

Programmable-logic directory

**THE SECOND ANNUAL EDN PLD
DIRECTORY HIGHLIGHTS THE
ARCHITECTURES AVAILABLE FOR YOUR
NEXT DESIGN. FIND OUT WHAT'S NEW,
WHAT'S OBSOLETE, AND WHAT'S
EVOLVED IN PALs, PLDs, AND FPGAs.**

**AND CHECK THIS OUT:
WE'VE POSTED COMPREHENSIVE
TABLES OF DEVICES AND FEATURES
IN THE WEB VERSION OF THIS
ARTICLE AT
www.ednmag.com**

THE UMBRELLA TERM "logic devices" subdivides into several categories: discrete logic, simple and complex PLDs, FPGAs, and standard- and custom-cell ASICs. FPGAs, SPLDs/PALs, and CPLDs are all programmable-logic devices, although their internal architecture implementations differ.

Programmable-logic devices are the fastest growing segment of the logic-device family, for two fundamental reasons. For one thing, their ever-increasing per-device logic-gate count "gathers up" functions that might otherwise spread over a number of discrete-logic and memory chips, improving end-system size, power consumption, performance, reliability, and cost. Equally important, you can in a matter of seconds or minutes configure and, in many cases, reconfigure these devices at your workstation or in the system-assembly line. This capability provides powerful flexibility to react to last-minute design changes, to prototype ideas before implementation, and to meet time-to-market deadlines driven by both customer need and competitive pressures.

Programmable-logic devices lack the long lead-times, up-front NRE charges, minimum-order quantities, and inventory complexity of ASICs. As per-gate cost decreases and the number of gates per component increases, programmable-logic devices are making significant inroads into gate-array-ASIC territory. System designers and manufacturers are only beginning to explore and exploit in-system reprogrammability, either to correct errors and upgrade functions once the end system is in users' hands or to use a fixed number of logic gates to implement multiple functions. This technique is known as "reconfigurable computing."

The programmable-logic industry is relatively young and highly varied. Just as companies use programmable logic's flexibility to differentiate themselves, a number of semiconductor vendors have developed unique PLDs and FPGAs to address an intersection of performance, power, integration, and cost targets. This diversity is perhaps the most complex challenge facing you, because, in many cases, you must deeply understand each programmable-logic architecture before you can select one that meets your needs. The market leaders are increasingly driving de-facto industry standardization, thus simplifying the selection task.

Highly complex programmable-logic architectures rely extensively on design-automation software to produce optimum results for end-system parameters. Pri-

oritizing these parameters depends on the application. Often, for example, designs targeting low power, high performance, or minimal gate count differ significantly from each other. Ideal design-automation software:

- isolates you from the internal device-architecture details,
- enables you to prioritize your design goals and optimizes the software's operation based on the order you choose,
- efficiently uses silicon resources,
- requires little to no manual intervention,
- quickly compiles and recompiles a design, and
- minimizes or eliminates timing and pinout changes between compilations.

The technical superiority of a programmable-logic vendor's silicon products and the comprehensiveness of the vendor's documentation are not the only determinants of the vendor's success. Equally important are the depth and breadth of the company's internally developed and third-party software-tool support.

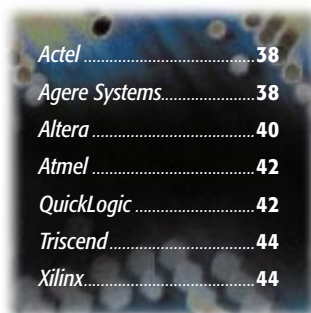
Burgeoning amounts of on-chip RAM and single-die ASIC/programmable-logic hybrid devices, along with predictable Moore's Law integration trends, are contributing to the explosion of effective gate counts. These factors are finally making a reality of the long-held vision of systems on chips. To exploit silicon capability in a time frame that still meets time-to-market requirements, many designers are turning from traditional low-level state-machine and schematic-entry synthesis to high-level languages, such as VHDL and Verilog, and even to traditional software languages, such as C. These new

language approaches provide the additional benefit of enabling design reuse. Yet, just as with high-level versus assembly-language software development, high-level logic design decreases development time but produces lower performance and less efficient gate usage.

Another technique that has become more popular over the last few years involves leveraging the already-completed designs, or IP (intellectual property), of another company instead of designing your own circuits. The convergence of accelerating silicon gate count, increasing system functions and standardization, and faster time-to-market requirements is driving this approach. Perhaps the biggest IP hurdles still to overcome are legal rather than technical, although robust test and verification suites and core interoperability among vendors and among silicon architectures are important.

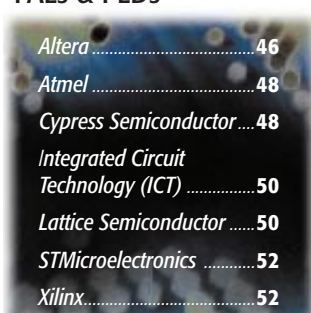
Although "pure" programmable-logic devices dominate industry shipments, this fast-moving product category is evolving in several intriguing directions. Vendors historically known as programmable-logic suppliers are adding ASIC-plus-programmable-logic hybrid devices to their portfolios in an attempt to optimally address the complicated task of balancing power, performance, and cost. ASIC vendors and foundries, too, are developing hybrid-device capabilities, as are standard-product suppliers, such as network-processor manufacturers. And IP developers are now offering not only "soft" logic cores that implement specific functions in programmable logic but also "hard" embedded programmable-logic cores that can implement any function you desire.

FPGAs



<i>Actel</i>	38
<i>Agere Systems</i>	38
<i>Altera</i>	40
<i>Atmel</i>	42
<i>QuickLogic</i>	42
<i>Triscend</i>	44
<i>Xilinx</i>	44

PALs & PLDs



<i>Altera</i>	46
<i>Atmel</i>	48
<i>Cypress Semiconductor</i>	48
<i>Integrated Circuit Technology (ICT)</i>	50
<i>Lattice Semiconductor</i>	50
<i>STMicroelectronics</i>	52
<i>Xilinx</i>	52

Embedded programmable-logic cores



<i>Actel</i>	54
<i>Adaptive Silicon</i>	54
<i>Agere Systems</i>	56
<i>Atmel</i>	56
<i>Integrated Circuit Technology (ICT)</i>	56

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to www.ednmag.com. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

Actel

1-408-739-1010
www.actel.com
 Enter No. 322

Adaptive Silicon

1-408-399-8900
www.adaptivesilicon.com
 Enter No. 323

Agere Systems

1-610-712-6011
www.agere.com
 Enter No. 324

Altera

1-408-544-7000
www.altera.com
 Enter No. 325

Atmel

1-408-441-0311
www.atmel.com
 Enter No. 326

Cypress Semiconductor

1-408-943-2600
www.cypress.com
 Enter No. 327

Integrated Circuit Technology (ICT)

1-408-434-0678
www.ictpd.com
 Enter No. 328

Lattice Semiconductor

1-503-268-8000
www.latticesemi.com
 Enter No. 329

QuickLogic

1-408-990-4000
www.quicklogic.com
 Enter No. 330

STMicroelectronics

1-781-861-2650
www.st.com
 Enter No. 331

Triscend

1-650-968-8668
www.triscend.com
 Enter No. 332

Xilinx

1-408-559-7778
www.xilinx.com
 Enter No. 333

OTHER COMPANIES MENTIONED IN THIS ARTICLE

ARM

www.arm.com

Chartered Semiconductor

www.csminc.com

Chip Express

www.chipexpress.com

Clear Logic

www.clear-logic.com

IBM

www.ibm.com

LSI Logic

www.lsillogic.com

MIPS Technologies

www.mips.com

Motorola

www.mot.com

National Semiconductor

www.national.com

Philips

www.philips.com

Synopsys

www.philips.com

TSMC (Taiwan Semiconductor Manufacturing Co)

www.tsmc.com

UMC

www.umc.com

SUPER INFO NUMBER

For more information on the products available from all of the vendors listed in this box, go to www.ednmag.com, click on the Reader Service link, and enter no. 334.

ACTEL

You probably know Actel as an antifuse-based-FPGA supplier, but the company is broadening its focus to include other programmable-logic technologies, as well as IP. Four main device families form the foundation of today's Actel product line: the 3.3V SX, the closely related 2.5V SX-A and eX, and the 5V MX. Actel also offers the flash-memory-based ProASIC FPGA architecture and an SRAM-based embedded FPGA called VariCore (see "Embedded programmable-logic cores"). And, by the end of 2001, the company hopes to begin sampling its first flexible-interface BridgeFPGA devices.

Antifuse technology delivers low-impedance—therefore, low-power and high-speed—signal interconnection and more robust immunity to high-radiation operating environments than other configuration technologies provide. Actel devices are available in both standard commercial- and extended-temperature options and in screened high-reliability, radiation-tolerant, and radiation-hardened versions.

Actel's parts are also nonvolatile. Unlike SRAM-based FPGAs, they require no separate memory to store their configuration data, and their functions are available immediately upon system power-up. The single-chip nature of antifuse FPGAs also makes them nearly impossible to reverse-engineer or clone. This characteristic becomes increasingly important as decreasing costs and higher capacities

AT A GLANCE

- ▶ Actel's antifuse technology delivers numerous benefits at the expense of limited flexibility.
- ▶ The company's second-generation antifuse structure doesn't consume substrate area.
- ▶ Flash memory builds on antifuse strengths and, according to the company, is finally ramping into production.

result in the parts' use in high-volume consumer products.

The company's first-generation PLICE (programmable low-impedance circuit element) antifuse-technology approach employs a metal-to-metal interconnect structure comprising polysilicon and a diffused N+ region, separated by a high-impedance oxide-nitride-oxide barrier. Applying high voltage during programming ruptures this barrier. The PLICE antifuse structures,

which reside on the same base layer as active circuit elements, consume die area that could otherwise find use in constructing additional logic blocks, embedded memory arrays, and other circuits.

As a result, beginning with the SX family, Actel has migrated to a second-generation-technology approach that locates antifuses directly between metal layers, above the logic. Whereas MX-series FPGAs use a relatively generic multiplexer-plus-register logic block, SX, SX-A, and eX devices employ a sea-of-modules ratio of two logic structures: C-cells and R-cells. C-cells contain a dual-level, two-input multiplexer structure plus input-inversion capability, and the company claims that the C-cells can implement more than 4000 functions of five inputs. R-cells contain multiple-function flip-flops with numerous signal-input, clocking, reset, and clear options.

The high voltage necessary to configure an antifuse FPGA usually means that you program it before installing it on

your system board, and, because antifuse creation is irreversible, in-system reconfiguration is impossible. In response to customers' requests for a more flexible FPGA technology (in the lab, the manufacturing line, and the field) that retains antifuse's fundamental benefits, Actel first partnered with and then acquired Gatefield Corp and its ProASIC line of flash-memory-based FPGAs.

In addition to being onboard-programmable and -reprogrammable, ProASIC devices use an extremely fine-grained three-input, one-output logic cell. Actel claims that this logic cell not only provides you with an intuitive ASIC-prototyping vehicle, but also delivers a smooth learning curve for budding FPGA designers who are already experienced with ASICs and their design tools. ProASIC also brings embedded-SRAM and FIFO capabilities to Actel's arsenal—features missing from all but a few of the company's antifuse FPGAs.

Actel's Libero and Designer tool sets support antifuse FPGAs, and several variants supply multiple combinations of device and family, design entry, synthesis, simulation, placement and routing, and programming support. The Silicon Explorer II verification and logic-analysis tool enables you to observe and analyze internal device nodes without, in most cases, iterating your design. The Silicon Sculptor single-site and multisite device programmer supports programming of FPGA prototypes right at your PC. For ProASIC FPGAs, the vendor-supplied tools include ASICmaster (for placement and routing) and Memorymaster (for embedded-memory-function generation). IP comes both from Actel and from the company's strategic partners.

AGERE SYSTEMS

Agere Systems, formerly Lucent Technologies (and before that, AT&T Microelectronics), first entered the FPGA business as an alternative supplier of Xilinx's XC3000 FPGAs, which Agere referred to as its ATT3000 series. From those humble beginnings, the company has made significant progress in establishing its own identity as a programmable-logic vendor and taking advantage of its internal FPGA and ASIC capabilities. Given its

AT A GLANCE

- ▶ FPGAs focus on networking needs.
- ▶ Logic-structure innovations require synthesis support or core-generator intercession.
- ▶ Hybrid devices handle high-speed interfaces, and the latest architecture hints at CPU-integration plans.

parent company's telecommunications heritage and the high percentage of FP-

GAs sold into the digital-communications market, Agere's product-line and technology focus in this area isn't a big surprise.

Within each product family (Orca Series 2, 3, and 4), multiple product generations exist. These subdivisions reflect different manufacturing-process technologies and, sometimes, corresponding differences in core and I/O voltages. Series 2 devices group four four-input LUTs and four registers into each PFU (pro-

grammable function unit). As in Xilinx's, and unlike Altera's architecture, each LUT grouping can alternatively serve as a synchronous or an asynchronous, single- or dual-port, RAM or ROM block. Each PFU also contains eight tristate buffers for implementing internal bus structures.

Enhancements with Series 3 include the SLIC (supplemental-logic-and-interconnect cell), an upgraded version of the tristate-buffer structure that now also supports an as-much-as-10-bit decoder and PAL-like AND-OR-INVERT logic. A built-in microprocessor interface enables parallel programming and configuration read-back and provides a glueless connection to i960 and PowerPC CPUs. The PCM (programmable clock manager), a PLL variant, enables adjustment of input clock phase and duty cycle. The size of each Series 3 PFU, which now includes eight four-input LUTs and nine registers, more than doubles that of Series 2 PFUs. Lucent supplies 5 (3C), 3.3 (3T), and 2.5V (3L) variants of the architecture.

Agere's Series 4 devices focus attention on signal routing. The company based this decision upon recognition that delays in this area—not in logic—are increasingly defining the upper limit of design performance. Agere fabricates its initial Series 4 devices with a 0.16-mi-

cron-drawn (0.13-micron-effective) process that provides abundant metal layers. The company employs these layers—along with other uses—to create a dedicated clock-distribution network throughout the chip. Active repeaters prevent signal-quality and performance degradation across long routes and multiple pass-gate interconnect elements. Agere's ASIC division also offers the Series 4 programmable-logic structure as an embedded core (see "Embedded programmable-logic cores").

Regardless of its logic density, each Series 4 device includes six general-purpose and two application-specific PLLs. Lucent supplements its FPGAs' previous-generation LUT-derived embedded-memory capability with dedicated 512×18-bit, quad-port (two read, two write) discrete RAM blocks, including built-in write-port arbitration, a FIFO, a multiplier, and CAM logic. Series 4 I/O buffers' optional LVDS terminating resistors are on-chip. Like previous-generation Orca architectures beginning with Series 2, Series 4 devices are partially reconfigurable. Lucent revamped the logic-cell structure with modular, hybrid design in mind. An on-chip, ARM-derived, multimaster peripheral bus both simplifies the interconnection of multiple log-

ic blocks and gives a glimpse of the company's probable future CPU-integration plans.

In 1996, both Actel and Lucent announced their intentions to produce hybrid ASIC-plus-FPGA devices. Lucent's promise turned into reality in May 1998, when the company introduced its first FPSC (field-programmable system chip). The Orca Series 3-based OR3TP12 combines an 18×18 PLC (programmable-logic-cell) array and an ASIC-housed, 64-bit, 66-MHz PCI core. This core interconnects via two 64×32-bit master FIFOs and two 64×16-bit target FIFOs. The follow-on OR3LP26B doubles both the amount of on-chip FPGA logic and the ASIC-to-FPGA interconnect bandwidth, and the ORT4622 replaces the PCI core with a four-channel, 622-Mbps, full-duplex synchronous interface with built-in CDR. The first Orca Series 4-based hybrid chips, the ORT8850 series, are derivatives of the ORT4622 that include an eight-channel, 850-Mbps CDR macro. The ORT82G5 is a 1.25-, 2.5-, or 3.125-Gbps backplane-interface FPSC, and the ORLI10G is a 10-Gbps line-interface device. Lucent's Orca Foundry back-end software handles placement and routing for all of the company's FPGAs and ASIC-plus-FPGA devices.

ALTERA

For legal and marketing reasons, you'll never find the acronym "FPGA" in any of Altera's literature, even though the rest of the world knows the chips as FPGAs. Altera prefers to call its parts "LUT-based programmable-logic devices." A long-line, routing-dominated FPGA architecture provides timing predictability; in addition, Altera's routing approach enables on-chip redundancy, particularly during the early stages of process and device production, to improve yields.

From the Flex 8000 and follow-on Flex 10K architectural foundations, Altera has taken its FPGA product line in several directions, all based on a common, essentially unchanged LAB (logic-array-block) structure. Flex 6000, the Flex 8000

AT A GLANCE

- ▶ Aggressive migrations to smaller lithography processes have enabled lower prices, higher gate counts, reduced power, and higher speeds.
- ▶ Copper interconnect strives to keep internal logic delays from falling behind devices' I/O-buffer performance capabilities.
- ▶ Hybrid chips with industry-standard ASIC-housed CPU cores, along with a soft CPU core, are finally here.

successor, targets moderate-density and -performance designs that require no on-chip memory or esoteric packaging. Altera eliminated the Flex 10K PLLs and decreased the amount of global signal routing on Flex 6000 from that on the Flex 10K. Instead, the Flex 6000 substitutes dedicated local LAB-to-LAB and LAB-to-I/O-buffer interconnect. The Acex 1K family is a Flex 10KE variant built on a smaller lithography process that is optimized for low cost. Like Flex 6000, Acex 1K comes in smaller gate counts and with packaging options that are more restricted and less expensive than its Apex 20K and other Altera big brothers.

Apex 20K includes as many as four on-chip PLLs. The device family also pro-

vides larger and more numerous EABs (embedded array blocks) and flexible I/O buffers that support numerous protocols and electrical standards. As with Xilinx's chips, Altera's devices' EABs are the key factors behind the vendor's claims of exponentially growing gate counts. If you can use all that memory, great. If not, you won't come close to squeezing into one chip the design sizes that Altera's marketing indicates are possible. Apex 20KE enables you to use the EABs for implementing not only single- and dual-port RAM and FIFOs, but also very small CAMs. Apex 20KC migrates the Apex 20KE architecture to a process employing low-impedance copper-interconnect lines for all metal layers.

Development of Altera's CPU-inclusive hybrid chips took longer than expected, but Altera is now sampling its ARM-based Excalibur XA devices, built on an Apex 20K foundation, with MIPS-based XM parts soon to follow. Combine Apex 20KE with eight to 18 channels of

ASIC-housed CDR circuitry and slightly modify the logic block structure, and you have the Mercury family, supporting data-throughput speeds of 125 Mbps to 1.25 Gbps. And double up the RAM-to-ESB (embedded-system- block) ratio in Apex 20KC; increase the number of PLLs; and speed up, increase the number of, and make more flexible the I/O

buffers, and you have the latest Altera architecture, Apex II.

Altera has also developed an optimized 8- and 16-bit RISC processor it calls Nios, which the company aims to have reside in FPGA logic as a soft core instead of in ASIC gates. Over time, Altera plans to port Nios to all of its FPGA architectures, along with its other internally developed

and partner-developed cores. The same Max1 software you'd use to design with Altera's CPLDs also supports Acex 1K and Flex FPGAs. For Acex 2K and Apex devices, you'll want to fire up Altera's more advanced Quartus II development-tool environment, along with the Signal-Tap logic analyzer.

ATMEL

Atmel's two FPGA architectures, the AT6000 and newer AT40K families, both offer dynamic partial-reprogramming capabilities that the company defined with reconfigurable-computing applications in mind. The two families' logic structures are quite different, though, and the AT40K devices include features that broaden their applicability to general-purpose designs.

AT6000 logic cells consist of several fixed-function logic gates plus a register. In AT40K FPGAs, you find the more common LUT-plus-register combination but with a twist. Instead of a single four-input LUT, Atmel pairs each flip-flop with dual three-input LUTs, a combination that, in some cases (if the design tools take advantage of it), offers greater design flexibility. Also, ahead of the LUTs is a dedicated two-input AND gate. Combined with diagonal routing, this logic tweak enables AT40K FPGAs to better support the matrix-multiplication operation so common in DSP functions.

Unlike with some FPGAs, you can't alternatively use the AT40K LUTs to im-

AT A GLANCE

- ▷ Second-generation FPGAs have multiplication in mind.
- ▷ An 8-bit AVR controller makes system-on-chip designs a reality.
- ▷ A security-minded architecture spin makes design duplication difficult.

plement distributed on-chip memory arrays, FIFOs, or other SRAM-derived functions. However, at the point at which four logic-block clusters (each containing 16 logic blocks) intersect, you'll find a 128-bit dedicated RAM array. This approach is an intermediary step between the small, 16-bit, LUT-derived memory arrays and the much larger dedicated SRAM blocks in other vendors' architectures. Atmel has also come up with an interesting packaging twist: The parts are pinout-compatible with some Xilinx XC4000 offerings, conceptually offering you an alternative silicon source for your designs.

AT40K is also the silicon foundation for Atmel's FPSLIC (field-programmable system-level-integration-circuit) devices. These hybrid parts contain both an FPGA array and an ASIC-housed AVR RISC microcontroller, with data and program memory and comprehensive peripherals. Atmel has focused on optimizing the intercommunication link between the CPU and FPGA array, enabling the FPGA array to, for example, implement hardware-accelerated coprocessor functions for

the CPU. An integrated design-tool environment lets you develop software and hardware in parallel, simulating and co-verifying the multiple pieces of your design.

The latest spin on the FPSLIC architecture, AT49S, is a single-chip, dual-die device with a specially designed, non-volatile-memory array. Containing two data buses, one going to the system and the other running only within the package between the memory and AT49K, the Secure FPSLIC's memory also includes a security bit. Once you set this bit, the chip will respond only to a full-chip-erase command. You must set the security bit to initiate memory-to-AT49K communication. The approach doesn't completely eliminate the security threat; although the external bit-stream trace between memory and AT49K is no longer present, stripping back the multidie package lid exposes the internal bus for probing. However, Atmel's tactic is an example of low-cost security that, in many cases, will be good enough to foil the efforts of would-be thieves. Atmel's ASIC group includes the AT40K programmable-logic structure in its IP portfolio, with FPSLIC as the proof of concept (see "Embedded programmable-logic cores").

QUICKLOGIC

Antifuse advocate QuickLogic has, since its earliest pASIC 1 architecture, employed ViaLink, a metal-to-metal antifuse technology, above the logic grid. Only recently has Actel begun to match ViaLink's advantages. QuickLogic's chips have all of the inherent antifuse advantages (see the Actel entry under "FPGAs"), including the availability of extended-temperature and military-screened device variants. Of all the programmable-logic manufacturers, Quick-

Logic has also been the most enthusiastic about embracing the hybrid ASIC-plus-FPGA approach and the most

AT A GLANCE

- ▷ Logic cells promote design flexibility.
- ▷ Embedded standard products integrate RAM and system-interface modules.
- ▷ ASIC-housed DSP, PCI, Fibre Channel and SERDES functions are now available, and MIPS CPUs are on the horizon.

aggressive about rolling out corresponding devices.

The company builds its pASIC 2 and 3 product families on lithographies smaller than those of the original pASIC 1 technology, with correspondingly higher logic counts, lower operating voltages, and higher speeds. The pASIC parts employ a novel logic cell comprising two six-input AND gates, four two-input AND gates, multiple two-to-one multiplexers, and a D flip-flop. The logic cell's numerous inputs allow it to implement—in one

logic level—functions that in other approaches might require multiple performance-sapping logic cells. Multiple logic-cell outputs allow the synthesis and place-and-route software to pack unrelated logic functions into one cell, maximizing silicon use.

QuickLogic's newest FPGA architecture, Eclipse, focuses first on improving the I/O buffers. These buffers now contain input, output, and output-enable registers and support a variety of voltages, including differential standards, on a per-bank basis with eight I/O banks per chip. QuickLogic has also doubled the number of registers in each logic cell and added a multiplexer and now provides as many as six outputs. Eclipse also has four PLLs, a beefed-up clock- and control-signal network, and multiple embedded dual-port-RAM blocks.

The pASIC FPGAs are the foundation of the company's QuickRAM, QuickPCI, and QuickFC devices, the first few in a series of the company's ESPs (embedded standard products). QuickRAM chips

add embedded-RAM blocks, and QuickPCI parts embed PCI cores alongside various sizes of user-programmable logic gates. The company offers numerous PCI-core flavors: 32- and 64-bit; 33-, 66-, and 75-MHz; and master and slave. A high-performance, beefy set of FIFOs connects the PCI core to your design's logic. QuickFC employs ASIC logic to construct a Fibre Channel encoder/decoder with data rates as fast as 2.5 Gbps and a 32-bit synchronous-FIFO system interface.

QuickLogic turned to the Eclipse architecture to implement its next ESP families. QuickDSP uses ASIC gates to implement dynamically reprogrammable dedicated-arithmetic circuits that the company calls ECUs (embedded computational units). Each ECU can implement several single-pass asynchronous and registered functions (8×8-bit multiply, 16-bit add, or accumulate with carry); multiple passes through the ECU support the common multiply-accumulate function. QuickSD devices integrate

as many as eight SERDES data channels and two clock channels, all bus-LVDS for high-current and long-signal-drive capability. And, like Altera, QuickLogic has licensed the MIPS32 4Kc core (with product availability slated for this year) and an option for the MIPS64 5Kc.

Design-software support comes from QuickLogic's Quick-Works for PCs and QuickTools for workstations. QuickLogic was one of the first programmable-logic companies to embrace synthesis-based design and bundle compilers with its tool set. The online WebAsic program enables you to upload your design's bit-stream file and receive free samples from QuickLogic; North American customers can receive samples in as little as 24 hours. The company also supplies the QuickPro desktop device programmer, and various PCI-development boards. And, for QuickDSP devices, the QuickFilter tool lets you create and analyze digital-filter designs. The tool then generates coefficients, computes magnitude and phase responses, and creates stimulus test files.

TRISCEND

Triscend originally intended the programmable-logic arrays on its E5 and A7 CSOCs (configurable systems on chip) to implement only customer-specific microprocessor peripherals. However, recent revisions of both Triscend's FastChip tools and its partners' design software enable the arrays' use as general-purpose programmable logic. Highlights of the E5 family include a performance-accelerated 8051 microcontroller core and as much as 64 kbytes of on-chip, dedicated system RAM. According to the company, 3200 program-

mable-logic cells translate to as many as 40,000 ASIC gates. E5 devices provide two dedicated DMA channels, and an on-chip breakpoint unit provides debugging capability.

Triscend's long-delayed A7 family brings the CSOC concept to the 32-bit processor world. In this case, the CPU core is the ARM7TDMI, and A7 CSOCs also feature an ASIC-housed memory controller, four-channel DMA controller, JTAG interface, 16-input interrupt controller, dual timer/counters, dual serial ports, and watchdog timer, and other circuits. Based on a joint-development

AT A GLANCE

- ▶ Configurable logic that once targeted only microcontroller peripherals now has broader general-purpose use.
- ▶ ARM- and 8051-based hybrid-chip families address 8- and 32-bit computing needs.
- ▶ SuperH-based devices are on the horizon.

agreement forged with Hitachi in January 2001, SuperH-based CSOCs are now on Triscend's road map.

XILINX

Since Xilinx's founding in 1984, the LUT-plus-register combination has been a consistent element in its devices, but the logic cell and peripherals have evolved over the years. Xilinx also briefly flirted with antifuse-based FPGAs. An architecture for reconfigurable computing has faded into the sunset, but Xilinx has resurrected some portions of the techniques this architecture pioneered in the company's latest product families.

Each XC4000-series CLB (configurable

AT A GLANCE

- ▶ Increased LUT and register integration within each logic block matches the pace of Xilinx's technology progression.
- ▶ Virtex brought embedded discrete-memory arrays to the company's FPGA product line.
- ▶ Foundry-friendly Virtex-EM FPGAs were the first to incorporate copper routing, and PowerPC-based hybrid chips are on the way.

logic block) comprises two four-input LUTs, a three-input LUT, and two registers. I/O buffers also contain multiple registers to ease setup-time restrictions and to boost clock-to-valid-output speeds. Over time, this product family has undergone several process migrations, with corresponding voltage decreases, growth in per-chip logic capacity, and higher performance. Although XC4000 devices contain no large embedded-RAM blocks, you can reconfigure the LUTs for use as distributed-RAM arrays.

ASIC-replacement Spartan and Spartan-XL parts trace their lineage back to the XC4000E. Spartan is a follow-on to Xilinx's first stab at less expensive FPGAs, the XC5200, which, like Altera's Flex 6000 and 8000, didn't support embedded memory. Spartan forgoes the XC4000E's parallel-interface-configuration option and dedicated, on-chip, wired-AND decode and comes only in low-cost plastic packaging. Xilinx offers both 5 and 3.3V versions of the first-generation Spartan architecture.

Compared with the baby steps that preceded it, the XC4000-to-Virtex transition was a giant leap forward in architecture. Xilinx found that front-end synthesis tools rarely take advantage of the three-input LUT. The Virtex logic block eliminates it, switching to a combination of four four-input LUTs and four registers. Xilinx supplements the distributed LUT memory with multiple 4-kbit dual-port discrete-SRAM arrays. For a third level of memory, Virtex provides high-speed, flexible I/O buffers to—among other things—interface to external SRAM and DRAM. First-generation Virtex devices contain four DLLs. The second-generation Virtex-E and Virtex-EM devices don't alter the Virtex logic-cell structure but dramatically boost the

amount of discrete SRAM, leading to an exponentially higher gate count that you can't exploit unless your design relies heavily on embedded memory.

Virtex-E parts also double the number of on-chip DLLs to eight and incorporate a more-than-300-Mbps-per-pin, chip-to-chip, buffered, double-data-rate communication protocol that the company calls SelectLink. Virtex-EM further bloats the memory-to-logic proportion (at four gates per SRAM bit). This device family marks the emergence of copper routing at the upper two metal layers. The copper routing provides power and low-latency clock distribution throughout the chip. Virtex also provides the foundation for the Spartan II family, which migrates to a smaller lithography and focuses on a lower cost packaging subset. Unlike the redesign that marked the XC4000E-to-Spartan transformation, Spartan II retains almost all of Virtex's features.

Xilinx's latest Virtex II family, another copper-enhanced architecture, incorporates discrete-RAM blocks that include parity and are 4.5 times larger than those in Virtex. Virtex II adds support for two new block-RAM operating modes. In moving from XC4000 to Virtex, Xilinx doubled the per-logic-block number of LUTs and registers. The company con-

tinues this trend with Virtex II, which offers eight LUTs and eight registers per CLB. With all of that available metal, Xilinx boosted the amount of long-line, LUT-to-LUT, and CLB-to-CLB routing. The company includes dedicated 18-bit, fast-multiplier logic, which it claims can perform more than 600 billion 8-bit multiply-accumulate operations/sec. Virtex II also adds Digitally Controlled Impedance Technology. This feature provides optional internal termination resistors, whose values automatically match those of the external reference resistors you supply (two resistors for each of the eight I/O blocks on a device).

Xilinx has finally embraced the ASIC-plus-FPGA concept and plans to by the end of 2001 introduce Virtex II Pro parts with integrated IBM PowerPC CPU cores. Design support for Xilinx's FPGAs comes from the Alliance tool set, which interfaces to third-party front-end design software, and from the full-featured Foundation suites. In line with the partially reprogrammable capabilities of Virtex devices, Xilinx has also developed a portfolio of Java-based software products that enable chip programming and reprogramming at the design bench, on the manufacturing line, and in the field.

PALs & PLDs

ALTERA

In the late 1980s, Altera pioneered the concept of the CPLD, a device with numerous logic blocks, each consisting of a PAL- or SPLD-like group of macrocells. These logic blocks interconnected to each other and the outside world via a fully or partially populated switch matrix. Altera's early CPLDs used PROM or EPROM cells as switch-configuration elements, but the company's Max devices have migrated to in-system-programmable EEPROM technology.

Max 7000 is the primary workhorse of today's Altera CPLD-product line, and the architecture has remained essentially unchanged through multiple process evolutions. Smaller semiconductor lithographies often translate not only to lower cost per macrocell but also to higher speed, lower operating voltage and power consumption, and higher macrocell

AT A GLANCE

- ▶ Over the years, Altera has only slightly fine-tuned its Max 7000 flagship CPLD-product line.
- ▶ Hierarchical routing can't keep up with Moore's Law and comes with timing trade-offs.
- ▶ Relaxed and restricted testing, along with plastic packaging, lead to low-cost Altera devices.

counts. Along the way, Altera has fine-tuned the Max 7000 in many ways. Perhaps the most significant is the addition of in-system programming beginning with the 5V S series. The 3.3 and 2.5V variants also include this in-system programming. Altera is also particularly proud of its Multivolt I/O technique, in which the I/O buffers drive output and handle input voltages both lower than

and exceeding the device's core operating voltage for no-glue system interfacing.

In an attempt to radically grow the macrocell capacity of its CPLDs, Altera introduced the Max 9000 family in 1994. This CPLD architecture migrates from a monolithic logic-block-to-logic-block interconnect matrix to a multistage approach. This approach is reminiscent of the company's Flex devices and more generally of any segmented routing FPGA. The advantage of a hierarchical interconnect structure is that, because it is distributed throughout the device, it doesn't exponentially grow with increasing macrocell and, therefore, logic-block count, as a global matrix tends to do.

The disadvantage of a hierarchical approach, though, is that pin-to-pin and logic-block-to-logic-block timing depend on placement and are therefore unpredictable. Performance predictability

and the ability to more easily implement logic circuits with many product terms and comparatively fewer registers have always been key advantages of PALs and PLDs over FPGAs. Search the Web version of this article at www.ednmag.com for **Table 2** and note that the logic density of the Max 7000 family, which has undergone aggressive cost reductions, now comes close to matching that of the largest Max 9000 device and at a lower price per macrocell.

Altera's introduction of its Max 3000A product family may be—at least partially—a response to the aggressive prices of products from CPLD-replacement vendor Clear Logic. Compared with 3.3V Max 7000 devices, Altera tests its Max 3000A parts to relaxed specifications, allowing higher power consumption. The company offers the parts with fewer packaging options, and the 3000A doesn't support all of the Max 7000 features.

Design-tool support comes from Altera's Max+Plus II development software, which is available in both single- and floating-node versions and for multiple operating systems. Altera's tool sets—like tool sets from all of the vendors in this directory—come in multiple versions; support various families of products and devices within each product family; and offer multiple design-entry, simulation, synthesis, fitting, and programming options.

ATMEL

Atmel specializes in both industry-standard devices and backward-compatible supersets of common PALs and CPLDs. The vendor's 16V8, 20V8, and 22V10 parts come in multiple power, voltage, and package variants. If you need to squeeze additional logic into a 22V10 pinout, consider the ATF750; the ATV-2500B delivers even more logic capacity in a 44-pin footprint.

With the ATF1500 series, Atmel has Altera's Max 7000 architecture in the bull's eye. By beefing up both the number of inputs (40) into each logic block and the global-routing switch matrix, the vendor

AT A GLANCE

- ▶ Atmel's 22V10 superset squeezes additional logic onto a low-pin-count device.
- ▶ An increased number of logic-block inputs improves probability of pinout- and performance-locking.
- ▶ An enhanced switch matrix augments larger-macrocell-count CPLDs.

claims that this series can handle designs that wouldn't fit into competitors' devices with comparable macrocell counts. Enhanced connectivity also means that you can more easily maintain your cre-

ation's pinout and performance through multiple design iterations.

Other ATF1500 features include in-system programming, individual I/O-buffer selection, enable capability, optional latch modes for the macrocell flip-flops, independent combinatorial and registered options for macrocell outputs and feedback terms, and three global clock inputs. Atmel offers conversion tools that automatically migrate designs originally targeting other vendors' architectures. The company also provides design-software options that support Abel, schematic-, and VHDL-synthesis design-entry alternatives.

CYPRESS SEMICONDUCTOR

The Ultra37000 family, like Atmel's ATF1500 CPLDs, builds on an Altera Max 7000 foundation with some aggressive assertions. These claims include no fan-out delays, no expander delays, no additional delays for I/O pins versus dedicated pins, no penalty for using the full 16 product terms per macrocell, and no delay for local product-term steering or sharing. There's also no additional delay through the programmable interconnect matrix, because all signals from all macrocells route through the matrix. The device specifications reflect this fact. Available in both 5 and 3.3V variants, the devices range from 32 to 512 macrocells, are in-system-reprogrammable, and come in a variety of packages.

With its Delta39K architecture, Cypress switched from EEPROM to SRAM to take advantage of SRAM's high degree of logic compatibility and leading-edge-lithography status. Attempting to blend the best aspects of CPLDs and FPGAs, Delta39K uses a hierarchical-routing

AT A GLANCE

- ▶ Cypress' traditional-looking Ultra37000 CPLD has nontraditional timing specifications.
- ▶ SRAM-based chips beef up the internal memory, sometimes include the configuration chip, and form the foundation of communications-focused hybrids.
- ▶ Low-cost, high-quality design software completes the picture.

scheme. Because the SRAM's already there for device-configuration memory, Cypress includes substantially more for your use. Each 16-macrocell logic block is nearly identical to those of the Ultra37000, including 36 inputs from the interconnect matrix. Groups of eight logic blocks combine to form a logic-block cluster, which also contains two 8-kbit single-port-RAM arrays.

Outside each logic-block cluster and closely connected to multiple sets of global-routing tracks is an additional 4 kbits of specialty SRAM. This SRAM in-

cludes all of the logic necessary to create dual-port memory or a synchronous FIFO. Cypress also includes PLL circuitry that can multiply (as much as 43 and 266 MHz) or divide (as little as $1/16$) an incoming 25- to 133-MHz clock. This circuitry can also de-skew and phase-shift the clock. The FPGA-reminiscent I/O buffers support numerous electrical protocols, have programmable slew rates and bus hold and contain multiple dedicated registers. Some Delta39K versions combine the CPLD and a configuration flash memory in a dual-die, single-chip package.

Cypress reduced the cost of Delta 39K to develop the Quantum 38K family. The company eliminated the single-port cluster memory and dedicated FIFO logic, along with the embedded PLL, and simplified the I/O buffers. Delta 39K also forms the foundation of the company's ASIC-plus-programmable-logic PSI (programmable-serial-interface) chips, which consolidate SERDES capability. Signaling speeds are 200 Mbps to 1.5 or

2.5 Gbps per serial link, and the family provides a maximum duplex serial bandwidth of 12 Gbps. Supported standards include PCI, SONET, Gigabit Ethernet, Fibre Channel and InfiniBand.

The company is proud of its silicon (which also includes several SPLDs), but it's also pleased with its software. For \$99, you get free technical support and life-

time upgrades for Warp, which supports the entire Cypress product line and includes VHDL and Verilog synthesis, a finite-state-machine editor (for Windows), a static timing analyzer and dynamic timing simulator, fitter software, and reference documentation. You can order and download Warp from Cypress' Web site. The \$175 Warp 2 in-sys-

tem-reprogrammable version includes a programming kit and a demo board. The \$495 Warp Professional adds a flow manager, a block diagram editor, and a language assistant, and \$2995 Warp Enterprise adds a Code2Graphics HDL converter, source-level behavioral simulation, a debugger, and testbench-generation capability.

INTEGRATED CIRCUIT TECHNOLOGY (ICT)

ICT's fortunes lie in the low-gate-count end of the programmable-logic business, where other suppliers fear to tread. ICT's PEEL devices replace 20- and 24-pin PALs with alternatives that have supersets of functions. Factors that differentiate the company's devices include low standby and active power consumption, additional device inputs, additional macrocell product terms, and additional register functions. Speeds are as fast as 5 nsec, with most devices in the 15- to 25-nsec range.

Higher complexity, proprietary PEEL arrays employ an architecture that squeezes maximum logic density into 24- to 44-pin packages. The central logic-ar-

AT A GLANCE

- ▶ Proprietary variants of industry-standard PALs target designs requiring low power and greater flexibility.
- ▶ A PLA and multiple registers create an unconventional CPLD.
- ▶ Tiny PEEL arrays aim to fill a niche between discrete logic chips and today's low-end PALs, with embedded PLDs acting as their on-ASIC equivalent.

ray structure is a PLA, allowing full sum-of-products logic functions to feed—in groups of four—each logic-control cell. The four-input (plus clock, preset, reset, and register-type signals), three-output logic cell contains three multiplexers and

a multipurpose flip-flop that can implement buried register functions. Input and I/O cells also contain registers.

On the other end of the complexity spectrum are ICT's TPAs (tiny PEEL arrays). These small, low-power devices will become available for sampling by the end of 2001. ICT is also pursuing embedded-PLD technology, based on SRAM instead of EEPROM (see "Embedded programmable-logic cores"). The company's Place development software is free to qualified users. It includes an architectural editor, a logic compiler, a waveform simulator, a documentation utility, and a programmer interface. Separate fitters are also available should you prefer to use third-party front-end tools.

LATTICE SEMICONDUCTOR

Lattice first differentiated itself from the competition by bringing in-system programming to PALs. Before this advancement, PALs were based on less flexible ROM, PROM, EPROM, and fuse technologies. In the CPLD arena, the company divides its products into SuperFast, SuperWide, and SuperBig categories. The ispLSI 2000 family ranges from 32 to 192 macrocells; comes in 2.5, 3.3, and 5V variants; and specifies propagation delays as fast as 3 nsec. At first glance, this product range might seem to overlap with the lower end of the Mach 4A device family, which Lattice adopted when it took over Vantis, but dig a little deeper into the specifications.

You can achieve the highest ispLSI 2000 performance when each macrocell consumes no more than four product terms and when the corresponding registers bypass the XOR gates at their inputs and remain in their default D-type

AT A GLANCE

- ▶ In-system programming remains Lattice's key corporate-marketing motto.
- ▶ Mach 4A overlaps SuperFast and SuperWide CPLDs but adds a timing twist.
- ▶ SPLDs, programmable switch and programmable analog devices complete the large product line.

state. Additional product terms and register redefinition require use of the product-term sharing array and incur a timing penalty. With Mach 4A, on the other hand, Lattice guarantees near-ispLSI 2000 speeds with as many as 20 product terms per macrocell—a feature the company calls Speed-Locking.

The ispLSI 5000 family also overlaps Mach 4A, this time at the Mach 4A's higher macrocell counts. This family is notable for the wide 68-input fan-in to each logic block. This specification sounds impressive until you realize that

each ispLSI 5000 logic block also contains 32 macrocells. The input-to-macrocell ratio is actually lower than that of some other CPLD architectures.

Lattice's SuperBig ispLSI 8000 architecture switches from a global to a dual-level hierarchical routing matrix, with local routing interconnecting six groups of 20 macrocells. Lattice calls these groups Megablocks, and global routing stitches together the Megablocks. Each logic block comprises 20 macrocells and has 44 inputs. The ispLSI 8000 devices also contain an embedded 108-line tristate bus. I/O buffers include flip-flops for optional registered input, output, and bidirectional control.

In addition to the previously mentioned programmable-analog devices, Lattice offers a series of SPLDs and the logic-deficient—compared with CPLDs—but routing-rich ispGDX programmable-interconnect architecture. For you fans of nondigital design, the

company also sells ispPAC programmable analog chips. Lattice's PC design software, ispDesignExpert, comes in numerous variants. For workstations, you need

to obtain ispExpert Compiler for Lattice-developed architectures and DesignDirect Summit for the Mach4A. In line with its heritage as the in-system-program-

ming pioneer, Lattice also supplies a variety of programming software, source code, and hardware to help you get your design up and running.

STMICROELECTRONICS

Waferscale Integration's PSDs (Programmable System Devices), now part of STMicroelectronics, began with a simple premise: Combine EPROM with logic to re-create port pins lost when you couple an 8-bit microcontroller with a conventional memory. From those humble beginnings, the product line has expanded in many directions. For example, the company now offers in-system-programmable and -reprogrammable flash memory, including a separate boot array, instead of EPROM.

PSDs also now combine in one chip nonvolatile-code, nonvolatile-data (EEPROM), and volatile-data (SRAM) partitions. The integrated logic has become

AT A GLANCE

- ▶ Memory-plus-logic hybrid chips began with the vision of preserving microcontroller ports while expanding code memory
- ▶ Latest variants combine multiple memory types and increase programmable-logic versatility.
- ▶ Support for 8- and 16-bit processors is now supplemented with DSP-targeted chips, and development tools speed your time to market.

more flexible, useful for general-purpose functions, not just port re-creation, and the same part, after design-specific configuration, can bolt to numerous 8- and 16-bit embedded-controller and -micro-

processor buses. You can program PSDs either offboard using a PROM programmer, onboard under system-processor control, or via a JTAG interface.

Smaller lithographies have boosted read and write speeds, as well as reduced operating voltages and currents. So-called "zero-power" variants employ address-transition detection and other design techniques to further reduce power consumption in standby modes. The latest PSD, the DSM2180F3, mates with Analog Devices' ADSP-218x DSPs and includes eight 16-kbyte flash-memory sectors. STMicroelectronics supplies logic-design tools, device programmers, and evaluation boards for all its PSD devices.

XILINX

Just as Altera is a CPLD vendor that later added FPGAs to its product portfolio, Xilinx is an FPGA supplier that subsequently bought and developed several CPLD-product families. Available in 2.5, 3.3, and 5V versions, XC9500 devices comprise 36 to 288 macrocells and incorporate a fairly mainstream global-routing architecture. One of the parts' most notable characteristics is that they're based on flash memory, not EEPROM. The vendor claims that this setup results in lower cost and higher reprogramming-cycle capabilities. (Most applications have little need for this reprogramming-cycle feature.)

In mid-1999, Xilinx purchased Philips' PAL and CPLD product lines. The acquired devices included a 22V10, a range of 32- to 128-macrocell, low-power but otherwise-conventional CPLDs and two high-macrocell-count, SRAM-based CPLDs conceptually similar to Cypress Semiconductor's Delta39K devices. Going forward, Xilinx is enthused about the CoolRunner XPLA3 architecture, which

AT A GLANCE

- ▶ Xilinx's second-generation architecture makes inroads into CPLD applications.
- ▶ Acquired CPLDs keep their cool and strive for a perfect fit.
- ▶ Hit the Web for free software and other stuff.

combines low power, high performance, and design flexibility. The parts deliver standby power of less than 100 μ A without the need for special power-down bits or pins that negatively affect device performance. Pin-to-pin propagation delays are as low as 5 nsec for the 32-macrocell device, and macrocell counts range from 32 to 384.

As with ICT's PEEL arrays, Xilinx's XPLA3 incorporates a 36348 PLA full-programmable-AND/programmable-OR structure. Along with a variable-function multiplexer and foldback NAND logic, this combination conceptually delivers robust logic optimization and device fitting. However, all of this silicon potential will go to waste without

equally robust software support, which has been a key focus for the product line since the Xilinx acquisition.

The company has two interesting—and free—software products. The Internet-based WebFitter tool accepts designs in VHDL, Verilog, ABEL, EDIF (electronic-design-interoperability format), and XNF (Xilinx Netlist Format). The tool includes competitive architecture-design file-conversion utilities. WebFitter provides links to fitting, timing, and log files; targeted-device data sheets; online price quotes; simulation and device-programming files; online tutorials; and complete online help. The downloadable WebPack-ISE modules include Abel and HDL synthesis, entry-level-simulation and testbench generation, schematic and graphical-state-diagram entry, fitting and place-and-route algorithms, and device-programming utilities. WebPack ISE supports both the XC9500 and CoolRunner CPLD product lines as well as the Virtex V300E, Virtex II (to 300,000 gates), and complete Spartan-II FPGA families.

ACTEL

Actel has chosen a conventional logic-cell approach, consisting of dual three-input LUTs feeding a register, for the VariCore embedded FPGA architecture it acquired when it purchased Prosys Technology in 2000. Reflecting the fact that FPGA compilers are tuned for the more common four-input LUT structure, the Actel logic cell can optionally but less efficiently group the two three-input LUTs into one four-LUT equivalent.

A review of Actel's VariCore data sheet reveals some curious omissions: hardware multipliers or other arithmetic-optimized structures, three-state buffers or bidirectional buses, and partial-reconfiguration capability. The embedded FPGA arrangement does, however, include built-in vertical-carry chains that optionally connect to one of the three LUT inputs. Each functional group contains four pairs of three-input LUTs and four registers. The registers share common control inputs (enable, preset, reset, and clock).

The next level of hierarchy is the PEG (programmable-embedded-gate) array, an 8×8 matrix of functional groups that Actel estimates represents a multiple-design average of 2500 ASIC gates, excluding RAM. Actel, along with partners Chartered, TSMC, and UMC, initially plans to offer 2×1 , 2×2 , 4×2 , and 4×4

AT A GLANCE

- ▶ Embedded programmable logic enables Actel to partner with ASIC companies that it competes against with its FPGAs.
- ▶ Conventional logic-cell structure is programmable-logic-synthesis-friendly.
- ▶ Multiple logic-array configurations address numerous space, power, and performance constraints.

PEG layouts on 0.18-micron process technologies. You can orient these multi-PEG arrangements in square, rectangular, and even L-shaped structures.

According to Actel, the ideal VariCore shape is square, because it minimizes the internal delays within the FPGA core. However, in some applications, a different-shaped core provides a more efficient implementation at the physical level. Also, each of the PEG blocks with external edges provides 48 inputs and outputs per horizontal edge and 32 inputs and outputs per vertical edge. A 2×2 -PEG array offers a maximum of 640 inputs and outputs, whereas a 4×1 equivalent contains 704 inputs and outputs.

The VariCore place-and-route software is parameterized to support multi-PEG-matrix sizes and orientations, thereby not forcing the HDL source code that you develop at the beginning of the

design process to explicitly specify the final physical layout. Actel recommends that you obtain and use a separate programmable-logic-optimized synthesis compiler for the portion of your design that goes into VariCore.

The architecture arranges both the logic structures and the routing between them in a hierarchical manner. Actel believes that this approach offers flexibility, predictability in timing and logic usage, and faster place-and-route times. The vendor claims that a design with 70% three-LUT usage on a 4×4 -PEG FPGA completely compiles on a 500-MHz PC in about nine minutes and that designs routinely and easily use close to 100% of the functional groups. Each multi-PEG array also contains JTAG circuitry and a BIST interface but not a BIST controller.

Because of the fixed silicon cost of overhead circuitry, Actel doesn't anticipate offering single-PEG arrays, at least at the 0.18-micron-process generation. PEG arrays can also include optional cascadable RAM blocks of 9 kbits each with both 9- and 18-bit data-interface options and built-in FIFO flag logic. You might use these RAM structures, for example, as bridges between ASIC and FPGA logic running at different bus widths and clock speeds. The company's future plans also include offering its ProASIC flash FPGA technology in embedded-IP form.

ADAPTIVE SILICON

Adaptive Silicon's approach to embedded programmable logic employs an ALU-based logic cell whose heritage extends back to the days when some of the company's founders were employed by National Semiconductor and developed the CLay (Configurable Logic Array) and NAPA (National Adaptive Processing Architecture) FPGAs. The fundamental building block of Adaptive's MSA (multiscale-array) 2500 is an enhanced 74LS181 4-bit ALU with a modular carry-look-ahead scheme and optional output registering. Four ALUs group to form a Quad Block that connects to other Quad Blocks by signal routing and ALU-control resources, and

AT A GLANCE

- ▶ MSA 2500's ALU-based logic cell is particularly efficient at implementing arithmetic-dominated circuits.
- ▶ Adaptive Silicon strives to make as few changes as possible to your ASIC design flow.
- ▶ LSI Logic is ready for production, with TSMC and other as-yet-unnamed partners waiting in the wings.

16 Quad Blocks form a Hex Block.

Adaptive Silicon estimates that each Hex Block represents an average of 1500 ASIC gates. The company points out that the logic cell can also implement more generic LUT-like functions via both gate-

and ALU-level-mapping capabilities. The interface between programmable- and standard-cell partitions comprises a PLC (programmable-logic-core) control structure in the programmable partition (along with an application-circuit interface) and a matching PLC adapter in the ASIC partition. The vendor has 0.18-micron test chips in-house from initial partner and investor LSI Logic (which calls its implementation of Adaptive Silicon's technology the LiquidLogic core) and TSMC.

Adaptive Silicon's Millennium PLC design flow is ASIC-designer-friendly. The synthesis tool set for the programmable-logic partitions is the same one you use for standard-cell-targeted logic:

Synopsys' Design Compiler or an equivalent. Adaptive Silicon relies on DesignWare libraries that it developed to perform the necessary translation and architectural optimization of VHDL or Verilog source code. The registers in each logic cell, with scan chains throughout the array, aid in silicon verification. Adaptive Silicon and its ASIC partners tailor the PLC adapter for each application with a soft macro that manages the programming and test interfaces, a five-

pin test-access port, and multiplexed test and data I/O signals. The programmable-logic core comprehends full-chip scan BIST to ensure programming integrity, quiescent power testing, and full-speed testing.

Both Adaptive Silicon and competitor Actel strive to seamlessly integrate their embedded FPGA partitions into your chip's layout, functional verification, and static-timing-analysis flows, with GDS II floor-planning layout frames, SDF tim-

ing models, and the like. Adaptive Silicon supports partner LSI Logic's CoreWare design flow, including the layout-versus-schematic, transistor-level check before tape-out, even though a dynamically configured programmable-logic partition uses few of the available transistors. The MSA 2500 pricing scheme includes a technology-licensing fee that can be as little as \$300,000 and volume-dependent per-chip royalties.

AGERE SYSTEMS

Agere Systems has a multipronged strategy for addressing its customers' needs for design flexibility. As discussed in the FPGA section of this directory, the company offers both "pure" programmable-logic devices and ASIC-plus-programmable-logic hybrid devices. Mask-programmable versions of the company's FPGAs give you a cost-reduction path once your design reaches high production volumes.

Agere acts as both a standard product supplier and, through its ASIC group, a more general-purpose silicon foundry.

Among the numerous IP cores available to you from the ASIC division is a Series 4 FPGA macrocell building block comprising 800 logic cells. You can either fully embed the FPGA macrocell or leave one of the four sides open to external I/O buffers. Integration of multiple embedded FPGA blocks is possible, and if FPGA logic doesn't meet your cost or performance requirements and you don't require ultimate flexibility, you might instead be interested in employing the laser-programmable embedded logic cores that Agere offers via a licensing agreement with Chip Express.

AT A GLANCE

- ▶ Field- and mask-programmable logic provides options for different stages in your products' life cycles.
- ▶ The Series 4 FPGA core forms the programmable foundation of many of Agere's hybrid chips and is available for use in your custom designs.
- ▶ Laser-programmable-logic cores might not be user-configurable, but they tend to be more silicon-efficient and offer higher performance than FPGAs, all other factors being equal.

ATMEL

As a provider of both FPGAs and ASICs, Atmel (along with Agere Systems) is a rarity among programmable-logic suppliers. The company not only produces FPGAs and FPGA-plus-ASIC hybrid chips of its own, but also supplies embedded programmable-logic cores for your ASIC designs. **Table 3** at the Web

version of article at www.ednmag.com shows that a range of logic-cell counts and associated RAM-array sizes are possible, a reflection of the fine-grained AT40K FPGA architecture that the company envisioned embedding from the beginning.

AT A GLANCE

- ▶ Embedded cores bring programmable logic to your ASICs.
- ▶ Fine-grained logic-cell and routing structures lead to a high degree of design flexibility.
- ▶ FPSLIC (field-programmable system-level-integration-circuit) devices are proof of concept of Atmel's hybrid capabilities.

INTEGRATED CIRCUIT TECHNOLOGY (ICT)

PLA vendor ICT hopes to migrate its programmable AND/OR PAL technology to the world of IP cores. The company points out that the portions of a design most likely to change, such as state machines, control-signal translation, and address decoders, are ideal applications for PLAs. ICT isn't trying to convince customers to embed large chunks of product-term-based logic on their chips. Such an approach would quickly become

AT A GLANCE

- ▶ ICT brings timing predictability to the world of embedded programmable logic.
- ▶ Exponentially growing routing-matrix sizes make the approach uneconomical with large logic densities.
- ▶ SRAM-based configuration allows for low cost and maximum foundry portability.

cost-prohibitive, particularly in a global-routing-matrix structure. Instead, the

company wants you to scatter small PLAs across the ASIC.

With its embedded cores, ICT is moving away from the EEPROM-based technology it uses in its stand-alone chips toward an approach employing SRAM to store configuration-bit values. Embracing SRAM in this manner enables the company to leverage the most advanced processes at multiple ASIC vendors and foundries.

ACRONYMS

Abel: Advanced Boolean Equation Language
AGP: accelerated graphics port
ALU: arithmetic-logic unit
ASIC: application-specific integrated circuit
BGA: ball-grid array
BIST: built-in self-test
BLVDS: bus low-voltage differential signaling
BQFP: bumped quad flatpack
CABGA: ceramic ball-grid array
CAM: content-addressable memory
CDIP: ceramic dual inline package
CDR: clock and data recovery
CG: plastic ball-grid array
Clay: configurable logic array
CLB: configurable logic block
CLCC: ceramic leadless chip carrier
CPGA: ceramic pin-grid array
CPLD: complex programmable-logic device
CPU: central processing unit
CQFP: ceramic quad flatpack
CS: chip scale
CSOC: configurable system on chip
CSP: chip-scale package
DIP: dual in-line package
DLL: digital delay-locked loop
DMA: direct-memory access
DRAM: dynamic random-access memory
DSP: digital-signal processor
DVB: digital-video broadcasting
EAB: embedded array block
EBGA: enhanced ball-grid array
ECU: embedded computational unit
EDIF: electronic-design-interoperability format
EEPROM: electrically erasable

programmable read-only memory
EPROM: erasable programmable read-only memory
ESB: embedded system block
ESP: embedded standard product
FBGA: fine-pitch ball-grid array
FC: fibre channel
FCBGA: flip-chip ball-grid array
FEC: forward error correction
FG: fine-pitch ball-grid array
FIFO: first in/first out
FPBGA: fine-pitch ball-grid array
FPGA: field-programmable gate array
FPSC: field-programmable system chip
FPSLIC: field-programmable system-level integration circuit
GTL: Gunning transceiver logic
HDL: hardware-description language
HQ: high-heat-dissipation quad flatpack
HSTL: high-speed transceiver logic
IEEE: Institute of Electrical and Electronics Engineers
I/O: input/output
IP: intellectual property
ISP: in-system programmable
JTAG: Joint Test Action Group
LAB: logic-array block
LQFP: low-profile quad flatpack
LUT: look-up table
LVMOS: low-voltage complementary metal-oxide semiconductor
LVDS: low-voltage differential signaling
LVPECL: low-voltage positive emitter-coupled logic
LVTTL: low-voltage transistor-to-transistor logic

MSA: multiscale array
MIPS: millions of instructions per second
NA: not applicable
NAPA: National Adaptive Processing Architecture
NRE: nonrecurring engineering
PAL: programmable array logic
PBGA: plastic ball-grid array
PBGAM: plastic ball-grid array, multilayer
PCI: Peripheral Component Interconnect
PCM: programmable clock manager
PDIP: plastic dual inline package
PECL: positive emitter-coupled logic
PEG: programmable embedded gate
PFU: programmable function unit
PGA: pin-grid array
PLA: programmable-logic array
PLC: programmable-logic core
PLCC: plastic leaded chip carrier
PLD: programmable-logic device (product-term-based)
PLICE: programmable low-impedance circuit element
PLL: phase-locked loop
PROM: programmable read-only memory
PQ: plastic quad flatpack
PQFP: plastic quad flatpack
PSD: programmable system device
PSI: programmable serial interface
QDR: quad data rate
QFP: quad flatpack
RAM: random-access memory
RISC: reduced-instruction-set computer
ROM: read-only memory

RQFP: power quad flatpack
RTL: register-transfer level
SDF: standard data format
SDH: synchronous digital hierarchy
SERDES: serializer/deserializer
SLIC: supplemental-logic-and-interconnect cell
SMPTE: Society of Motion Picture and Television Engineers
SOIC: small-outline integrated circuit
SONET: synchronous-optical network
SOP: small-outline package
SPLD: simple programmable-logic device
SQFP: shrink quad flatpack
SQFP2: power shrink quad flatpack
SRAM: static random-access memory
SSOP: shrink small-outline package
SSTL: stub-series-terminated logic
TPA: tiny PEEL arrays
T_p: propagation-delay time
TQ: thin quad flatpack
TQFP: thin quad flatpack
TSSOP: thin shrink small-outline package
UART: universal asynchronous receiver-transmitter
VHDL: very-high-speed integrated-circuit hardware-description language
VQFP: very thin quad flatpack
VTQFP: very thin quad flatpack
XAU1: extended attachment-unit interface
XNF: Xilinx Netlist Format

REFERENCES

1. Dipert, Brian, "Do combo chips compute (or even compile)?," *EDN*, Feb 15, 2001, pg 101.
2. Dipert, Brian, "EDN's first annual programmable-logic directory," *EDN*, Aug 17, 2000, pg 54.
3. Dipert, Brian, "Cunning circuits confound crooks," *EDN*, Oct 12, 2000, pg 103.
4. Dipert, Brian, "Low-cost programmable logic: How low should you go?," *EDN*, March 16, 2000, pg 123.
5. Dipert, Brian, "The best (or worst?) of both worlds," *EDN*, Nov 11, 1999, pg 139.
6. Dipert, Brian, "Figuring out reconfigurable logic," *EDN*, Aug 5, 1999, pg 103.
7. Dipert, Brian, "Lies, damn lies, and benchmarks: The race for the truth is on," *EDN*, May 27, 1999, pg 54.
8. Dipert, Brian, "Synthesis shoot-out at the *EDN* corral," *EDN*, Sept 11, 1998, pg 95.
9. Dipert, Brian, "Counting on gate counts? Don't count on it," *EDN*, Aug 3, 1998, pg 52.
10. Dipert, Brian, "Moving beyond programmable logic: if, when, how?," *EDN*, Nov 20, 1997, pg 77.
11. Dipert, Brian, "Programmable logic: Beat the heat on power consumption," *EDN*, Aug 1, 1997, pg 57.
12. Dipert, Brian, "Shattering the programmable-logic speed barrier," *EDN*, May 22, 1997, pg 36.