.091mW per megahertz, but the clock rate plunges to 50MHz.) In other words, the speed-synthesized 570T is about one-third slower than the Cortex-R4X, approximately one-half as large, and consumes about one-third as much dynamic power.

If the Cortex-R4X were optimized for area instead of clock frequency, in the same TSMC 65LP process, it would fare better against Tensilica's Diamond 570T. ARM's processor could easily match the 400MHz clock speed of Tensilica's core while reducing both core area and dynamic power. (The synthesizable Cortex-R4 is about 0.7mm^2 and consumes 0.26mW dynamic power per megahertz under those conditions.) Both the Tensilica Diamond 570T and ARM Cortex-R4X have about the same Dhrystone throughput per clock cycle (1.5–1.6Dmips/MHz).

Tensilica's FLIX instructions could give the Diamond 570T an advantage in applications that have exploitable parallelism. Under ideal conditions, the 570T can execute three FLIX operations per clock cycle. But the Cortex-R4X is capable of similar tricks. Like the Cortex-R4, it can issue, execute, and retire two instructions per clock cycle, even though it doesn't have end-to-end superscalar pipelines. Only comparative benchmarking could settle the question of which processor actually delivers better throughput. Tensilica has published certified EEMBC benchmarks showing that the Diamond 570T easily beats an ARM1026EJ-S or ARM1020E; ARM hasn't published EEMBC scores for the Cortex-R4 or Cortex-R4X. If high throughput were paramount, a Tensilica developer could fit two 570T processors into the same space as one Cortex-R4X while staying within the ARM processor's power envelope.

ARM's Cortex-M1 for Low-Power Actel FPGAs

While Intrinsity has been developing a faster version of ARM's Cortex-R4, Actel has been working on a lower-power version of ARM's Cortex-M1. On September 17, Actel announced that the Cortex-M1 will soon be available on Igloo FPGAs—Actel's lowest-power family of programmable-logic devices.

As Microprocessor Report noted earlier this year, the Cortex-M1 is the first ARM processor core specifically designed and licensed for deployment on FPGAs. (See MPR 3/19/07-01, "ARM Blesses FPGAs.") Until now, the Cortex-M1 was available only on Actel's larger Fusion and ProASIC3 devices. By porting the processor to the lower-power Igloo family, ARM and Actel are pitching FPGAs as practical solutions for battery-powered mobile systems. Actel estimates that the market opportunity for programmable-logic devices in portable products will be more than $500 million by 2010.

Actel fabricates Igloo devices in an older UMC 0.13-micron CMOS process—two generations behind the latest FPGAs from market leaders Altera and Xilinx. Nevertheless, Igloo chips consume less power than conventional FPGAs, because they implement the programmable-logic fabric in flash memory instead of SRAM. Flash memory typically has only one transistor per bit cell instead of six transistors. Core voltage is relatively low (1.2V–1.5V), as is current leakage. Unlike SRAM, flash memory is nonvolatile, so Actel FPGAs don't have to burn power configuring their fabrics after a cold start.

In addition, a flash-based FPGA can stop its clock and enter a very low power state without erasing the fabric.

Portable systems often spend most of their time in a sleep mode, and Actel's low-power states are very power-stingy. Actel says the Cortex-M1 will draw less than 24 microamps in static mode, less than 20 microamps in "Flash-Freeze" mode, and less than 3 microamps in deep-sleep mode (stop clock).

Igloo-series FPGAs have 30,000 to three million programmable system gates, depending on the device. To put those capacities in perspective, the Cortex-M1 requires about 200,000 system gates. Although the processor is too large for the lowest-priced Igloo devices (which start at $1.50), it's practical for a 600,000-gate Igloo M1AGL600, which will start at $3.70. Actel is sampling the first ARM-capable Igloo parts now and will begin shipping the M1AGL600 in 4Q07. More ARM-capable Igloo devices are coming next year.

As MPR reported last March, a big advantage of ARM's partnership with Actel is that customers don't need to license the Cortex-M1 from ARM, pay upfront licensing fees, or pay chip royalties. All those costs are built into the prices of the FPGAs. Customers get a black-box implementation of the Cortex-M1, ready for deployment. This arrangement can easily slash 18 to 24 months off a project schedule, compared with developing an ASIC.

For more information about Actel's Igloo family and the Cortex-M1, see:

- *www.actel.com/products/igloo/*
- *www.arm.com/products/CPUs/ARM_Cortex-M1.html*

Table 1 summarizes the features and specifications for all these competing processor cores. If it's not obvious by now, our analysis is splitting hairs. All these cores have very similar features. All are so tiny in a 65nm, or even a 90nm, process that they would occupy only a small corner of an SoC; on-chip memory and peripherals would use most of the silicon. Likewise, the power estimates depend on so many variables that all these processors are suitable for "low-power" embedded applications needing high throughput—especially if developers are willing to trade some clock speed for lower dynamic power. Intrinsity is eager to apply Fast14 technology to any processor from any company (for a price, of course), so it's tantalizing to speculate about possible FastCore implementations of processors from ARC, MIPS, and Tensilica.

By working with Intrinsity to create a Fast14 implementation of the Cortex-R4, ARM is offering customers new flexibility. The same up-to-date core is available in hard and soft versions, with significantly different performance characteristics. Even if the Cortex-R4X doesn't win every comparison made with competing cores, it's a drop-in macrocell that can save many months of development time, and that's worth a lot. Its high throughput could make the difference between creating an easily programmed single-core chip or a more problematic multicore design. In addition, it provides an upgrade path from chips based on the soft version of the processor. All these considerations weigh heavily in time-to-market and product life-cycle calculations. MPR considers the Cortex-R4X a worthwhile addition to ARM's lineup and an endorsement of Intrinsity's Fast14 technology. ◇