Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

Philips Shows Flashy MCUs

ARM7-based Microcontrollers Have Embedded Flash Memory

By Tom R. Halfhill {5/19/03-01}

Eight-bit chips still account for 56% of microprocessor sales by volume and 40% of revenues, according to *World Semiconductor Trade Statistics*. Unit sales of the most popular 8-bit architecture, Intel's disco-era 8051, average 3.3 billion parts per year—about 30 times the volume of 32-bit PC processors. For that matter, even the lowly 4-bit chips that hearken back to the first microprocessor in 1971 persistently sell in volumes only 15% below their all-time peak. Embedded-systems developers keep using these puny chips because they are unbeatably cheap, sip miniscule amounts of power, and are small enough to add a dab of silicon intelligence to almost anything large enough to see.

Hoping to displace 8- and 16-bit chips with more-powerful (and more-profitable) devices, Philips Semiconductors is introducing a new line of ARM7-based 32-bit MCUs. Philips knows it can't beat the smallest MCUs on their own terms, but the company figures that a wisely downsized 32-bit MCU might deliver enough extra performance to lure some developers away from their frugal habits. To sweeten the bait, Philips is fabricating the new MCUs in a special 0.18-micron CMOS process that offers the option of embedded flash memory.

Not a Flash in the Pan

All the MCUs in Philips's new LPC2100 family use an ARM7TDMI-S core that includes 16-bit Thumb instructions, debug extensions (including a real-time monitor, real-time trace, and EmbeddedICE), and a 32-bit multiplier. The ARM7—with its simple three-stage pipeline and von Neumann memory architecture—is the smallest 32-bit RISC core that enjoys broad popularity and industry-wide tool support. Still, at 31K gates, an ARM7TDMI-S is three times the size of an 80C51, which has about 9.5K gates. However, that difference isn't very significant when using 0.18-micron design rules, which allow densities of 100K gates per square millimeter and SRAM cells measuring only 4.65 square microns.

The size difference becomes even less significant when the processor core is integrated with peripherals and on-chip memory, the real silicon hogs. In an embedded system that requires network connectivity (an increasingly popular feature), the Ethernet media-access controller alone can require 60K–90K gates. Likewise, any useful amount of on-chip SRAM will probably dwarf the amount of silicon the processor occupies.

To cut costs and power consumption, Philips is manufacturing the initial LPC2100 chips in a 0.18-micron process that allows 60MHz operation at 1.8V. Philips claims it's the industry's first 0.18-micron CMOS process with embedded flash memory. The zero-wait-state flash is based on two-transistor cells with access times in the 50ns range, and it connects to the processor core over an uncommonly wide 128-bit interface. This allows the processor to retrieve four 32-bit words at a time without the long latencies normally associated with flash. The flash controller also does some intelligent prefetch buffering, so instructions stored in flash memory are immediately available when the processor must handle an interrupt service routine. Philips provides a flash loader boot-block program that can download a user program over a serial port and update the system in the field.

Price & Availability

The first three Philips LPC2100-series 32-bit MCUs are available now at prices ranging from about $5 to about $8 in 10,000-unit quantities. For more information, visit *www.semiconductors.philips.com/news/content/file_949.html*.

So far, all three devices in the LPC2100 family have 128K of embedded flash, but future chips will offer as little as 64K or as much as 1MB. Philips says parts with 256MB of flash will be available later this year. The 1MB parts will arrive when Philips begins production in a smaller 0.13-micron process, scheduled for 2004.

Charter members of the LPC2100 family differ from each other largely in their amounts of on-chip SRAM. The LPC2104, LPC2105, and LPC2106 have 16K, 32K, and 64K of SRAM, respectively. Larger on-chip memories make the LPC2105 and LPC2106 more suitable for networking applications that handle protocol stacks. Philips says future chips in the family will add logic for Ethernet, USB, 802.11, control-area networks (CAN), and analog-to-digital converters.

Two Ways to Integrate Flash

Philips isn't the first company to make ARM-based MCUs, of course. It's not even the first company to integrate flash memory. Atmel, Hynix, and Oki also offer general-purpose MCUs based on the ARM7TDMI core with flash, and all their parts have more flash than the new Philips MCUs have. The most formidable competitor is Atmel, whose populous AT91 family includes four such chips—some with as much as 2MB of flash, or 16 times more than Philips's best LPC2100-series device now available. The Atmel MCUs also tend to have more SRAM, and two of them run at higher clock speeds than the Philips chips.

However, there's an important difference among the ARM7-based MCUs from Atmel, Hynix, Oki, and Philips: only the ones from Philips and Hynix actually embed the flash memory in the same die as the processor core. Atmel and Oki integrate a separate flash chip with the processor in a multichip package. Naturally, this has important consequences for memory performance, package size, power consumption, and cost.

In Atmel's MCUs, the flash interface is only 32 bits wide, not 128 bits as in the Philips MCUs, and the access time is about 110ns, only half the speed of the Philips flash memory. One result is that Atmel's MCUs can deliver their full performance only when executing code from SRAM, not from flash—unlike the Philips MCUs, which can run programs from zero-wait-state flash memory without a performance penalty. Oki's MCUs are even more constricted, because the multichip interface linking the processor to the flash memory is only 16 bits wide. Hynix uses a 16-bit interface, too, even though the flash is embedded with the processor core in a single chip. The access time for Hynix's flash memory is about 90ns—about half as fast as the flash memory in the Philips MCUs.

Another benefit of Philips's 0.18-micron embedded-flash process is a smaller chip, which should reduce power consumption, package size, and cost. At only 7 × 7mm, the Philips MCUs are less than half the size of their nearest competitors, Atmel's 10 × 10mm AT91FR4042 and AT91FR40162. Oki's ML67Q400x/500x devices have by far the largest packages in this group: at 20 × 20mm for the low-profile quad flat-pack (LQFP), they are four times larger than the Philips devices.

Although Atmel's MCUs tend to have more SRAM than the Philips MCUs, they need it to compensate for their slower flash memory—and SRAM is larger and more expensive than flash. Unfortunately, we were unable to obtain power-consumption ratings, package sizes, and quantity pricing for every device in this comparison, but the core voltages and other parameters leave little doubt that Philips's 0.18-micron embedded-flash process will give the LPC2100 family an advantage in very small embedded systems. Table 1 summarizes the features of these general-purpose ARM7 MCUs with integrated flash memory.

All these ARM7-based MCUs are bargain priced for 32-bit processors, but keep in mind that 8-bit chips average about $1.40 in the same volumes. A few bucks may seem a trifling difference, but it's significant for low-end embedded systems.

Smaller ARM, Larger Thumb?

ARM could help narrow the price gap between 8- and 32-bit MCUs by introducing a smaller core especially designed for MCU integration that also expands the subset of Thumb instructions. Right now, ARM programs must switch back and forth between 16- and 32-bit modes because Thumb isn't a complete instruction set: among other things, it can't handle exceptions or some system-control functions or access the entire register file. (See *MPR 3/27/95-01*, "Thumb Squeezes ARM Code Size.") If Thumb were more comprehensive, developers could write entire programs in compact 16-bit code while retaining most of the advantages of a 32-bit architecture.

An ARM core with comprehensive Thumb instructions might resemble Hitachi's first-generation SuperH architecture, which couples a 32-bit RISC architecture with a space-saving 16-bit instruction set. (See *MPR 3/6/95-02*, "Hitachi SH-3 Hits 100 MIPS.") Last year, ARC International took a novel approach with its new ARCompact instruction-set architecture, which allows programmers to write 16- or 32-bit code or freely mix together both types of instructions. (See *MPR 2/18/03-06*, "Soft Cores Gain Ground," and the sidebar, "ARCompact: An Elegant 16/32-Bit ISA.") The base-case configuration of an ARCtangent core is about 16K gates, or about half the size of an ARM7, so it's possible for a 32-bit RISC core to compete on nearly equal terms with 8- and 16-bit architectures.

	Philips LPC2104	Philips LPC2105	Philips LPC2106	Atmel AT91FR40xx	Atmel AT91Fxxxxx	Hynix HMS39C70x	Oki ML67Q400x	Oki ML67Q500x
CPU Core	ARM7TDMI-S	ARM7TDMI-S	ARM7TDMI-S	ARM7TDMI	ARM7TDMI	ARM7TDMI	ARM7TDMI	ARM7TDMI
Frequency	60MHz	60MHz	60MHz	82MHz	40–47MHz	50MHz	33MHz	60MHz
Performance	54 mips	54 mips	54 mips	63 mips	36 mips	45 mips	30 mips	54 mips
Integrated Flash	128K On chip	128K On chip	128K On chip	512K–2,048K In package	1,024K–2,048K In package	192K–512K On chip	256K–512K In package	256K–512K In package
Flash I/F	128-bit	128-bit	128-bit	32-bit	32-bit	16-bit	16-bit	16-bit
SRAM	16K	32K	64K	256K	8K–136K	4K–12K	32K	32K
ROM	No	No	No	No	No	No	4K	4K
RT Clock	Yes	Yes	Yes	No	No	No	No	No
Timers	2 x 32-bit + watchdog	2 x 32-bit + watchdog	2 x 32-bit + watchdog	3 x 16-bit + watchdog	3 x 16-bit + watchdog	6 x 16-bit + watchdog	7 x 16-bit + watchdog	7 x 16-bit + watchdog
I²C	1	1	1	No	No	No	1	1
SPI	1	1	1	No	No	No	No	No
UART	2	2	2	2 USART	2 USART	2	2	2
GPIO	32	32	32	32	32	70–81	42	42
Other	—	—	–	Para I/O Ctrl	Para I/O Ctrl	ADC, DAC	DMA, PWA x 2	DMA, PWA x 2
IC Process	0.18µm	0.18µm	0.18µm	N/A	N/A	0.35µm	0.22µm	0.25µm
Core Voltage	1.8V	1.8V	1.8V	1.65V–1.95V	2.7V–3.6V	3.0V–3.6V	3.6V	3.6V
Power	54mW	54mW	54mW	N/A	N/A	~99mW	NA	NA
Operating Temp*	0°–70°C -40°–85°C	0°–70°C -40°–85°C	0°–70°C -40°–85°C	-40–85°C	-40–85°C	0–70°C	-40–85°C	-40–85°C
Package (size)	LQFP-48 7 x 7mm	LQFP-48 7 x 7mm	LQFP-48 7 x 7mm	BGA-121 10 x 10mm	BGA-120 11 x 17mm	TQFP-100 14 x 14mm	LQFP-144 20 x 20mm	LQFP-144 20 x 20mm
Package (size)	HVQFN-48** 7 x 7mm	HVQFN-48** 7 x 7mm	HVQFN-48** 7 x 7mm	—	—	—	LFBGA-144 11 x 11mm	LFBGA-144 11 x 11mm
Availability	Now	Now	Now	Now	Now	Now	3Q03	3Q03
Price (units)	<$5 (10K)	~$6 (10K)	<$8 (10K)	N/A	N/A	N/A	$5–$6 (100K)	$6–$7 (100K)

Table 1. The new ARM7-based MCUs from Philips have less embedded flash memory and SRAM than some MCUs from other vendors, but their flash interfaces are faster. Moreover, their lower core voltages and smaller packaging should give the Philips chips an advantage in low-power applications. N/A: data not available. *Each Philips MCU is available as two part numbers certified for different operating temperatures. **HVQFN (heatsink very thin quad flat-pack no-leads) is a chip-scale package that uses Amkor's MLF (MicroLeadFrame) packaging technology.

Until a smaller ARM core designed for MCUs becomes available, the ARM7 is the best choice for companies that prefer the world's most popular 32-bit embedded architecture. By integrating flash memory on chip, ARM7-based MCUs enable the design of very small self-contained embedded systems that can be easily upgraded on site or over networks. ◇