CIRCUIT CELLAR®
THE MAGAZINE FOR COMPUTER APPLICATIONS

SILICON UPDATE

Tom Cantrell

Atmel Gets Huge

Being involved with both print and Internet publications keeps Tom busy, but it also gives him a double-barrel opportunity to cover topics like Atmel's latest releases. Read on to find out why their new $20 8-bit MCU is much more than just an AVR on steroids.

No doubt the emergence of the Internet as a publishing medium is a big deal. I don't think it spells the end of paper, but I do see both electronic and print media coexisting with and complementing each other.

Print media remains my favorite for the under-the-hood stuff. It just seems easier to ponder weighty material on a page than on a screen. Perhaps that's because print works best for detailed artwork, such as schematics and high-resolution photos that are critical for getting a complicated message across. Paper isn't going away, at least not until the Wizards come up with a color screen you can roll up and shove in your pocket.

The web does have some compelling advantages. Notably, the ability to deliver software and other intellectual property (not to mention take your money) electronically. Old-timers may remember early efforts by magazines to publish software using barcodes and the like. I seem to recall an especially laughable scheme involving flimsy, paper-backed 45-rpm records you were supposed to try (over and over…) to playback through your cassette interface.

By far, the biggest advantage for the web is it's much faster and cheaper to squirt bits onto your screen than it is ship them to your front door.

However, from my perspective as both print and online columnist, the reduced latency of the web can lead to interesting time warp effects. I have to be careful not to write something that will appear on the web tomorrow that refers to something I wrote for paper yesterday, which won't get printed until the day after tomorrow. Got that?

About the time you read this on the web, my print column entitled "Atmel Gets Tiny" will hit the streets. I actually wrote the words in that article some time ago; and in the meantime, Atmel hasn't been standing still. Fortunately, thanks to the web, I can deliver my words as fast as they can deliver the chips.

FEELING PRETTY SLIC

"Atmel Gets Tiny" is all about the tiny parts with low prices that are fueling growth in the high volume mainstream of the "billions served" 8-bit MCU market.

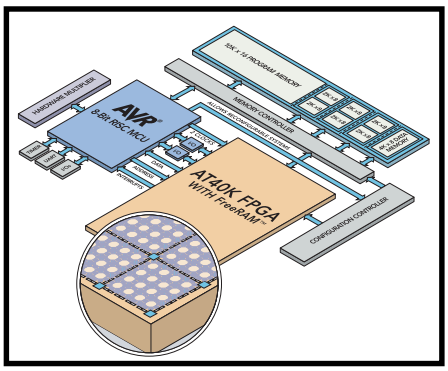

Figure 1—*By combining an AVR MCU, SRAM, and FPGA, Atmel delivers the AT94K, a System-on-Chip for the masses.*

Device	AT94K10	AT94K20	AT94K40
FPGA Gates	10K	20K	40K
FPGA Core Cells	576	1024	2304
FPGA SRAM Bits	4096	8192	18432
FPGA Registers (Total)	864	1408	2880
Max FPGA User I/O	144	192	288

Table 1—*The first three members of the AT94K family differ only in the characteristics of the integrated AT40K FPGA.*

"Atmel Gets Huge" is about an 8-bit MCU that costs $20 and may well be worth it. Why? Because it isn't just an AVR on steroids, it represents a new way of thinking for system designers. Short and sweet: the AT94K FPSLIC (field programmable system-level integrated circuit) combines the AVR with an AT40K FPGA (see Figure 1). Quite slick indeed!

So far, three versions (see Table 1) of the AT94K have been announced that differ in terms of the FPGA gate (10 k, 20 k, 40 k) and I/O count (144, 192, 288).

The FPGA is literally a drop-in of Atmel's existing AT40K parts, which feature look-up table (LUT) cell architecture and distributed 10-ns SRAM. Indeed, the AT94K is pin compatible with like-sized members of the AT40K family and features the same 3.0–3.6-V power supply and PCI-compliant 24-mA I/O structures. It's really as much an FPGA with a CPU built in as it is a CPU with an FPGA.

MEMORY MIPS

The onboard AVR shown in Figure 2 isn't just a core, but includes the full complement of peripherals found in a standalone MCU including fancy timers and counters, UARTs, and all the rest. However, although an MCU typically relies on OTP or flash memory for program storage, the AT94K goes with on-chip SRAM.

The good news is, the SRAM is fast enough (15 ns) to deliver one-cycle operation at up to 40 MHz, which translates into about 30 MIPS for the AVRs modern pipeline RISC design. For flexibility, the 36 KB of MCU SRAM is able to be partitioned between x16 program store and x8 data RAM. For example, one setup might dedicate 20 and 16 KB for program and data respectively, another 32 and 4 KB.

The bad news is, the SRAM (both MCU program/data and FPGA configuration) has to be initialized at powerup, typically via external serial EEPROM. Note that the volatility of the SRAM can be exploited as a security feature, making it difficult for a would-be cracker to get at your code.

The FPGA features the Atmel cache-logic partial FPGA reconfiguration scheme that allows a portion of the FPGA logic to be dynamically changed at runtime, without disrupting overall operation or pin state. It's in-operation, not just field or in-system programmable.

MAKING A CONNECTION

As important as the individual features of the MCU and FPGA, is the way the two are connected. Based on the limited information at hand, it appears Atmel has adopted a simple approach that treats the FPGA much like another onboard 8-bit peripheral. There's an address decoder for generating up to 16 pseudochip selects into the FPGA and, going the other way, 16 interrupt lines that are fed from the FPGA into the MCU.

A loosely coupled scheme makes sense in the context of Atmel's cut-and-paste approach to building chips. Besides putting together an MCU and FPGA, they also offer FPGA with gate array, and I expect they'll offer more combinations. Keeping the interconnect between modules simply facilitates reuse by minimizing the tweaks necessary to glue together the pieces.

In the case of the AT94K, some accommodation is made to support synchronism between the MCU and the FPGA. The MCU has access to the FPGA's eight global clocks and can drive two of them relying on its own combination of internal and

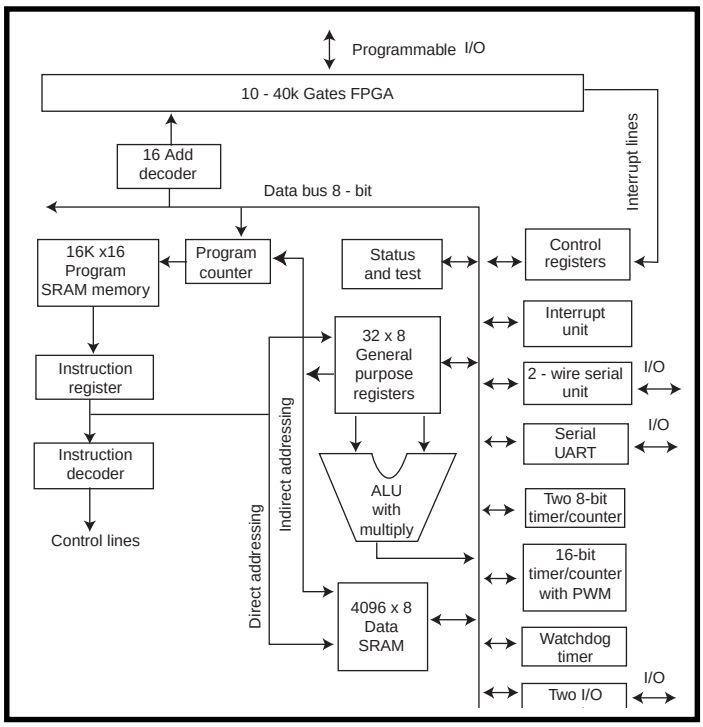

Figure 2—*Plenty of fast on-chip SRAM, including 36 KB for the AVR and distributed SRAM in the FPGA, makes for speedy performance.*

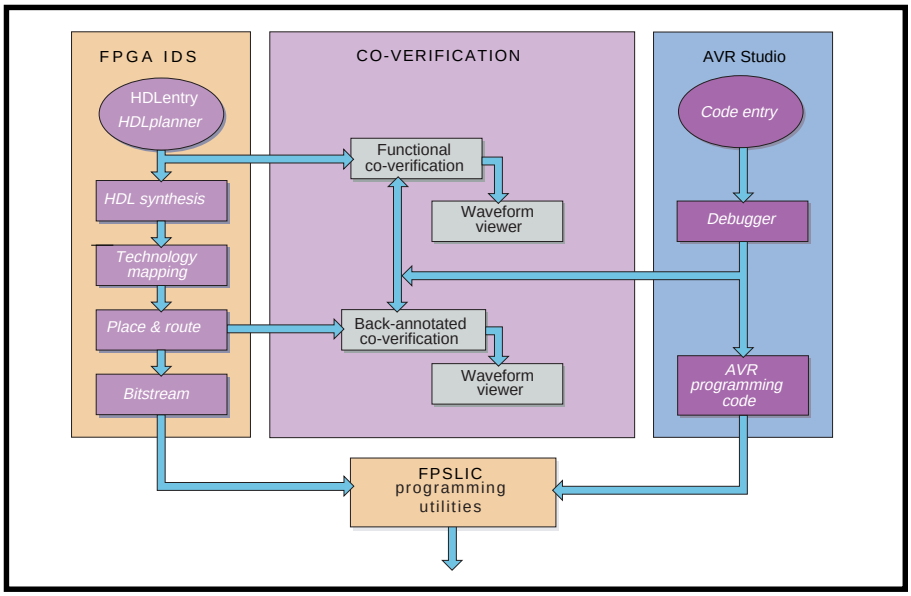

Figure 3—*At only $495 for a yearly subscription (just think of it as a new approach to upgrades), System Designer delivers codesign capability to cost-conscious developers.*

external oscillators, clock dividers, timer/counters, and so on.

CO(DE)SIGNER

Getting MCU software working is hard. Getting an FPGA working is hard. Needless to say, getting an MCU and FPGA working together isn't easy.

Now, when you're talking "Codesign" EDA Tools, the big question is who's the cosigner for the loan. Only the most well-heeled customers can afford a full-blown ASIC-class EDA capability.

Atmel comes to the rescue with their System Designer tool suite (see Figure 3) that coordinates MCU and FPGA development with source-level debug and full hardware visibility. For implementation, the package includes place-and-route, floor planning, macro generators, and bitstream utilities.

The AT94K press materials also allude to third-party tool support from the likes of Mentor Graphics. It seems likely the combined lineup of existing AVR and AT40K FPGA tool providers will weigh in with more support, perhaps even pairing up to combine their hardware and software tools into integrated codesign packages.

HUGE IDEA

Atmel isn't the first to come up with the idea of combining hard and soft logic on a single die. However, I think they have a good shot at moving the concept forward.

A key factor in the success or failure of a particular hard and soft IC is the degree to which each portion is popular in its own right. For instance, sometime ago Motorola announced, then quickly abandoned, the so-called Core Plus that combined a ColdFire CPU with their SRAM-based FPGA design. I think this is a case where neither side of the equation was very compelling, especially their FPGA, which never really got off the ground at all.

Although neither the AVR nor AT40K is a dominant front-runner in their respective markets, both are worthy, credible, and proven contenders. Furthermore, Atmel is on quite a roll as demonstrated by 1999 third-quarter revenues up 9% over second quarter, and a whopping 24% over the third quarter in 1998. Atmel's been acquiring companies left and right (one example is Temic) and even made a deal to buy Hitachi's state-of-the-art 8² wafer fab in Texas. In short, Atmel knows how to execute and has quite a head of steam.

What does the new millennium mean for designers? Gaze in the crystal ball, and I bet you'll find a field-programmable System-On-Chip in your future.

Tom Cantrell has been working on chip, board, and system design and marketing in Silicon Valley for more than ten years. You may reach him by e-mail at tom.cantrell@circuitcellar. com, by telephone at (510) 657-0264, or by fax at (510) 657-5441.

SOURCE

AT94K FPSLIC
Atmel Corp.
(408) 441-0311
Fax: (408) 436-4200
www.atmel.com

Circuit Cellar, the Magazine for Computer Applications. Reprinted by permission. For subscription information, call (860) 875-2199, subscribe@circuitcellar.com or www.circuitcellar.com/subscribe.htm.