


## Hybrid-logic chips

# The **best** (or **worst?**) of both worlds

**PROGRAMMABLE-LOGIC DEVICES DELIVER DESIGN, MANUFACTURING, AND AFTER-SALE-SERVICE FLEXIBILITY THAT ASICs CAN'T MATCH, BUT CPLDs AND FPGAs ALSO RUN MORE SLOWLY, BURN MORE POWER, AND COST MORE PER GATE. DOES THE EMERGING ONE-CHIP ASIC/PROGRAMMABLE-LOGIC HYBRID ELIMINATE—OR JUST ACCENTUATE—THE SHORTCOMINGS OF BOTH PRECURSORS?**



*At a glance* ..... 140  
*Where's everyone else?*..... 144  
*Variations on a theme*..... 146  
*For more information* ..... 150

**C**PLDs AND FPGAs, IN THEIR WIDESPREAD INDUSTRY adoption, are achieving success mostly at the expense of gate-array—and, to a smaller extent, standard-cell—ASICs. To the degree that flip-flop toggle rates and fast-route delays approximate the results you see on real designs, programmable logic is getting faster. If you can use the embedded memory that programmable logic provides,

the logic is getting bigger and, correspondingly, cheaper on a per-gate basis, too (**Figure 1** and **Reference 1**).

Programmable-logic trends are particularly welcome news if your volumes are too low to attract an ASIC vendor's or foundry's interest or if you can't suffi-

ciently amortize mask and NRE costs, which are exponentially increasing with each process generation. ASICs also require you to develop detailed test-vector suites for per-die functional and timing verification, something that CPLD and FPGA manufacturers themselves handle.

You also can't ignore the time-to-market angle: A multiweek or multimonth delay occurs from ASIC-design completion to getting the first samples in your hand. Compare that delay with your ability to buy a PLD from a distributor's or a supplier's inventory and begin debugging your design in minutes. Make a few design changes to fix bugs or add features before production begins, and the ASIC NRE costs and long manufacturing cycles become even more problematic.

If your design *will* enter high volume, however, and if you can handle the fabrication delays, standard-cell ASICs represent your lowest cost-per-gate option (Table 1 and Reference 2). ASICs also deliver much higher speeds than FPGAs and CPLDs and at lower power consumption. Just as some flash-memory users wish that they could switch to cheaper mask ROMs, many programmable-logic customers, especially those not using the field-reprogrammable aspect of the technology, look with envy at the ASIC alternative that, for whatever reason, they never get to incorporate in their designs.

### THE MIDDLE PATH

Some vendors, however, think there's a third option: combining programmable logic and gate-array or standard-cell ASICs on the same device. Their approaches differ in the percentages of the die devoted to each type of logic and in what—if anything—they put into the ASIC portion. But the basic motivation is the same: to blend the best attributes of both technologies. How realistic is this goal?

If your objective is solely to reduce cost, the viability of the hybrid-logic approach is unclear and fast-evolving. Take a look, for example, at the ubiquitous PCI core. Just two years ago, programmable-logic vendors were struggling to create a 32-bit target-only core small enough to leave the users with *any* meaningful space for their own custom logic, even using the largest devices in the vendors' arsenals. Today, a 64-bit initiator-plus-target core (minus the FIFO buffers, whose size and number are design-dependent) takes up only an estimated 2.4% of Altera's (www.altera.com) upcoming EP20K1500E logic resources and 1% of those in Xilinx's (www.xilinx.com) upcoming XCV3200E.

### AT A GLANCE

- ▶ Programmable-logic, gate-array, or standard-cell-ASIC approaches alone cannot deliver the perfect combination of high speed, low power, low cost, and flexibility.
- ▶ Combining multiple technologies on a chip results in yet another set of trade-offs for you to consider.
- ▶ The cost savings you achieve with a hybrid-logic chip depend on your production volumes; performance motivations are more compelling.
- ▶ Today's ASIC-plus-programmable-logic-device alternatives incorporate standard-cell and gate-array ASICs and antifuse- and SRAM-based FPGAs. More options are on the way.
- ▶ Silicon availability is insufficient without corresponding software support to bridge the ASIC- and programmable-design methodologies, including hardware/software co-verification.
- ▶ As designs grow more complex and you spend more of your development time simulating, programmable logic's time-to-market advantage over ASICs will diminish.

Granted, these are expensive chips. But take another look at Figure 1's cost-per-gate trends over the last several years and at the predictions for the next few years. Manufacturers also ship these generic logic devices in high volumes and across a range of customers and applications. A hybrid alternative might have a smaller die, depending on the ratio of ASICs to programmable logic, but that die size represents only part of the device cost. Other factors, such as custom testing flows, burgeoning line-item management, and lack of volume-manufacturing efficiencies, play an equivalent or

greater role in increasing the price.

To implement large chunks of logic, such as microprocessors, using ASIC gates is still a better approach (Reference 3). When large CPLDs and FPGAs will become feasible homes for these types of cores depends to a degree on the amount of on-chip memory the cores can use. Xilinx, for example, claims that the 0.18- $\mu\text{m}$  XCV2000E Virtex-E FPGA has twice as many "system" gates as the 0.22- $\mu\text{m}$  XCV1000 Virtex device. However, the logic-cell count grows only 50% from Virtex to Virtex-E; the remainder of the incremental gate count comes mostly from a doubling of the digital delay-locked loops (DLLs) and a fivefold increase in the amount of on-chip Block SelectRAM.

Conceptually, embedded memory could find use as a CPU core's Level 1 and Level 2 cache, but optimal cache design would require larger arrays than those appropriate for other logic circuits. Cache performance might also suffer, and the amount of available generic logic gates would decrease unless the vendors included dedicated cache controller logic. This on-chip logic is analogous to the DLLs and PLLs, dual-port-RAM, content-addressable-memory, FIFO-buffer, and other dedicated-function resources that programmable-logic vendors provide. Unlike those circuits, however, cache logic has more limited applicability. The more customized a chip becomes, the shorter the application and customer list and the smaller the potential volume.

### THE NEED FOR SPEED

Performance is a more compelling reason to use hybrid chips. If, for example, your CPU core needs to run at speeds equivalent to a standard product, ASICs are your only options. The programmable-logic vendors all claim that their 64-bit PCI cores hit 66-MHz speeds, and the cores probably do in some cases, espe-

**TABLE 1—LOGIC TECHNOLOGIES VERSUS STANDARD-CELL ASICS**

	Laser-programmable gate array	Metallized gate array	Programmable logic
Time to first silicon	Days	Weeks	Minutes or hours
Speed	1.1 times slower	1.25 times slower	Three to five times slower
Area*	3.3 times larger	1.6 times larger	10 to 20 times larger

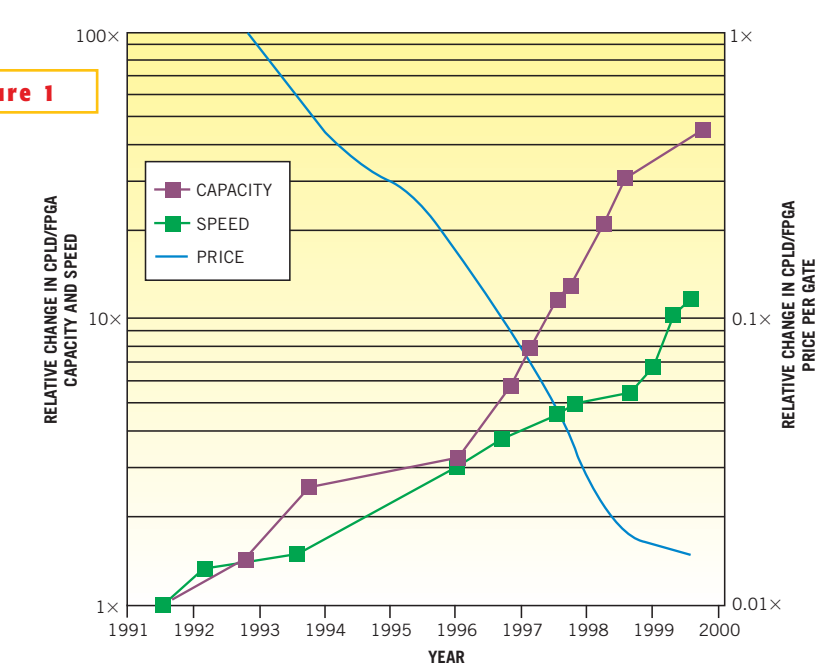
\*For gate count derived from logic cells, not embedded memory.

cially in non-zero-wait-state configurations. However, according to newsgroup postings, when the vendors take this approach at least some of the cores place inflexible placement demands on the front- and back-end tools. As a result, the user-customized portion of the design can't meet gate-count and performance goals (references 4, 5, and 6).

Not too long ago, though, designers raised the same issues with 33-MHz PCI. So, with history as a guide to future trends, robust 66-MHz operation will almost certainly happen in time. I/O-buffer electrical and speed gaps between ASICs and programmable-logic devices are also quickly narrowing; reference DynaChip's (www.dyna.com) DY8000 devices, along with the previously mentioned Altera and Xilinx architectures. However, the interconnect-rich ASIC logic's top speed will always exceed that of routing-matrix-constrained CPLDs and pass-transistor-limited FPGAs and to a smaller extent even that of antifuse FPGAs (Reference 7).

What other high-speed ASIC-based logic cores would make sense in an application that also values programmable logic's flexibility? Although FPGA vendors are beginning to talk with confidence about supporting memory controllers for 133-MHz and double-data-rate (DDR) synchronous DRAM (SDRAM), as well as those for no-latency, DDR, and quad-data-rate synchronous SRAM, they all evade a discussion of Rambus (www.rambus.com) DRAM (RDRAM) controllers (RACs). RACs have high performance requirements, and their external-to-internal bandwidth fan-out would wreak havoc with programmable-logic-routing resources.

RDRAM's single-memory-width channel interface enables smaller system granularity than the minimum density it takes to achieve the same peak bandwidth using wide-bus SDRAM. This factor is especially valuable in embedded designs. Manufacturers of PCs and other



Programmable-logic suppliers are doing their utmost to capture business you might historically have given to an ASIC supplier (courtesy Xilinx Corp).

consumer products, such as Sony's (www.sony.com) Playstation 2, are ramping up volumes. This ramp-up will help to make today's RDRAM incremental cost less than that of asynchronous DRAMs and SDRAMs. Don't be too surprised, therefore, if a hybrid-chip vendor in the not-too-distant future offers a device combining programmable logic and a RAC.

Many high-speed networking- and telecommunications-interface protocols also require speeds that programmable logic can't currently handle. One other requirement that many of them share, leading to another opportunity for ASIC-plus-programmable-logic hybrids, is mixed-signal integration. Analog capability, with the exception of PLLs, is notably absent from today's CPLDs and FPGAs. Specialized clock-recovery circuits and wireless baseband processing are examples of analog functions that

standard-cell and custom-ASIC processes uniquely encompass.

Power consumption remains a nagging concern in large CPLD- and FPGA-based designs (Reference 8). In many cases, the vendors, through packaging advances and lithography shrinks that reduce operating voltages, have been able to stay (barely) below the operating-temperature threshold beyond which device function and long-term reliability suffer. I/O-buffer flexibility lets you interface these parts to legacy devices with higher supply voltages and input levels and wider output swings. However, extreme operating frequencies, coupled with high toggle percentages for internal logic nodes, demand heat sinks, forced air cooling, or both. By putting at least some fast logic into ASICs, you lower the overall device power consumption, though you still need to watch out for die "hot spots."

TABLE 2—FIPSOC CURRENT AND FUTURE PRODUCT FAMILIES

Digital-macrocell-array size	No. of configurable analog blocks	No. of programmable I/Os	Configuration memory (kbytes)	Package
8×8	One	56	4	120-pin QFP
8×12	One	56	6	160-pin QFP
16×12	Two	96	12	208-pin SQ
16×16	Two	112	16	208-pin SQ

THE PIONEERS

PCI's popularity, not just in PCs but also in a variety of embedded designs, makes PCI-based products a natural first step for vendors taking a stab at the hybrid-chip business (see sidebar "Where's everyone else?"). In developing the OR-3TP12, Lucent Technologies replaced 72

of the OR3T55's 324 logic cells with an 85,000-gate gate-array-ASIC-based, 64-bit, 66-MHz PCI initiator/target core. The OR3TP12 retains 252 programmable-logic cells, which Lucent estimates can hold 30,000 to 60,000 gates of custom logic.

The 0.35- $\mu$ m-fabricated OR3TP12 includes two 64 $\times$ 32-bit initiator and two

64 $\times$ 16-bit target FIFO memories. Four 16- or two 32-bit signal paths create the bridge between ASIC-core and FPGA partitions. Lucent has just started shipping the OR3LP26B, an enhanced version of the OR3TP12 that migrates to a 0.25- $\mu$ m process for faster FPGA-logic performance and doubles the amount of interconnects between the PCI core and

FPGA subarray. Lucent also doubled the amount of FPGA logic on the OR3LP26B and better integrated the clock-distribution networks running between the PCI core and programmable logic.

Lucent and Atmel alone among silicon vendors have both ASIC and programmable-logic divisions under the same corporate "roof." Using its hybrid-logic

## WHERE'S EVERYONE ELSE?

As you've probably figured out by now, combining ASIC and programmable logic on a chip involves few technical limitations. True, an antifuse-, EE-PROM- or flash-based CPLD or FPGA requires more processing steps than a standard logic process requires. However, an SRAM-based programmable-logic device is essentially 100%-compatible with a standard logic process, especially if the SRAM-based device uses six-transistor memory cells.

The design-tool issues are somewhat more complicated, but not insurmountable. To the programmable-logic software, the ASIC-based circuit appears as a "hard," or location-fixed, core. Analogous to a pinout-frozen chip, fixed intellectual property (IP) does place challenges on the programmable-logic array's routing resources and fitting algorithms; "soft," or location-flexible, IP would reduce these issues. Depending on the functions that the ASIC-based core performs, a design may require additional development-tool support, such as Atmel's co-verification software.

Most vendors in both the ASIC and programmable-logic segments have so far shied away from participating in the hybrid-chip opportunity. CPLD and FPGA companies predictably cite ASIC limitations. They also claim that the combination of large, fast, or large and fast ASIC-based cores alongside a CPLD or FPGA

represents too small a business opportunity for them. That is, if they can't now serve this opportunity with a fully programmable device now, they'll undoubtedly be able to do so soon.

Other factors are probably also at work, however. The programmable-logic companies are set up to ship high volumes of relatively few products and few speeds, packages, and other options. The ASIC business model is radically different, comprehending large numbers of smaller volume custom chips. Shifting from programmable logic toward ASICs requires a transformed corporate organization and infrastructure, especially for companies that rely on foundries instead of owning their own fabrication facilities. Trying to simultaneously service both businesses would create a multiple-personality situation that few companies can successfully pull off. Plus, with the CPLD and FPGA manufacturers doing all they can to obsolete ASICs, offering a combo chip would just send mixed messages to you.

Given this situation, you'd think that the ASIC vendors and foundries would be more likely to migrate toward the hybrid-chip model. However, these companies' business models depend on revenues obtained from numbers of design starts and from redesigns to fix bugs. On-chip programmable logic reduces the likelihood of the occurrence of either of these

scenarios. With this reality in mind, Atmel's and LSI Logic's embedded-programmable-logic-core plans are, depending on your perspective, either visionary or fatal. Or perhaps they simply reflect the vendors' desires to remain relevant to the vast middle ground of potential customers whose volumes are too high to go to production with programmable logic but too low for standard-cell ASICs.

Speaking of visionary, Actel ([www.actel.com](http://www.actel.com)) was the first programmable-logic company to publicly unveil, in late 1996, plans for ASIC-plus-programmable hybrid devices, which the company called system-programmable gate arrays. The chips, however, never saw the light of day because the company stumbled with its first stab at SRAM-based FPGA technology and the corresponding design-tool set.

Motorola ([www.motorola.com](http://www.motorola.com)) also early last year briefly dipped a toe into the ASIC-plus-programmable-logic water with its MPACF250. This chip combined a Pilkington ([www.pilkington.com](http://www.pilkington.com)) MPA1000 FPGA partition with an ASIC-based ColdFire CPU core and peripheral set. Several weeks after announcing the chip, Motorola acquired Pilkington and canceled the chip and, shortly thereafter, the entire programmable-logic line. The MPACF250 was a great idea, and the CPU core Motorola integrated benefited from more

widespread usage than Atmel's AVR, but the implementation didn't survive budget and headcount cutbacks and redirections.


Several companies other than Adaptive Silicon have also expressed interest in licensing their programmable-logic cores, but implementations have yet to emerge. These suppliers include FPGA vendor Gatefield ([www.gatefield.com](http://www.gatefield.com)), perhaps in partnership with its foundries (and, not coincidentally, ASIC suppliers), Infineon Technology ([www.infineon.com](http://www.infineon.com)), and Rohm ([www.rohm.com](http://www.rohm.com)), as well as PLD manufacturer ICT ([www.ictpld.com](http://www.ictpld.com)). Xilinx's ([www.xilinx.com](http://www.xilinx.com)) recent announcement of its acquisition of Philips' ([www.philips.com](http://www.philips.com)) CoolRunner PLD product line also contained an interesting quotation:

"We look forward to working with Xilinx on future joint developments, such as system-on-chip designs," said Arthur van der Poel, Philips Semiconductor's chairman and CEO. Coming close on the heels of Philips acquisition of VLSI Technology ([www.vlsi.com](http://www.vlsi.com)), this statement may indicate that Philips plans to leverage Xilinx's CPLD and FPGA expertise. Then again, product plans change quickly in an emerging market such as this one. Make sure you check out the version of this article on EDN's Web site, where updated versions of this sidebar will contain any late-breaking news.

capability, Lucent plans to focus beyond PCI toward other aspects of the communications-interface market. The company's ORT4622, due to enter production in the first quarter of 2000, is the first example of these plans (Figure 2). This chip contains four 78-MHz, 8-bit, bidirectional interfaces, translating to four full-duplex serial channels of 622-Mbps synchronous-optical (SONET) capability.

In other words, the approach lets you simultaneously stream 2.5 Gbits of data into and out of the ORT4622. The chip

**Figure 2**




The ORT4622 blends SRAM-based FPGA logic with a four-channel, 622-Mbps backplane transceiver implemented in standard-cell ASICs (courtesy Lucent Technologies).

## VARIATIONS ON A THEME

QuickLogic's first hybrid chip, the QL5030, actually didn't contain ASICs at all. The company put a 32-bit, 33-MHz, target-only core into antifuse FPGA gates. Implementing the core in programmable logic may not be as die-efficient as the ASIC-based alternative, but, on the other hand, QuickLogic used an off-the-shelf FPGA, not a custom chip, to construct the QL5030. This approach has some merit for a nonvolatile CPLD or FPGA, and it wouldn't be surprising if QuickLogic or some other company again takes this route.

What happens if you require faster time to first silicon than a standard-cell ASIC can give you but portions of the design still require standard-cell speed, cost, and power, and the remainder has more stringent requirements than programmable logic can deliver? Both Texas Instruments (www.ti.com) and VLSI Technology (www.vlsi.com) sell ASICs that blend standard-cell and quicker-turnaround gate-array technologies. As with a programmable-logic device, most of a gate array contains generic logic, customer-customized via the last few metal-interconnect layers at the end of fabrication.

Along the same lines, Lucent Technologies complements its FPGA-plus-standard-cell chips by offering embedded arrays based



**Figure A**

The LPGA logic cell, incorporated as an embedded core in Lucent Technologies' 0.25- $\mu$ m, standard-cell-ASIC technology, presents yet another proprietary coarse-grained structure for synthesis vendors to support (courtesy Chip Express).

on Chip Express' laser-programmable gate arrays (LPGAs) for Lucent's five-layer-metal and higher standard-cell processes. LPGA densities and speed more closely mimic those of standard-cell ASICs than do FPGAs (Table 1). Although you can't program an LPGA on your testbench or in the manufacturing line, using the parts means even less of a time lapse occurs than with metalized gate arrays between design completion and first silicon. However, as your designs grow more complex and you spend

more time simulating them, fab-facility throughput for different logic technologies will decrease in importance.


Like FPGAs, LPGAs use coarse-grained logic cells to minimize the amount of required programmable interconnect (Figure A). LPGA-derived logic is intrinsically slower than that in standard-cell ASICs, but the vendors claim that oversized routing-drive transistors make up much of the difference. Chip Express and Lucent's goal is that the design of the LPGA portion

of the hybrid device be no different to you from that of the ASIC partition. However, the companies admit that, as with FPGAs and CPLDs, work still remains in convincing the synthesis-tool vendors to directly support not only an ASIC's two-input NAND cell, but also the vendors' proprietary logic structures. Until then, Chip Express and Lucent will extract as much efficiency as they can using their back-end tool sets.

also integrates analog-centric circuits, such as clock and data recovery, framing functions, and the pointer interpreter, along with corresponding FIFO buffers. The same device can support numerous back-end interfaces through its generic 259-pin FPGA partition, which the company claims can hold a 60,000- to 120,000-gate circuit. Lucent uses standard-cell ASICs to construct the SONET core, and the company will probably also use standard-cell processes for all future ASIC-plus-FPGA devices. The company plans to ship three more communications-centric hybrid chips in the first half of 2000 and is investigating DSP-core integration.

QuickLogic is also tackling PCI with the QL5064 and QL5032 (Figure 3). The 5064 integrates a 64-bit initiator/target core (which you can also use as 32 bits), running at frequencies as high as 75 MHz, whereas the second chip incorporates a 32-bit-only initiator/target core and operates as fast as 33 MHz. QuickLogic's 75-MHz option is not strictly PCI-compliant but may be acceptable for closed-box systems that could use the performance boost. For the same reason, Lucent offers a 50-MHz option for its OR3TP12. The QL5064 contains 12.7 kbits of on-chip RAM and approximately 30,000 gates of user-configurable antifuse-based logic, and the device comes in a 456-pin PBGA package.

QuickLogic is evaluating a high-speed, 32-bit-only initiator/target device, which the company would derive from the QL5064 but could offer in a package with fewer pins. The QL5032 comes in 208-pin PQFP and 256-bump PBGA packages, contains 16 kbits of on-chip RAM, and offers 145,000 gates of programmable logic, again based on the vendor's pASIC3 FPGA family. QuickLogic has just released to production the QL5030-based QL5130 with 57 kbits of RAM and 138,000 FPGA gates, along with the QL5032-derived QL5232, with 25.3 kbits of RAM and 122,000 FPGA gates (see



Antifuse-based logic and an interconnect configuration make up the user-programmable portion of the QL5064 (courtesy QuickLogic Corp).

sidebar "Variations on a theme"). The company touts the robust number of interconnections between the PCI core and the user-programmable logic that the antifuse technology delivers. QuickLogic's plans include products that cover the range of PCI functions, including new standards, such as PCI-X, as well as high-speed serial-bus and communications applications.

#### MAKE MINE A MICRO

A number of companies have combined or are planning to combine ASIC-based microprocessor or DSP cores with programmable-logic arrays (Reference 9). Vendors include Chameleon Systems ([www.chameleonsystems.com](http://www.chameleonsystems.com)), Malleable Technologies ([www.malleable.com](http://www.malleable.com)), Morphics Technology ([www.morphics.com](http://www.morphics.com)), QuickSilver Technology ([www.quicksilverttech.com](http://www.quicksilverttech.com)), and Triscend ([www.triscend.com](http://www.triscend.com)). The on-chip programmable logic could serve multiple purposes: gobbling otherwise-discrete "glue logic" on the system board, acting as an adaptive coprocessor, or implementing a customer-specific set of peripheral functions.

Like Triscend, Semiconductor Design Solutions bases its field-programmable

systems on chips (FIP-SOCs) on 8051 cores (Table 2). Both the analog and the digital peripheral subsystems are user-programmable. Each digital macro cell includes four four-input look-up tables and four flip-flops. Configurable analog blocks enable custom functions, such as differential amplification, comparison, and data conversion. The company has tailored these blocks for use as signal-conditioning front ends instead of for use in more general-purpose analog applications. A triple-array configuration-memory plane enables you to download a new bit stream while the FIPSOC is operating and to rapidly perform partial or complete reconfiguration of the device.

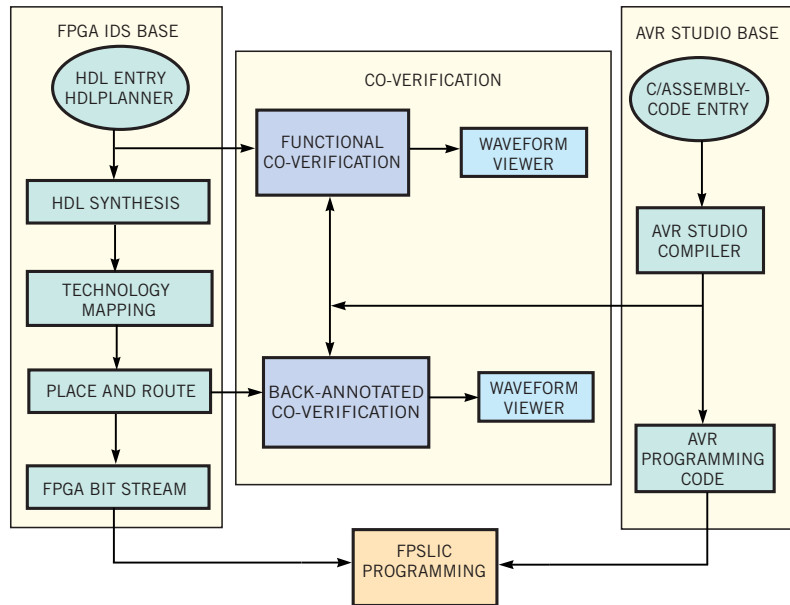
Atmel becomes the latest entrant in this increasingly

crowded race with its first 2.7V field-programmable system-level IC (FPSLIC) family. All devices include a greater-than-30-MIPS AVR RISC microcontroller core to which Atmel has added a dedicated hardware multiply unit, 32 kbytes of user-partitionable code and data RAM, a serial-peripheral interface on high-pin-count packages, two UARTs, two 8-bit timer/counters, a 16-bit timer/counter, a watchdog timer, and a real-time clock. The programmable-logic array, having 576 to 2304 logic cells (10,000 to 40,000 estimated gates) and 4.6 to 18.4 kbits of dedicated RAM, derives from Atmel's AT40K FPGAs. The communication channel between the CPU and the FPGA includes as many as four FPGA-generated interrupts.

Design-tool support is at least as important as silicon capability in turning the systems-on-chips hype into reality. Atmel's System Designer tool set combines the vendor's IDS FPGA synthesis and place-and-route software with the Studio AVR compiler. System Designer also bundles a Mentor Graphics ([www.mentor.com](http://www.mentor.com)) functional co-verification tool, along with Model Technologies' ([www.model.com](http://www.model.com)) simulator and waveform viewer (Figure 4). The

AT40K FPGA architecture is a natural fit for adaptive hardware acceleration, and System Designer lets you quickly evaluate multiple hardware-versus-software partitioning scenarios without an in-circuit emulator or a prototype board.

**Figure 4**



The ideal hybrid-chip-development suite combines hardware and software design with cosimulation capability (courtesy Atmel Corp).

**A MORE GENERIC APPROACH**

Selecting an application-specific standard product, such as one of Atmel's, Lucent's, or QuickLogic's devices, gives you different benefits from those you derive with the alternative approach of designing your own custom chip. You don't have to worry about NRE charges, minimum volume requirements, expensive design tools, time to market, and first-time-functional risk. You also avoid a plethora of the potential issues involving purchasing, modifying, and integrating intellectual property (IP) with the rest of your design. However, an off-the-shelf chip sometimes just doesn't work for you; it may lack some features your design requires or contain capabilities you don't need and don't want to pay for.

ASIC vendor LSI Logic is betting that many of you fit this description, and, in response, the company has partnered with programmable-logic IP provider

Adaptive Silicon (Reference 10). LSI Logic offers programmable-logic embedded-core capability beginning on its

0.18- $\mu$ m G12 process. The Adaptive Silicon logic cell's heritage extends back to National Semiconductor's CLay and NAPA architectures (www.national.com/appinfo/milaero/napa1000), as well as to Concurrent Logic's technology, which became the basis for Atmel's AT6000 line.

LSI Logic and Adaptive Silicon are currently alpha-testing the silicon and design software with a short list of customers, and the two companies hope to broaden the program in the first half of next year. Unlike the more general-purpose multiplexer- or look-up-table-based logic cells, LSI Logic chose to go with Adaptive Silicon's half-adder cell, which works well with arithmetic-centric circuits. More general-purpose designs, which also require robust routing resources, probably won't efficiently map to the Adaptive Silicon structure, however.

The two companies see a variety of applications for the embedded cores. These uses include a general-purpose ASIC that could serve multiple system configurations by putting a subset of the logic into programmable gates or a field-repro-

**FOR MORE INFORMATION...**

For more information on products such as those discussed in this article, circle the appropriate numbers on the Information Retrieval Service card or use EDN's InfoAccess service. When you contact any of the following manufacturers directly, please let them know you read about their products in EDN.

**Adaptive Silicon Inc**  
1-408-399-8900  
www.adaptivesilicon.com  
Circle No. 484

**Fujitsu Microelectronics**  
1-408-988-2445  
www.fujitsu.com  
Circle No. 487

**QuickLogic Corp**  
1-408-990-4000  
www.quicklogic.com  
Circle No. 490

**Atmel Corp**  
1-408-441-0311  
www.atmel.com  
Circle No. 485

**LSI Logic Corp**  
1-408-433-8000  
www.lsil.com  
Circle No. 488

**Semiconductor Design Solutions**  
1-415-398-0110  
www.sidsa.com  
Circle No. 491

**Chip Express**  
1-408-988-2445  
www.chipexpress.com  
Circle No. 486

**Lucent Technologies**  
1-610-712-4331  
www.lucent.com  
Circle No. 489

**SUPER CIRCLE NUMBER**

For more information on the products available from all of the vendors listed in this box, circle one number on the reader service card. Circle No. 492

grammable ASIC with which you could respond to evolving standards or fix bugs. Reconfigurable computing is another possible usage model, and even if your production goal is a 100%-ASIC-based chip, you could use a hybrid device to quickly get your product to market. Bug fixing is the least feasible of these scenarios, because, when you complete the ASIC, you don't know where problems will later arise. However, you could place the programmable logic near particularly high-risk ASICs and hope that the on-chip interconnect is sufficient to patch the FPGA array to the rest of the design.

Historically, design-software-development time has slowed the roll-out of new programmable-logic architectures. Thanks to multiyear grants from the Defense Advanced Research Projects Agency of the Department of Defense, however, the Adaptive Silicon engineers have enjoyed no shortage of either time or money at National Semi-

## HISTORICALLY, DESIGN-SOFTWARE-DEVELOPMENT TIME HAS SLOWED THE ROLL-OUT OF NEW PROGRAMMABLE-LOGIC ARCHITECTURES.

conductor, their previous employer. Atmel also promotes FPGA IP capability to its foundry customers, and, depending on the size of the ASIC-based circuits, company officials believe they can cost-effectively integrate programmable-logic cores having 10,000 to 80,000 gates (576 to 4608 logic cells). As with the company's AVR-based hybrid chips, Atmel will use its AT40K FPGA as the foundation for its embedded FPGA-on-ASIC offerings. □

### REFERENCES

1. Dipert, Brian, "Counting

on gate counts? Don't count on it," *EDN*, Aug 3, 1998, pg 52.

2. Dipert, Brian, "Moving beyond programmable logic: if, when, how?" *EDN*, Nov 20, 1997, pg 77.

3. Dipert, Brian, "The end of the road?" *EDN*, May 13, 1999, pg 54.

4. Dipert, Brian, "Getting a handle on HDLs," *EDN*, May 7, 1998, pg 71.

5. Dipert, Brian, "Synthesis shoot-out at the *EDN* corral," *EDN*, Sept 11, 1998, pg 95.

6. Dipert, Brian, "Lies, damn lies, and benchmarks: The race for the truth is on," *EDN*, May 27, 1999, pg 54.

7. Dipert, Brian, "Shattering the programmable-logic speed barrier," *EDN*, May 22, 1997, pg 36.

8. Dipert, Brian, "Programmable logic: Beat the heat on power consumption," *EDN*, Aug 1, 1997, pg 57.

9. Dipert, Brian, "Figuring out reconfigurable logic," *EDN*, Aug 5, 1999, pg 103.

10. Dipert, Brian, "ASICs get 'embed' with FPGAs," *EDN*, Aug 19, 1999, pg 20.

