

# M I C R O P R O C E S S O R

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

## ARM-BASED MCUs FLEX MUSCLES

*Actel, Oki, and Philips Launch Innovative 32-Bit Microcontrollers*

*By Tom R. Halfhill {4/4/05-02}*

ARM-based chips continue to gain strength in the fast-growing microcontroller market. At the recent Embedded Systems Conference, Actel announced FPGAs with specially optimized ARM7 processors encrypted in programmable logic; Oki Semiconductor unveiled

the world's smallest ARM-based MCUs; and Royal Philips Electronics introduced the first ARM9-based MCUs fabricated in a 90nm CMOS process.

All these milestones provide more reasons to upgrade from 8- or 16-bit MCUs to 32-bit ARM-based devices. It's not just about 32-bit processing power and memory addressing, although those advantages are compelling enough for many customers. ARM's MCU licensees are offering fresh bait. For example, Actel's new ARM7-enabled FPGAs will tempt some designers to bypass a custom ASIC or SoC project in favor of these versatile off-the-shelf parts. The Philips 90nm ARM9 chips will cut costs and reduce power consumption, overcoming two perennial objections to 32-bit MCUs. And Oki's minuscule ARM7 chips—about one-fourth the size of a dime, fully packaged—will go a long way toward defeating the board-space argument in favor of less powerful devices.

ARM makes no secret of its ambition to rule the MCU market, dominated for so long by low-end chips with aging 8- and 16-bit architectures. ARM seems to have an inherent disadvantage: all things being equal, an 8- or 16-bit MCU should always consume less power, fit into a smaller package, and cost less to manufacture than a 32-bit MCU. But all things are rarely equal. At deep-submicron scales, a 32-bit processor core becomes a mere microdot on the die, without sacrificing its superior capabilities. No one expects 32-bit MCUs to completely displace the 8- and 16-bit bottom-feeders, but ARM can capture enough market share to fuel many years of strong growth.

### **It's Not Hard...Not Soft...It's Firm**

Actel's twist on 32-bit MCUs will include a new series of FPGAs with integrated ARM7 processor cores, scheduled to ship later this year. This isn't the first time a vendor has integrated an ARM processor into an FPGA, but it is the first time ARM has collaborated with an FPGA vendor to optimize an ARM core for "soft" integration. We qualify the term "soft," because this particular ARM implementation is part-way between a synthesizable core and a prehardened core—let's call it a firm core.

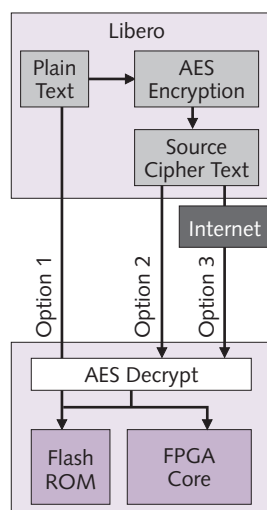
Until now, ARM hasn't licensed its soft cores for FPGA integration (other than design prototyping) because it feared the theft of its intellectual property (IP). A moderately clever engineer could copy the processor's gate-level netlist from the FPGA and derive a synthesizable model that unscrupulous designers could use anywhere. ARM makes almost all its money by licensing processors, so the company fears IP thievery the way the record industry fears *Kazaa*. A few years ago, ARM allowed Altera to introduce some Excalibur FPGAs with integrated ARM922T processor cores, but the processor was safely embedded in a silicon layer as a hard core, not in programmable logic as a soft core.

Actel is taking a wholly different approach. The new ARM-enabled FPGAs—based on Actel's ProASIC3 architecture—will encrypt a synthesizable ARM7 core in the programmable logic. Actel believes its security technology can protect the gate-level netlist from prying eyes and hacker tools. In addition to guarding the ARM core, the security is useful for

encrypting any IP that Actel's customers might add to the FPGA. Indeed, the main reason for implementing the ARM7 in programmable gates instead of embedding it as a hard core is to allow customers full access to the processor's I/O interfaces, so they can more easily attach their application-specific logic. ProASIC3 security doesn't interfere with the normal operation of either the processor or the custom IP. As Figure 1 shows, Actel's technology also allows customers to deploy chips that are upgradable in the field, even remotely over a network.

To reduce the loss of performance that typically results when a synthesizable processor runs in programmable logic, Actel and ARM are optimizing an ARM7TDMI-S processor core specifically for ProASIC3 devices. ARM says this special implementation will run faster in programmable logic than will the standard soft version, while maintaining full software compatibility. Even so, some loss of performance is inevitable. ARM estimates the firm-core ARM7TDMI-S will run at about 40MHz in a ProASIC3 device. Although that's fast enough to compete with low-end ARM-based MCUs that lack programmable logic, it's slower than some alternatives.

For instance, the hard-core ARM922T in Altera's Excalibur runs at clock frequencies up to 200MHz, even though Excalibur is an older-generation line of FPGAs. (Altera hasn't embedded any ARM cores in its newer Stratix and Cyclone families of FPGAs.) In addition to Excalibur, Altera has a soft-core processor called Nios II, designed specifically for FPGA integration, and it runs at 140–180MHz in a fast Stratix or Stratix II device. Nios II is a 32-bit RISC processor with a deeper pipeline than the ARM7TDMI has (six stages



**Figure 1.** Actel's security technology allows developers to encrypt and embed their IP in a ProASIC3 FPGA alongside the embedded ARM7TDMI-S processor core. The security is also built into Actel's Libero software-development tools. This technology uses the Advanced Encryption Standard (AES) with 128-bit keys to protect a bitstream in transit, so customers can securely modify the contents of a ProASIC3 device over a network or the Internet.

vs. three), and it has additional advantages: dynamic or static branch prediction, configurable instruction and data caches, and an extendable instruction set. However, Nios II has two disadvantages: all instructions are 32 bits long (the ARM7TDMI has 16-bit Thumb instructions for greater code density), and it's a proprietary Altera architecture, not the industry-standard ARM architecture. (See *MPR 6/28/04-02*, "Altera's New CPU for FPGAs.")

More alternatives are available from Xilinx and Stretch. Xilinx sells a Virtex-II Pro FPGA with an embedded PowerPC 405 hard core, and the company also offers a 32-bit RISC soft core known as MicroBlaze. The hardened PowerPC 405 is much more powerful than a firm-core ARM7TDMI, but Actel's ProASIC3 chips should be significantly less expensive. Xilinx's MicroBlaze soft core is similar to Altera's Nios II, sharing many of the same advantages and disadvantages. (See *MPR 11/5/01-03*, "FPGAs Catch Fire At MPF," and *MPR 6/28/04-02*, "Altera's New CPU for FPGAs.") Stretch's S5000 chips are a hybrid technology: they are standard parts based on a Tensilica Xtensa V processor core, but they also have some programmable logic for adding application-specific accelerators. (See *MPR 4/26/04-01*, "Stretching Performance.")

Further analysis must wait until later this year, when Actel and ARM release more technical information and pricing in advance of delivery. Actel's ARM-enabled FPGAs will likely be priced in the midrange to high end of the ProASIC3 family. At present, the family ranges from small 30,000-gate devices priced from \$1.50 to three-million-gate behemoths costing less than \$100. To leave enough gates for customers to add meaningful amounts of application-specific logic, the devices with ARM7 cores must be relatively large. For that reason, these devices will compete indirectly with other ARM-based MCUs and even less directly with conventional MCUs. Instead, Actel will probably nibble away at ASICs and SoCs, luring project managers who are reeling from the high costs of mask sets and other nonrecurring engineering (NRE) expenses.

### Oki's Incredibly Shrinking ARM

Oki's claim to fame is to offer the world's smallest ARM-based MCU. In some embedded systems, circuit boards are as crowded as Tokyo subway trains, forcing designers to use 8- or 16-bit chips instead of more-capable 32-bit processors. By June, there will be an alternative: the very smallest of Oki's new Advantage MCUs will cram a 33.3MHz ARM7TDMI processor core, 16KB of SRAM, and 64KB or 128KB of flash memory into a wafer-level chip-size package (WCSP) less than 25mm<sup>2</sup>—about one-fourth the size of a dime. Yes, that's the whole package, not just the die. Nearly two dozen of these packaged parts would fit on the unpackaged 596mm<sup>2</sup> die of Intel's recently announced Montecito Itanium-2.

So far, Oki has announced two lines of MCUs in the Advantage family: the tiny ML67Q405x series, packaged in thin quad flat packs (TQFP); and the tinier ML67Q406x series, packaged in the aforementioned WCSP, TQFP, or

low-profile fine-pitch ball-grid array (LFBGA). All chips in the 405x and 406x series have a 33.3MHz ARM7TDMI processor, 16KB of internal SRAM, 64KB or 128KB of flash memory, and numerous on-chip peripherals. Prices range from \$4 to \$5.50. The 405x chips are the luxury models, sporting an external bus for memory or general-purpose I/O (GPIO). In contrast, the 406x chips are so miniaturized they dispense with external memory altogether and must execute their software entirely in on-chip SRAM. Eliminating the memory bus dramatically reduces the pin count. Whereas the 405x chips have 144 pins, the 406x chips have only 64 or 84 pins.

These little MCUs would be even smaller if Oki weren't manufacturing them in a fabrication process imported from the last century: 0.22 micron, with aluminum interconnects. (To be fair, most vendors of 8- and 16-bit MCUs are using even older processes.) However, Oki's trailing-edge technology should make current leakage insignificant, and it provides some design flexibility. The cores run at 2.5V ( $\pm 10\%$ ) and the I/O rings tolerate 2.5V ( $-10\%$ ) to 3.3V ( $+10\%$ ), so designers can cut corners by supplying both rails with the same voltage, if they wish. Unfortunately, the relatively high voltage of these chips will probably keep their power consumption from being as impressive as their diminutive physical dimensions are. Oki is still characterizing the early samples and hasn't estimated power consumption, other than to say the chips will draw only 25 microamps in stop mode. (In that sleep-like mode, the internal SRAM retains data and the real-time clock can continue running, allowing instant power-up on demand.)

Anticipating that designers will use the 405x and 406x in hard real-time systems, Oki has taken extra steps to guarantee deterministic response. The ARM7TDMI cores have no caches and no prefetch buffers. Instead, the processor can store critical instructions and data in the single-cycle internal SRAM and zero-wait-state 32-bit flash memory. Oki claims these MCUs are more deterministic than are competing ARM7-based MCUs from Philips that have wider buses, prefetch buffers, and slower flash memory.

### Stuffed With Peripherals, Too

As mentioned before, I/O capabilities distinguish Oki's 405x chips from their smaller 406x cousins and account for the different packages. The 405x series has a 32-bit memory bus (which supports 8-, 16-, or 32-bit memory), 23-bit memory addressing, and more-flexible GPIO capabilities. The 23-bit addressing allows up to 16MB of memory space, with one bank each for RAM and ROM, plus two banks for memory-mapped I/O.


There are 40 dedicated GPIO pins, but designers can use another 68 of the memory-bus pins for GPIO as well. An Inter-IC Sound ( $I^2S$ ) interface supports audio applications with 24-bit samples and sampling rates as high as 48kHz. Two nine-bit UARTs are compatible with eight-bit designs and have FIFO buffers to offload some work from the processor. There's also a real-time clock, six 16-bit timers, a four-channel

analog-to-digital converter, a multimaster  $I^2C$  interface, a SPI interface, and other miscellaneous features.

For single-chip systems that don't need external memory, the smaller 406x series eliminates the memory bus to achieve a 64- or 84-pin package. Some parts in this series are available with 128KB of flash memory instead of 64KB, to compensate for their lack of external RAM. Their GPIO capabilities vary slightly, but, as a whole, the 406x chips are more limited in this respect than their 405x cousins are. Package types include the aforementioned 64-pin WCSP (the smallest, at 4.84mm  $\times$  5.09mm); a 64-pin TQFP (10mm  $\times$  10mm); and an 84-pin LFBGA (9mm  $\times$  9mm).

An interesting feature of the WCSP is its 8  $\times$  8 array of pins, which evenly covers the entire mounting side of the package without using the C4-style solder-bump technology seen in desktop PC and server processors. Oki's WCSP uses a proprietary technology that routes fine wires from the distributed pins to the die's pads, which are conventionally located around the edges of the die. This method requires several extra manufacturing steps, but it allows Oki to economically produce a very small package with manageable pin spacing (the pitch is 0.5mm). The packaged chip weighs only 0.03 grams. Figure 2 is a macrophotograph of WCSP pin wiring inside the package.

One drawback of the WCSP's unusual pin layout is that it requires more routing layers in the system board—perhaps as many as six or eight layers vs. one or two layers for the same chips in the larger TQFP and LFBGA packages. The extra cost of a more complex board is tolerable only in more-expensive systems that must fit into very small spaces—examples might be cellphones, wireless Bluetooth modules, and portable biometric devices. Nevertheless, space-conscious designers will find Oki's Lilliputian MCUs a compelling alternative to 8- and 16-bit chips.



**Figure 2.** Oki Semiconductor uses a proprietary manufacturing and packaging process to redistribute pins over the entire mounting surface of its wafer-level chip-size package (WCSP). Notice the tiny wires connecting the contact surfaces to the pads, which are located around the periphery of the die in a conventional manner. This technology allows Oki to make a packaged 64-pin chip requiring a mounting area only 24.64mm<sup>2</sup>. (Photo: Oki Semiconductor)

### Philips Targets Low Power, High Performance

With more than 250 ARM-based designs under its belt, Philips has more experience with ARM than does any other company and is already a major player in 32-bit MCUs with its LPC2000 family of ARM7TDMI-based chips. (See *MPR 5/19/03-01*, “Philips Shows Flashy MCUs.”) Now, Philips is introducing the LPC3000 family, the first ARM9-based MCUs manufactured in a 90nm process. By using a more powerful processor core and cutting-edge fabrication technology, Philips is aiming for both higher performance and lower power consumption.

All LPC3000 parts will use the ARM926EJ-S, a soft processor core with 32KB instruction/data caches, an MMU, translation lookaside buffer (TLB), 16-bit Thumb instructions, DSP extensions, and Jazelle extensions for Java. In the initial devices, the processor runs at 200MHz. Those are impressive features, but Philips also added a vector floating-point coprocessor licensed from ARM, a DMA controller, and a DDR memory controller. Consequently, LPC3000 chips are muscular MCUs, suitable for heavy-duty embedded applications. They can run Embedded Linux, Palm OS, Symbian OS, and Windows CE—certainly not the kind of operating systems found in the company of 8- and 16-bit

MCUs, or even most other 32-bit MCUs. Philips is validating the first silicon now and expects to start production in 3Q05.


Figure 3 is a block diagram showing the basic features of an LPC3000. Frankly, it looks more like somebody’s pet-project SoC than a general-purpose MCU. The vector FPU is IEEE-754 compliant, runs at the same clock frequency as the closely coupled ARM9 processor core, and executes as many as three 32- or 64-bit floating-point operations simultaneously. The DDR-1 memory controller has a 32-bit data interface and 32-bit addressing, allowing up to 4GB of memory. The real-time clock has its own power domain, so it can stay awake while the rest of the chip sleeps. Among the external I/O interfaces is a Memory Stick interface, included for a special customer Philips won’t name. (*Microprocessor Report* ventures out on a limb and guesses the customer’s name begins with “S” and ends with “Y.”) Everything is tied together with an AMBA high-speed on-chip bus (AHB).

The only scarcity on these MCUs is on-chip memory: with merely 64KB of SRAM and no flash, the first LPC3000 parts are less well endowed than some MCUs in much lighter weight classes. However, Philips says future members of the family will have 2MB to 4MB of on-chip flash memory. In addition, some future devices will have Ethernet controllers, Hi-Speed USB 2.0 host/device controllers, and other step-up features. Downsized family members may lose the MMU and perhaps some other blocks, though Philips says the FPU is so small in 90nm, it’s hardly worth discarding.

Considering all those features, one might wonder how Philips expects the LPC3000 to achieve its low-power goals, even when fabricated in 90nm. Philips says it designed the chips using a special low-power library and a 90nm process tuned for low power. The fabrication process was jointly developed by the Crolles2 Alliance, which includes Philips, Freescale, and STMicroelectronics. The CMOS process uses copper interconnects and low-*k* dielectrics but not silicon-on-insulator (SOI) technology. Philips will manufacture the chips on 300mm wafers in the Crolles2 Alliance pilot fab in Crolles, France. The specially tuned 90nm process is supposed to reduce current leakage to levels more typical of a 0.18-micron process.

Philips is still testing the first silicon, so typical power consumption isn’t yet confirmed, but the company is aiming for 250mA at a core voltage of 1.2V—or about 300mW at 200MHz. (I/O voltage is a compatible 3.3V, and the DRAM interface is 1.8V.) If Philips achieves that target, it will be an impressive feat for such heavily equipped MCUs.

In addition, the devices have a low-power mode that drops the core to 14MHz at 0.9V, which will draw only about 5mA. Nevertheless, the processor is still capable of executing about 20 MIPS in such a stupor—as much horsepower as some 8- or 16-bit MCUs can



**Figure 3.** Philips LPC3000 block diagram. Is it a microcontroller or an SoC? The first LPC3000 chips will have a 200MHz ARM926EJ-S processor core, a 64-bit vector FPU, a DMA controller, a 32-bit DDR-1 memory controller, 32-bit memory addressing, and other high-end features alien to most store-bought MCUs. Note the Memory Stick interface (upper right), added for an anonymous customer that couldn’t live without it.

muster while wide awake. The combination of high performance and low power consumption at 90nm will give the LPC3000 family a clear advantage over other 32-bit MCUs manufactured in older fabrication processes, and the 300mm wafers should produce economies of scale that keep prices low as well.

### ARM Will Dominate 32-Bit MCUs

Industry analysts expect the market for 32-bit MCUs to double in size before the end of this decade, and ARM has already staked out so much territory that it's hard to see how any other microprocessor architecture can muscle in. ARM penetrated this market early and continues to aggressively pursue it. Last year, ARM even tried to acquire Triscend, which would have made ARM a minor MCU vendor as well as the leading supplier of 32-bit IP for MCUs. (See *MPR 3/15/04-02*, "Xilinx Reconfigures Triscend.")

Certainly, there are other 32-bit microprocessor architectures suitable for MCUs. ARC International, MIPS Technologies, and Tensilica all have licensable 32-bit processor cores small enough and powerful enough to compete with ARM. Their cores are often more affordable as well. So far, however, the vast majority of ARC, MIPS, and Tensilica licensees are designing custom chips for special applications, not general-purpose MCUs sold as standard parts. ARM's growing strength in MCUs will hurt these competitors in three ways.

First, the rapidly rising 32-bit MCU market generates significant royalties for ARM and fuels the company's further expansion and new-product development. (See *MPR 9/7/04-01*, "ARM Extends Its Reach.") Second, the growing selection and variety of 32-bit MCUs as standard parts will reduce the need to develop custom ASICs and SoCs, thus shrinking the licensing opportunities for other IP companies. Third, the propagation of ARM-based MCUs boosts the popularity of the ARM architecture, thereby encouraging even more support from software vendors and developers. In sum, the MCU business accelerates ARM's momentum.

That momentum will increase as ARM's Cortex-M processors seep into the market. This new family of cores—announced late last year at ARM's Developers' Forum and In-Stat's Processor Forum Taiwan—uses Thumb-2 instructions exclusively. The first member of the Cortex-M family is

### Price & Availability

Actel's ProASIC3 FPGAs with embedded ARM7TDMI-S processor cores will ship this year, but the company hasn't announced specific availability or pricing. However, Actel notes that customers won't be required to take an ARM license or pay royalties for ARM-based ProASIC3 devices. For more information, see: [www.actel.com/products/ip/ARM7.html](http://www.actel.com/products/ip/ARM7.html)

Oki Semiconductor has some samples of its ARM7TDMI-based Advantage MCUs available now and plans to start production of all parts in May or June. The smallest ML67Q406x devices in wafer-level chip-size packages (WCSP) will cost about \$4.50 in 10,000-unit quantities. The slightly larger ML67Q405x chips with 64KB of flash memory in 64-pin thin quad flat packs (TQFP) will start at \$3.98 and rise to \$4.50–\$5.10 with 128KB of flash; in 144-pin packages, prices will range from about \$4.00 to \$5.50. For more information, see [www2.okisemi.com/site/productscatalog/armsolutions/Overview.html](http://www2.okisemi.com/site/productscatalog/armsolutions/Overview.html).

Philips plans to ship the first LPC3000-series MCUs with ARM926EJ-S processor cores in 3Q05. Initial parts will run at 200MHz and cost about \$10 in 10,000-unit quantities. Philips will announce more details later. For more information, see [www.semiconductors.philips.com/news/content/file\\_1139.html](http://www.semiconductors.philips.com/news/content/file_1139.html).

the Cortex-M3. It offers high code density and low power consumption while retaining nearly all the advantages of a 32-bit processor, including a software-compatible upgrade path to ARM's full 32-bit instruction set. (See *MPR 11/29/04-01*, "ARM Debuts Logical V7.")

Vendors of 8- and 16-bit MCUs will suffer from ARM's inroads into their market, especially as 32-bit MCUs replace the higher-value low-end parts. But 8- and 16-bit chips will survive for many years to come, because there will always be room at the bottom for dirt-cheap MCUs with just enough processing power for deeply embedded applications. Meanwhile, ARM and its innovative licensees will be more than happy to skim off the cream of the market. We call that a winning strategy. ♦

To subscribe to Microprocessor Report, phone 480.483.4441 or visit [www.MDRonline.com](http://www.MDRonline.com)