Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

EMBEDDED SYSTEMS CONFERENCE HIGHLIGHTS

News From the ESC Exhibition Floor and Meeting Rooms in San Jose

By Tom R. Halfhill {04/23/07-01}

This year's Embedded Systems Conference in Silicon Valley began with an unusual keynote address by former Vice President Al Gore—or, as he jokingly refers to himself, "the former next president." The San Jose Civic Center Auditorium was filled with a mostly receptive crowd of engineers, although there were skeptics, too.

Perhaps our expectations were low, because Gore isn't a techie and isn't known as a charismatic speaker. But we were surprised by his hour-long talk. He was amusing, self-effacing, and informative, and he didn't inflict a single PowerPoint slide on the audience. The thrust of his address was that global climate change is a wake-up call, much like the Sputnik launch in 1957. Just as the ensuing space race became an opportunity for great technological progress, Gore sees an opportunity for today's engineers to design a new generation of power-efficient systems for all applications. This effort will not only help the environment and boost the economy but could also spark the same excitement about science and engineering that the Apollo project did in the 1960s. As a result, those technical fields may attract more young people, which would improve the nation's competitiveness. Of course, not everyone agrees with Gore's views. Nevertheless, he made a logical argument for wiser investments in science and engineering.

After Gore's address, the crowd flowed across the street to the McEnery Convention Center, where practical engineering was on display in dozens of exhibition booths and meeting rooms. The following report isn't an exhaustive account of everything at ESC—the dailies and weeklies are better at that—but rather a collection of interesting microprocessor-related news we gathered. Highlights:

MIPS Technologies has negotiated a landmark licensing deal with STMicroelectronics that appears to resolve a long-running dispute with China over MIPS-like derivatives of the MIPS architecture. This deal opens the door for the Chinese to design fully compatible MIPS processors while protecting the intellectual property (IP) of MIPS Technologies.

The Power.org consortium has formed technical subcommittees to resolve differences among Power Architecture microprocessors and processor cores. This initiative is of particular interest to embedded-system developers, because it aims to resolve differences among bus protocols, on-chip interconnects, debug interfaces, and other things that could deter some developers from using Power processors.

ARC International announced a surprising acquisition of Teja Technologies. We suspect there's more to this deal than ARC disclosed in its press release.

NXP Semiconductor (formerly Philips Semiconductors) showed some fascinating preliminary results of tests with the power-consumption benchmarks that EEMBC introduced last year.

Innovasic Semiconductor, which specializes in satisfying demand for chips discontinued by other companies, wants to clone the Intel 386 processor, which Intel recently dropped from its product catalog.

Chinese Make Peace With MIPS

As *MPR* reported in two major articles in 2005 and 2006, Chinese engineers have designed a family of microprocessors that's about 90% compatible with the classic MIPS

instruction-set architecture (ISA). (See *MPR 6/26/06-02*, "China's Microprocessor Dilemma," and *MPR 7/25/05-01*, "China's Emerging Microprocessors.") Known in the West as the Godson family, these processors closely mirror the MIPS 32- and 64-bit architectures, omitting a few instructions and features protected by MIPS patents in the U.S. and elsewhere. However, MIPS Technologies hasn't been flattered by these imitations. For years, MIPS has unsuccessfully tried to sell the Chinese a MIPS architectural license that would allow Godson processors to become fully MIPS compatible.

One problem, for the Chinese, is that an architectural license is the most expensive type of processor-IP license. It allows the license holder to design an original microarchitecture fully compatible with the CPU architecture. An example of a MIPS architectural licensee is Cavium Networks, which designed its custom cnMIPS core for a rapidly growing line of 64-bit network processors. (See *MPR 10/5/04-01*, "Cavium Branches Out," and *MPR 2/6/06-01*, "Cavium Expands Octeon Family.") Most processor-IP licensees purchase a less expensive core license, which merely allows the holder to use an existing processor core in a new chip design. An architectural license from a company like MIPS typically costs several million dollars.

Another problem, for MIPS, is that China has a poor reputation for protecting IP. CD-ROMs and DVDs filled with the latest Windows software and Hollywood movies are widely available on the street for $5. That's a serious concern for a company whose business model is based almost wholly on licensing IP. As one MIPS executive puts it: "All our IP fits on a single thumb drive." If that IP—essentially, the RTL files of the processor cores—walks out the door and appears on a renegade website, it could seriously damage MIPS's business. MIPS hopes that as the Chinese develop more of their own IP, the risk of IP theft will decline over time.

Godson processors are the brainchild of China's leading CPU architect, Dr. Weiwu Hu, a professor at the Institute for Computing Technology (ICT) at the Chinese Academy of Sciences in Beijing. *MPR* has visited Hu and talked with him in detail about his Godson processors and plans for future designs. (See our previously cited articles.) It's not surprising that Hu would have trouble justifying the cost of a MIPS architectural license when he can make processors that are nearly MIPS compatible. It's also not surprising that MIPS worries about protecting its IP and is eager to expand its business in China.

After lengthy and difficult negotiations (one can only imagine), MIPS has broken the impasse with an innovative licensing deal. Instead of licensing the MIPS architecture directly to the Chinese, MIPS has sold a MIPS64 architectural license to STMicroelectronics. ST, in turn, will use ICT to design MIPS64-compatible processor cores for future SoCs. These new 64-bit microarchitectures will be fully compatible with the MIPS64 ISA and will be based on ICT's existing Godson processors and their derivatives, known as Loongson cores. ST can promote the resulting cores as "MIPS-Based" and "MIPS-Verified" (both terms are trademarked by MIPS), just as Cavium promotes its cnMIPS cores. ST can integrate the MIPS-compatible cores into SoCs and sell the chips anywhere.

The key to this deal is positioning ST as the middleman between MIPS and ICT. Of course, ST came up with the licensing money. But ST also is responsible for protecting the licensed IP, just as any licensee is responsible for supervising its subcontractors. In this case, ICT is the design house, and ST is a large, reliable company that MIPS can reasonably expect will run a tight shop. Security measures include audits of ICT's networks to find vulnerabilities and IP repositories on special intranets decoupled from the Internet.

No system is airtight, of course. But MIPS says the licensing deal brings ICT into the MIPS community and "is better than the situation we had before." For the Chinese, the deal legitimizes ICT's processor designs and relieves ICT of the responsibility to support a CPU architecture that wasn't 100% MIPS compatible. *MPR* is watching this novel arrangement with great interest.

Reunifying the Power Architecture

As *MPR* reported last year, the Power.org consortium is working to merge different features that crept into the Power Architecture during the years when IBM and Motorola/Freescale weren't fully cooperating with each other. (See *MPR 8/21/06-01*, "The New Power Architecture.") Power.org had a large booth at ESC and announced that new technical subcommittees are working toward specific goals that could encourage more embedded-system developers to use Power processors and cores in their projects.

Also at ESC, as *MPR* has already reported, Freescale announced that it will license some of its Power processor cores to chip developers for the first time. Power.org member IPextreme is Freescale's licensing agent. (See *MPR 4/2/07-01*, "Freescale Licenses Power Cores.") Freescale's new licensing strategy and the Power.org subcommittees are further evidence that the consortium is rejuvenating the Power Architecture, making it more competitive with rival architectures—especially in the embedded market.

Two examples of technical differences that Power.org wants to resolve are debug interfaces and on-chip interconnect standards. As noted in our article about Freescale's licensing, the Freescale Power e200 processor cores have Nexus debug interfaces, whereas IBM's licensable Power processor cores have JTAG debug interfaces. Likewise, Freescale's e200 cores use an AMBA-AHB bus to connect on-chip peripherals, whereas IBM's licensable Power cores use IBM's own CoreConnect bus. Power.org's goal is to smooth out those differences for tool vendors and developers.

This initiative won't necessarily choose a single standard. For instance, it's hard to imagine IBM abandoning CoreConnect, which is an IBM-native technology, just as it's hard to imagine Freescale abandoning AMBA, which is widely adopted

in the licensable-IP community. More likely, Power.org will define common gasket interfaces and perhaps urge Freescale and IBM to design future processor cores with interchangeable interfaces. Power.org acknowledges that some changes could take years, but the long-term goal is to make Power more competitive with rivals like ARM and MIPS.

One technical subcommittee will define the Embedded Power Architecture Platform Requirements (E-PAPR). These requirements describe the interactions between Power processors and the embedded system's firmware, hardware devices, and operating system. The objective is to standardize and streamline the boot process, device tree, interrupt management, memory management, and power management, among other things. Power.org members on this subcommittee include AMCC, Cadence, Ericsson, Freescale, HCL, IBM, IPextreme, P.A. Semi, Thales, and Wistron.

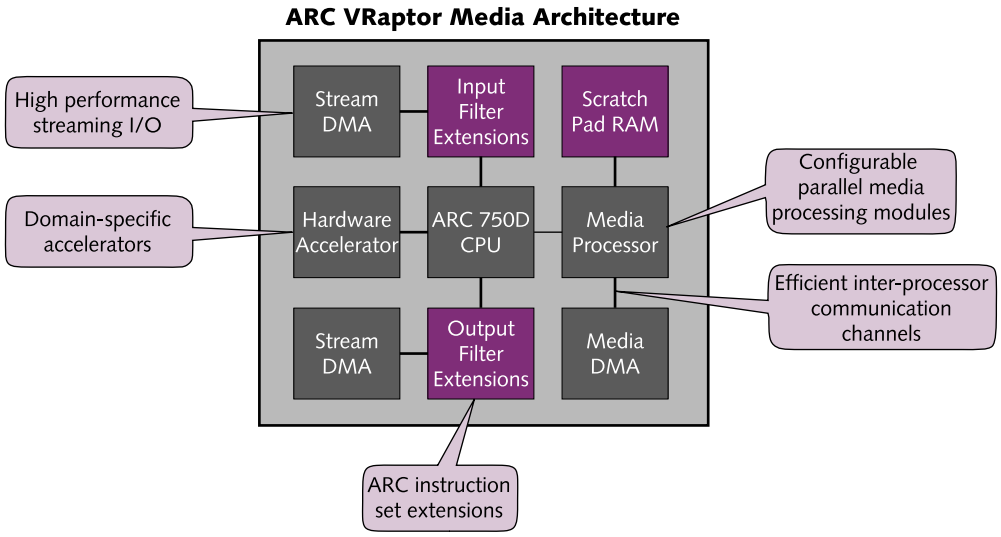

Figure 1. ARC VRaptor block diagram. ARC introduced this multicore multimedia engine at last fall's Microprocessor Forum. Its configurable architecture is based on one or more ARC 750D processor cores with media extensions, multiple 128-bit SIMD processors, media-specific accelerators, high-bandwidth on-chip interconnects, and high-performance streaming I/O. ARC is scheduled to reveal more technical details at the next Microprocessor Forum in San Jose (May 21–23).

Another subcommittee will define a "world-class SoC design ecosystem," including performance-analysis tools, design methodologies, and simulators. Subcommittee members include Cadence, Chartered, Denali, Ericsson, Freescale, HCL, IBM, IPextreme, Mentor Graphics, and Synopsys. This group will probably coordinate with the bus architecture subcommittee, which wants to define a three-level bus hierarchy for SoC interconnects. Members of that team include AMCC, Cadence, Denali, Ericsson, Freescale, HCL, IBM, Mercury Computer, P.A. Semi, and Synopsys. Additional subcommittees are working on the debug interfaces and common software frameworks for various types of embedded systems. For a complete list of Power.org subcommittees and their missions, follow the web link in the "For More Information" box in this article.

ARC Acquires Teja Technologies

Since its initial public offering in 2000, just as the tech bubble was bursting, ARC International has struggled financially and is still looking for its first profitable quarter. But under CEO Carl Schlachte, the company has steadily gained strength and is nearly breaking even. Now ARC feels affluent enough to make its first corporate acquisition since the 1990s. ARC's surprising purchase of Teja Technologies not only signals new financial confidence but is also an important strategic move.

True, the acquisition is worth only $5 million, but it's still important. ARC covets Teja's software-development team, which will work on new and existing tools for ARC's VRaptor Multicore Architecture. VRaptor is a powerful multimedia engine based on ARC's configurable 32-bit embedded-processor cores. (See Figure 1.) Thanks to previous acquisitions of MetaWare and Precise Software in the 1990s, ARC already has some software-development tools and the MQX real-time operating system (RTOS). The Teja acquisition brings additional programmers and engineers with experience in multicore design and performance analysis.

But more may be going on. As *MPR* reported last year, Teja also has a product called Teja FP (FPGA Platform). (See *MPR 4/3/06-02*, "Teja's FPGA Play.") This package of development tools, software, and hardware IP lets software programmers build and deploy a packet processor in an FPGA using ANSI C instead of a hardware-description language (HDL)—and without spinning custom silicon. Using Teja's performance profilers and analysis tools, programmers can partition their new or existing data-plane code written in C. The most compute-intensive code runs in optimized logic in the FPGA fabric, while less-critical code runs on soft processor cores synthesized in the fabric. Teja FP uses low-cost Xilinx MicroBlaze processor cores, but ARC's configurable cores are also suitable, if ARC pursues this path.

Teja FP could pave the way for ARC customers to deploy some of their SoC designs in FPGAs instead of merely using FPGAs for development and verification. ARC's biggest competitor, ARM, is moving in that direction. Last month, ARM introduced the Cortex-M1, its first processor core intended for FPGA deployment. (See *MPR 3/19/07-01*, "ARM Blesses FPGAs.")

For More Information

Chinese Make Peace With MIPS
- MIPS Technologies: *www.mips.com/PressRoom/PressReleases/2007-03-28*
- STMicroelectronics: *www.st.com/stonline/stappl/press/news/year2007/t2154.htm*

Reunifying the Power Architecture
- Power.org Consortium: *www.power.org*

ARC Acquires Teja Technologies
- ARC International: *www.arc.com/news/PressRelease.html?id=316*
- Teja Technologies: *www.teja.com/newsevents/pr_arc.html*

NXP Trumpets EEMBC's EnergyBench
- EEMBC: *www.eembc.org*
- NXP Semiconductor: *www.standardics.nxp.com/products/lpc3000/lpc3180/*

Innovasic Eyes the Intel 386
- Innovasic Semiconductor: *www.innovasic.com*

NXP Trumpets EEMBC's EnergyBench

As *MPR* reported last year, the Embedded Microprocessor Benchmark Consortium (EEMBC) finished a difficult project to define the industry's first standardized power-consumption benchmark. (See *MPR 7/17/06-02*, "EEMBC Energizes Benchmarking.") Unfortunately, few of EEMBC's member companies have published EnergyBench scores. That doesn't necessarily mean the companies are shunning EnergyBench. EEMBC allows members to share test results with customers under a nondisclosure agreement while keeping the scores private. Of course, openly published scores—which EEMBC must certify—are much more interesting (and believable). At its ESC booth, with EEMBC's permission, NXP Semiconductor was showing off some preliminary scores for one its ARM9-based microcontrollers.

NXP tested its LPC3180 MCU, which has an ARM926EJ-S processor core, ARM's VFP9 vector floating-point coprocessor, 32KB instruction and data caches, and an abundance of on-chip peripherals. (See *MPR 4/4/05-02*, "ARM-Based MCUs Flex Muscles.") NXP manufactures these chips at multiple fabs, using a 90nm CMOS process codeveloped with the Crolles alliance. The LPC3180 has two operating ranges: a nominal 1.2V mode (±0.1V) supporting clock frequencies from 20MHz to 208MHz, and a special low-power 0.9V mode supporting clock frequencies from 13MHz to 20MHz. The 0.9V mode is particularly impressive, because NXP had to use extensive clock gating and other power-reduction techniques throughout the chip.

As our previous article explained, EnergyBench periodically samples power consumption while running the individual tests in an EEMBC benchmark suite, then uses that data to calculate average (typical) power, maximum power, and the amount of energy consumed by that workload. NXP chose to use EnergyBench while running EEMBC's automotive suite—a good match for one of the LPC3180's target markets and the only benchmark tests that fit entirely within the LPC3180's on-chip memory. EEMBC created a custom application using National Instruments' LabVIEW software to support the graphical user interface of EnergyBench. The custom application displays the benchmark results and benchmark-specific configuration data.

NXP conducted multiple test runs to measure LPC3180 power consumption at various voltages and clock speeds with the FPU and caches turned on or off. Rarely would a real-world system turn off the FPU and caches, but NXP wanted detailed engineering data about the trade-offs of those features. Some ARM-based FPUs (e.g., those based on the ARM7TDMI processor core) lack an FPU and caches.

Even the throughput data from these test runs is interesting. The caches improved throughput by 23% at clock speeds from 13MHz to 104MHz and boosted throughput by 64% at 208MHz. NXP's logical explanation for that difference is that the memory subsystem generally keeps up with demand until the CPU frequency exceeds the bus frequency (104MHz). Then memory falls further behind, and the CPU relies more heavily on the caches. Not unexpectedly, these tests also showed that a hardware FPU (the ARM VFP9, in this case) is far better at executing floating-point workloads than software libraries are. In some test runs, the LPC3180's FPU improved throughput by 435%.

More illuminating are the EnergyBench results. Figure 2 shows the LPC3180's power consumption

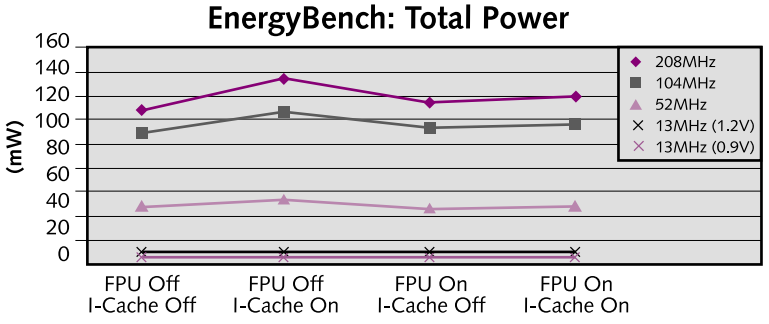

Figure 2. NXP Semiconductor measured power consumption of its LPC3180 microcontroller using EEMBC's automotive benchmark suite and EnergyBench. This chart shows preliminary data only for the Basic Floating-Point test, one of 16 kernels in the automotive suite. In all, NXP conducted 20 test runs, varying the clock frequency (13–208MHz), CPU voltage (0.9–1.2V), and power states of the FPU and instruction cache (FPU off, cache off; FPU off, cache on; FPU on, cache off; and FPU on, cache on). The power states of the FPU and cache made surprisingly little difference, especially considering their great effect on throughput. (Data source: EEMBC)

while running a floating-point workload with the FPU and instruction cache turned on or off. Power correlates with capacitance, frequency, and voltage ($P=CFV^2$), but the power states of the FPU and cache produced surprisingly small variations—perhaps a testament to the efficiency of their design. Another interpretation is that emulating an FPU in software is inefficient for both throughput and energy consumption.

Figure 3 shows the most interesting EnergyBench results. It graphs the energy (in microjoules) consumed per loop iteration while running the same floating-point workload. Again, NXP measured several test runs at different clock speeds and voltages, with the FPU and instruction cache turned on or off. Enabling the FPU made the biggest difference. It dramatically reduced energy consumption compared with test runs using the emulation libraries. Enabling the instruction cache significantly reduced energy consumption, too, though not quite as dramatically as the FPU did.

NXP's preliminary EnergyBench tests suggest a few conclusions. First, as the company claims, the LPC3180 is an efficient device. Even at its highest core voltage and clock frequency, with the FPU and instruction cache fully powered, this ARM9-based MCU draws only 120mW in this test. Second, the ARM VFP9 adds a little silicon cost to the chip, but it's much faster than floating-point emulation and more energy-efficient, too. Third, the LPC3180 can complete a finite task more quickly at a faster clock speed while using no more energy than it would use at a slower clock speed, revealing the fallacy of relying exclusively on "typical" power numbers measured at lower frequencies. *MPR* encourages other companies to use EnergyBench and to publish their EEMBC-certified scores.

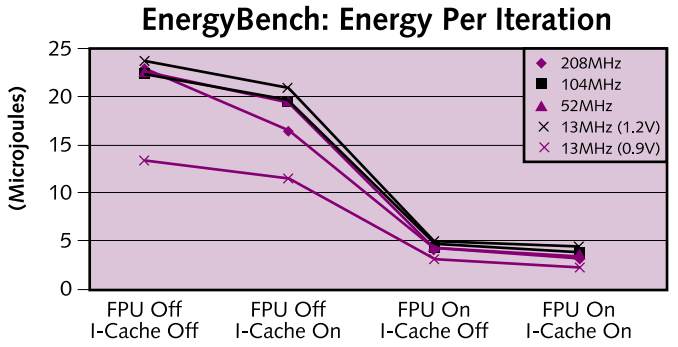

Figure 3. LPC3180 energy consumption per floating-point loop iteration, in microjoules. This graph of 20 test runs covers the same operating ranges and power states as the data in Figure 2. Switching on the FPU greatly reduces energy consumption versus emulating the FPU in software. Switching on the instruction cache also reduces energy consumption. However, note that the scale of this graph makes the fully powered test runs look too much alike. The difference between the two 13MHz tests (0.9V and 1.2V) is actually about 2 to 1. (Data source: EEMBC)

Innovasic Eyes the Intel 386

We may never be certain if there's an afterlife for mere mortals, but there's definitely one for microprocessors. Innovasic Semiconductor, a 15-year-old company based in Albuquerque, New Mexico, specializes in selling processors and other ICs that have reached their end-of-life stage at other companies. Innovasic extends the longevity of discontinued parts by cloning them—an approach that has also been proposed for humans but is less controversial when applied to chips. Thanks to Innovasic, perfectly good embedded systems that depend on old devices can live on, even after the original suppliers have dropped the parts from their catalogs in favor of shiny new products.

Until recently, Innovasic focused exclusively on small processors and MCUs, as well as on miscellaneous parts, such as digital/analog converters, physical-layer interfaces (PHY), real-time clocks, crystal oscillators, PLLs, and the like. Recently, Innovasic moved up the food chain by introducing the Fido 1100, a real-time 32-bit MCU based on a new architecture that's compatible with the Motorola/Freescale 68000 instruction set. (See *MPR 1/16/07-01*, "Fido Runs With the Big Dogs.") Now Innovasic is eyeing another classic CPU architecture: Intel's x86.

Specifically, Innovasic wants to clone the 386, the first 32-bit x86 processor, which Intel introduced in 1985. Why the 386? Because it's still widely used in embedded systems and was unceremoniously dropped from Intel's catalog last year. Intel has purged its product lines by discontinuing many old processors, including some or all members of the 186, 386, 486, 8051, i960, MCS90, and MSC251 families. Those may be crumbs to Intel, but Innovasic views them as a fresh business opportunity—particularly the 386, and perhaps even the 486.

Innovasic didn't exhibit at ESC, but *MPR* met with CEO Keith Prettyjohns there. Although Prettyjohns has approached Intel about reproducing the 386, he says Intel isn't too interested, for various reasons. Nevertheless, even without Intel's help (or permission), he believes Innovasic can make a legal clean-room clone of the 386. Prettyjohns says his engineers have become experienced at this kind of work, and he notes that most 386-related patents should be expiring right about now.

It's understandable that Intel could have sound business reasons for not participating in such a project. However, *MPR* hopes Intel won't actively discourage it. Handing off discontinued parts in an orderly fashion to companies like Innovasic seems like a good idea. Embedded-system developers would feel more comfortable about using Intel processors in their designs if there's more assurance the devices won't suffer an untimely death. Even for humans, the promise of an afterlife is a positive motivation.