By Brian Dipert, Technical Editor

EDN's PAL, PLD, AND FPGA 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.

VITAL STATS

ANNUAL PROGRAMMABLE-LOGIC DIRECTORY THIRD ANNUAL PROGRAMMABLE-LOGIC DIRECTORY THIRD ANNUAL PROGRAMMABLE-LOG

AND CHECK THIS OUT: WE'VE POSTED COMPREHENSIVE TABLES OF DEVICES AND FEATURES IN THE WEB VERSION OF THIS ARTICLE AT WWW.EDN.COM.

Programmable-logic devices are the fastest growing segment of the logic-device family, for two fundamental reasons. Their ever-increasing gate count per device gathers up functions that you might otherwise spread over a number of discrete-logic and memory chips. This approach improves end-system size, power consumption, performance, reliability, and cost. Equally important, you can in 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 needs and competitive pressures.

The term "logic devices" subdivides into discrete logic, simple and complex PLDs, FPGAs, and standard and custom-cell ASICs. FPGAs, SPLDs and PALs, and CPLDs are all programmable-logic devices, although their internal architecture implementations differ. I use the term "programmable-logic device" to refer to the entire range of products and specify a category via its abbreviation.

Compared with ASICs, programmable-logic devices have shorter leadtimes, lower upfront NRE charges, and no minimum-order quantities, all of which simplify inventory management. As per-gate cost decreases and as the number of gates per component increases, programmable-logic devices are making significant inroads into traditional gate-array—and to a smaller extent standard-cell—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 implement reconfigurable computing—that is, using a fixed number of logic gates to time-division-multiplex multiple functions.

Just as your company uses programmable logic's flexibility to differentiate itself from your competition, semiconductor vendors have developed unique PLDs and FPGAs to address specific intersections of performance, power, integration, and cost targets. This diversity is perhaps the most complex challenge you face, because, in many cases, you must analyze each programmable-logic architecture in detail before selecting 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. Prioritizing these parameters depends on the application. Often, for example, designs targeting low power, high performance, or minimal gate count significantly differ 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 this priority ordering,
- 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 or failure. Equally important are the depth and the breadth of the company's internally developed and third-party-provided 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 decreasing time to market 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 silicon architectures are important.

Although "pure" programmable-logic devices comprise a dominant percentage of industry unit 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 attempting to optimally address the complicated power-versus-performance-versus-cost balancing act. ASIC vendors and foundries, too, along with standard product suppliers such as communications-processor manufacturers, are developing hybrid-device capabilities. And companies 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

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. Five main antifuse device families form the foundation of today's Actel product line: the 5V MX, 3.3V SX, and closely related 2.5V SX-A, 2.5V eX, and 1.5V Axcelerator. Actel also offers the flash-memory-based ProASIC and ProASIC Plus FPGA architectures and the Varicore SRAM-based embedded FPGA (see the "Embedded Programmable Logic" section of this directory). And, by the end of the year (unfortunately, the same thing I said this time *last* year...), the company hopes to begin shipping samples of its first flexible-interface BridgeFPGA devices.

Antifuse delivers lower impedance and, therefore, lower power and higher speed signal interconnection and more robust immunity to high-radiation operating environments than other configuration technologies. 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 antifuse and, for that matter, flash parts are also nonvolatile. Unlike SRAM-based FPGAs, they require no separate memory to store their configuration data, and their functions are immediately available upon system power-up. The single-chip nature of antifuse and, to a smaller extent, flash FPGAs also makes them nearly impossible to reverse-engineer or clone. This characteristic becomes increasingly important as decreasing costs and higher capacities result in the parts finding use in high-volume consumer products.

The company's first-generation PLICE (programmable-low-impedance-circuit-element) antifuse-technology approach, which Actel used in the MX and earlier product families, employs a metal-to-metal interconnect structure comprising polysilicon and a diffused N+ region, separated by a high-impedance oxide-nitride-oxide barrier. A high programming voltage ruptures this barrier. The PLICE antifuse structures, which reside on the same base layer as active circuit elements, take up 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, eX, and Axcelerator devices employ a sea-of-modules ratio of C-Cell and R-Cell logic structures. C-Cells contain a dual-level, two-input multiplexer structure plus input-inversion capability. R-Cells contain multiple-function flip-flops with numerous signal-input, clocking, reset, and clear options. Latest generation Axcelerator parts include embedded-memory blocks, PLLs, and other circuits.

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 line of flash-memory-based FPGAs.

In addition to being onboard-programmable and -reprogrammable, ProASIC devices use an extremely fine-grained logic cell. Actel claims that this architecture 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, like Axcelerator, brings embedded-SRAM and FIFO capability to Actel's arsenal. "Plus" parts, built on more advanced processes, also offer a higher memory-to-logic-resource proportion.

Actel's Libero and Designer tool sets support antifuse FPGAs, with several variants supplying 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 device programmer supports configuration of FPGA prototypes right at your PC. For ProASIC FPGAs, the vendor-supplied tools include ASICmaster for placement and routing, MEMORYmaster for embedded-memory-function generation, and Flash Pro for programming. IP comes both from Actel and from the company's strategic partners.

AT A GLANCE: ACTEL

- Actel's antifuse technology delivers numerous benefits at the expense of limited flexibility.
- The company's second-generation antifuse structure consumes no silicon-substrate area.
- Flash memory builds on antifuse strengths and, after a lengthy gestation period, is in production on its second product generation.

ALTERA

From the Flex 8000 and follow-on Flex 10K architectural foundations, Altera has taken its FPGA, or, as the company prefers, LUT (look-up-table)-based programmable-logic-device, product line in multiple directions, all based on a common, essentially unchanged LAB (logic-array-block) structure. The Acex 1K family is a lower cost Flex 10KE variant, built on a smaller lithography process. Like its Flex 6000 predecessor, Acex 1K comes with fewer gates and with more restricted, less expensive packaging options than its Apex 20K and other Altera big brothers.

Apex 20K includes as many as four on-chip PLLs. The device family also provides 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. They're great if you can use all that memory. If not, you won't come close to squeezing into one chip the design sizes that Altera's marketing might indicate are possible. Apex 20KE enables you to use the EABs for implementing not only single- and dual-port RAM and FIFOs, but also small CAMs (content-addressable memories). Apex 20KC migrates the Apex 20KE architecture to a process employing low-impedance copper interconnect lines for all metal layers.

Altera's first generation of CPU-inclusive Excalibur XA hybrid chips, based on the ARM architecture, is now in production, although the planned MIPS-based XM parts are on indeterminate hiatus. Combining Apex 20KE with eight to 18 channels of ASIC-housed CDR (clock-and-data-recovery) circuitry and slightly modifying the logic-block structure give you the Mercury family, supporting data throughput speeds of 125 Mbps to 1.25 Gbps. And doubling the RAM-to-ESB ratio in Apex 20KC, increasing the number of PLLs, accelerating and adding I/O buffers, and making those buffers more flexible give you the Apex II.

Historically, Altera's FPGAs have employed a long-line routing-dominated architecture that provides timing predictability. Additionally, Altera's predominant routing approach enables on-chip redundancy, particularly at the early stages of process and device production, to improve yields. However, the company's latest device family, Stratix, adopts a three-length mix more significantly shifted toward short routes, with *no* lines spanning the die as in the past. The resultant "tile" architecture should speed the development of future Excalibur-like devices that blend "soft" programmable logic and "hard" diffused circuits. Stratix also provides numerous MACs (multiply-accumulators) and three onboard-memory structures, each with application-optimized array numbers, densities, and bus widths.

Altera has developed an optimized 16- and 32-bit RISC processor it calls Nios, which aims to 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 company's cost-reduction vehicle for high-volume designs, the HardCopy program, involves creation of mask-programmable FPGA variants. The same Max+ software you'd use to design with Altera's CPLDs also supports Acex 1K and Flex FPGAs. For Apex and Stratix devices, you'll want to fire up Altera's more advanced Quartus II development-tool environment, along with the SignalTap logic analyzer.

AT A GLANCE: ALTERA

- 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.
- Latest developments include hybrid chips with industry-standard ASIC-housed CPU cores, a soft CPU core, a next-generation FPGA with revamped routing and an internal ASIC-conversion program.

ATMEL

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

AT6000 logic cells comprise several fixed-function logic gates plus a register. In AT40K FPGAs, you find the more common LUT (look-up-table)-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, you find a dedicated two-input AND gate. This logic tweak combines with diagonal routing to enable AT40K FPGAs to better support the matrix-multiplication operation so common in DSP functions.

Unlike with some other FPGAs, you can't alternatively use the AT40K LUTs to implement 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 find a 128-bit dedicated RAM array. This approach is an intermediary step between the small, 16-bit, look-up-table-derived memory arrays and the much larger dedicated SRAM blocks in other vendors' architectures.

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 the 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 of the FPSLIC architecture, AT49S, is a family of single-chip, dual-die devices with a custom nonvolatile-memory array. Containing two data buses, one going to the system and the other running only within the pack-

AT A GLANCE: ATMEL

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

age between the memory and AT49K, the Secure FPSLIC's memory also includes a security bit. Once you set this bit, the chip responds only to a full-chip-erase command. You must set the security bit to initiate memory-to-AT49K communication. The approach doesn't eliminate the security threat; although the external bitstream trace between memory and AT49K is no longer present, stripping back the multidie package lid exposes it for probing. But it 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 intellectual-property portfolio, with FPSLIC as the proof of capability (see the "Embedded programmable logic" section of this directory).

LATTICE SEMICONDUCTOR

Lattice Semiconductor, like Altera before it, is a company whose programmable-logic roots lie in the world of PALs and PLDs. Also like Altera before it, Lattice is broadening its focus to include FPGAs, much as FPGA pioneer Xilinx has expanded into CPLDs. This broadening has come through acquisitions via the purchase of Agere Systems' (formerly, Lucent Technologies and, before that, AT&T Microelectronics) Orca product line and through internal development in the form of the ispXPGA product family.

Orca Series 2 devices group four four-input LUTs (look-up tables) and four registers into each PFU (programmable function unit). Like Xilinx and unlike Altera, each LUT grouping can alternatively find use 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 readback. The PCM (programmable clock manager), a PLL variant, enables adjustment of input clock phase and duty cycle. The size of each Series 3 PFU more than doubles that of Series 2, now including eight four-input LUTs and nine registers.

AT A GLANCE: LATTICE

- Orca logic-structure innovations require synthesis support or core-generator intercession.
- Hybrid devices handle high-speed interfaces.
- The internally developed ispXPGA offers multimemory integration and synthesis-friendly logic.

Series 4 devices focus attention on signal routing, based upon a recognition that delays in this area—not in logic—are increasingly defining the upper limit of design performance. Abundant metal layers create—along with other uses—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. Regardless of its logic density, each Series 4 device includes six general-purpose and two application-specific PLLs. Series 4 parts supplement 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 (content-addressable-memory) 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. An on-chip, ARM-derived, multimaster peripheral bus both simplifies the interconnection of multiple logic blocks and gives a glimpse of potential future CPU-integration plans. The Series 4-based ORT8850 hybrid chips contain an eight-channel, 850-Mbps CDR (clock-and-data-recovery) 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.

Within months of acquiring the Orca line, Lattice, somewhat surprisingly, rolled out an internally developed FPGA family, ispXPGA, which includes on a single die not only the power-dependent SRAM-based logic- and routing-configuration cells and memory blocks, but also the nonvolatile-EEPROM configuration-storage arrays. Aside from its volatile-plus-nonvolatile-memory amalgamation, Lattice's ispXPGA is in many other respects a conventional modern FPGA. Each PFU contains four synthesis-friendly, four-input LUTs, which can alternatively implement a 64-bit, single-port or a 32-bit, dual-port distributed-memory block; eight registers; dedicated hardware to implement counters, multiplexers, adders, and conventional and Booth multipliers; and input and output switch matrices.

Discrete block RAMs, configurable as either 512×9 or 25×18 bits, support FIFO and single- and dual-port-memory structures with 3-nsec access times and both width- and depth-cascading capability. Each ispXPGA programmable-I/O structure, residing between the PFUs and the I/O buffers, contains input, output, and output-enable registers and offers configurable output slew rate and input delay; this delay enables a zero-hold-time configuration. PLLs and associated 850-Mbps serializer/deserializer circuits, along with built-in 10B/12B support, form the heart of the chips' high-speed serial interfaces. Lattice's ispLever design software works either stand-alone or with third-party compilers and simulators, and the company's development-tool arsenal also includes numerous evaluation boards, programmers and intellectual-property 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 later did Actel begin to match ViaLink's benefits. QuickLogic's chips have all of the inherent antifuse advantages (see Actel entry in this directory, pg 46),

including the availability of extended-temperature and military-screened device variants. Of all the programmable-logic manufacturers, QuickLogic has also been the most enthusiastic about embracing the hybrid ASIC-plus-FPGA approach and the most aggressive about rolling out corresponding devices.

The company builds its pASIC 2 and 3 product families on smaller lithographies than 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 2-to-1 multiplexers, and a D flip-flop. The logic cell's 14 inputs allow it to implement—in one logic level—functions that might require multiple performance-sapping logic cells in other approaches. 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 Eclipse FPGAs focus first on improving the I/O buffers. They 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 for 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 32- and 64-bit; 33-, 66-, and 75-MHz; and master and slave PCI-core flavors. 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.

AT A GLANCE: QUICKLOGIC

- Logic cells promote design flexibility.
- ESPs integrate RAM and system-interface modules.
- ASIC-housed DSP, PCI, Fibre Channel and serializer/deserializer functions are now available, along with a MIPS processor and peripheral suite.

QuickLogic turned to the Eclipse architecture to construct its next ESP families. QuickDSP, now renamed EclipsePlus, 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) with multiple passes through the ECU supporting the common multiply-accumulate function. QuickSD devices integrate as many as eight serializer/deserializer data channels and two clock channels, as well as all bus-LVDS for high-current and long-signal-drive capability. QuickMIPS fills the hole created by the suspension of Altera's MIPS-based Excalibur program. A MIPS32 4Kc processor combines with a separate arithmetic unit, 16 kbytes each of instruction and data caches, and 32-bit Advanced High Performance and Advanced Peripheral Buses. Peripherals include dual 10/100-Mbps Ethernet ports; a 32-bit, 33/66-MHz PCI host; a high-bandwidth memory controller; 16 kbytes of high-speed SRAM; an interrupt controller; dual serial ports; and 32-bit timer-counters. QuickLogic's MIPS license also includes an option for the MIPS64 5Kc processor core.

Design-software support comes from QuickLogic's QuickWorks for PCs and QuickTools for workstations. Via the online WebASIC program, you can upload your design's bit-stream file and receive free samples in 24 hours (for North American customers). The company also supplies the QuickPro desktop-device programmer and various 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

The original intent of the programmable-logic arrays on Triscend's E5 and A7 CSOCs (configurable systems on chips) was to implement customer-specific microprocessor peripherals. However, recent revisions of both Triscend's FastChip tools and partners' design software enable their 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 programmable-logic cells translate to as many as 40,000 ASIC gates. E5 devices include two dedicated DMA channels, and an on-chip breakpoint unit provides debugging capability.

AT A GLANCE: TRISCEND

- Microcontroller-peripheral-intended configurable logic now has more general-purpose use.
- The 8051- and ARM-based hybrid-chip families address 8- and 32-bit computing needs.
- SuperH-based devices are on the horizon.

Triscend's A7 family brings the CSOC concept to the 32-bit processor world. In this case, the CPU core is the ARM7-TDMI, along with ASIC-housed memory controller, four-channel DMA controller, JTAG interface, 16-input interrupt controller, dual timer/counters, dual serial ports, a watchdog timer, and other circuits. Based on a joint-development agreement the company made in January 2001 with Hitachi, SuperH-based CSOCs are now on Triscend's road map.

XILINX

Since Xilinx's founding in 1984, the LUT (look-up-table)-plus-register combination has been a consistent element in its devices, but the logic cell and peripherals have gone through a number of evolutions 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.

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 decoding 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 architecture leap forward. 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, which, as in earlier architectures, you can reconfigure for use as distributed-RAM arrays, and four registers. Xilinx supplements the distributed-LUT memory with multiple 4-kbit true 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 devices don't alter the Virtex logic-cell structure but dramatically boost the amount of discrete SRAM, leading to an exponentially higher claimed gate count.

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 expands the memory-to-logic proportion at a four-gate-per-SRAM-bit counting estimate; this device family marks the emergence of copper routing at the upper two metal layers. This 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, 2.5V Spartan II and 1.8V Spartan IIE retain almost all of Virtex's features

Xilinx's 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 continues this trend with Virtex II, which offers eight LUTs and eight registers per CLB (configurable-logic block). With all of that available metal, Xilinx boosted the amount of long-line, LUT-to-LUT, and CLB-to-CLB routing. Virtex II includes dedicated 18-bit, fast-multiplier logic, which according to Xilinx can perform more than 600 billion 8-bit multiply-accumulate operations/sec. Virtex II also adds Digitally Controlled Impedance Technology; 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.

Like chief competitor Altera, Xilinx has embraced the single-chip ASIC-plus-FPGA concept with Virtex-II Pro, a product family containing high-speed serial transceivers and as many as four integrated IBM PowerPC CPU cores. And, also like Altera, Xilinx has a soft-CPU core, the 32-bit MicroBlaze, in its intellectual-property arsenal. However, unlike Altera, Xilinx doesn't believe in an *all*-ASIC HardCopy-like path to cost reduction. Turning its back on its own HardWire heritage, Xilinx is now advocating EasyPath, which evaluates parts with a customer design-specific subset of the normal testing flow. Partner IBM will also implement Xilinx-based FPGA cores within its ASICs by 2010 as an alternative path to hybrid devices.

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- and WindRiver-based software products that enable chip programming and reprogramming at the design bench, on the manufacturing line, and in the field. Partnerships with Cadence and the MathWorks simplify your task of implementing DSP functions in FPGAs.

AT A GLANCE: XILINX

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

PALs AND PLDs

ALTERA

In the late 1980s, Altera pioneered the concept of the CPLD, a device with numerous logic blocks, each comprising a PAL- or an SPLD-like group of macro-

AT A GLANCE: ALTERA

- Over the years, Altera has performed only minor fine-tuning of its Max 7000 flagship CPLD-product line.
- Hierarchical routing can't match Moore's Law and comes with timing trade-offs.
- Relaxed and restricted testing, along with plastic packaging, led to low-cost Altera devices.

cells. These logic blocks interconnected to each other and to 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 to not only lower cost per macrocell, but also higher speed, lower operating voltage and power consumption, and higher macrocell counts. Along the way, Altera has fine-tuned the Max 7000 in numerous ways. Perhaps the most significant tweak 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 increase the macrocell capacity of its CPLDs, Altera in 1994 introduced the Max 9000 family. This CPLD architecture migrates from a monolithic logic-block-to-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 timing and logic-block-to-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.

Compared with 3.3V Max 7000 devices, the company tests its cost-optimized Max 3000A parts to relaxed specifications, allowing higher power consumption. Altera offers the parts with fewer packaging options and 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, as with those of many 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.

ANACHIP

Anachip, which this spring acquired Integrated Circuit Technology's product line, believes that its fortunes lie in the low-gate-count end of the programmable-logic business where other suppliers fear to tread. 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 high as 5 nsec, and most devices tout 15- to 25-nsec speeds.

More complex proprietary PEEL arrays employ an architecture that squeezes maximum logic density into 24- to 44-pin packages. The central logic-array structure is a PLA, allowing full sum-of-products logic functions in groups of four to feed 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 Integrated Circuit Technology's TPLDs (tiny programmable-logic devices). 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 for those of you who prefer to use third-party front-end tools.

AT A GLANCE: ANACHIP

- Proprietary variants of industry-standard PALs target designs requiring low power and greater flexibility.
- A PLA and multiple registers create an unconventional CPLD.
- TPLDs aim to fill a niche between discrete logic chips and today's low-end PALs.

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.

AT A GLANCE: ATMEL

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

With the ATF1500 series, Atmel has Altera's Max 7000 architecture in the bull's eye. By beefing up the number of inputs to 40 into each logic block and the global-routing switch matrix, the vendor claims that this series can handle designs that wouldn't fit in competitors' devices with comparable macrocell counts. Enhanced connectivity also means that you can more easily maintain your creation'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. No additional delay occurs 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 have 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 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 that 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.

There are 4 kbits more of specialty SRAM outside each logic-block cluster and closely connected to multiple sets of global-routing tracks. This SRAM includes 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 4 and 266 MHz) or divide (as little as $^{1}/_{16}$) an incoming 25- to 133-MHz clock. This circuitry can also deskew 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 cost-reduced Delta 39K to come up with 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 for the company's ASIC-plus-programmable logic PSI (programmable-serial-interface) chips, which consolidate serializer/deserializer capability. Signaling speeds range from 200 Mbps to 1.5 Gbps or 2.5 Gbps per serial link, and the family provides a maximum duplex serial bandwidth of 12 Gbps. The devices support the PCI, SONET, Gigabit Ethernet, Fibre Channel, and InfiniBand standards.

The company is proud of its silicon, which also includes several SPLDs, but it's also pleased with its software. Warp supports the entire Cypress product line and includes VHDL and Verilog synthesis, a finite-state-machine editor for Windows, a static timing analyzer, a dynamic timing simulator, fitter software, and reference documentation. You can order and download Warp from the Cypress Web site. The Warp in-system-reprogrammable version includes a programming kit and demo board. Warp Professional adds a flow manager, block diagram editor, and language assistant, and Warp Enterprise adds a Code2Graphics HDL converter, source-level behavioral simulation, a debugger, and testbench-generation capability.

AT A GLANCE: CYPRESS

- 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 for communications-focused hybrids.
- Low-cost, high-quality design software completes the picture.

LATTICE SEMICONDUCTOR

Lattice first differentiated itself from the competition by bringing in-system programming to PALs. Before this advancement, manufacturers based PALs 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 3.3V ispLSI 2000 family ranges from 32 to 192 macrocells 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 state. Additional product terms and register redefinition require use of the product-term sharing array and incurs 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.

Lattice's ispMACH 4000 parts, architecturally based on ispMACH 4A, come in 1.8, 2.5, and 3.3V variants. The 3.3V ispLSI and 2.5V ispMACH 5000 families also overlap Mach 4A, this time at the Mach 4A's higher macrocell counts. These families are notable for the wide 68-input fan-in to each logic block. This specification sounds impressive until you realize that each logic block also contains 32 macrocells. The input-to-macrocell ratio is actually lower than that of some other CPLD architectures.

The newest architecture, ispXPLD, is SRAM-based, a trait it shares with Cypress' Delta39K and Quantum38K families. Unlike Cypress' devices, however,

Lattice includes the nonvolatile configuration memory on the die. Each ispXPLD multifunction block, when in product-term mode, is reminiscent of ispLSI 5000 devices, with 68 inputs and 32 macrocells. Alternatively, the multifunction block can construct an 8-kbit dual-port RAM, a 16-kbit single-port or pseudo (one read/write port, one read-only port) dual-port RAM, a 16-kbit FIFO with built-in control logic, or a 128×48-bit ternary CAM (content-addressable memory). Because you can preload all of the memory structures at power-up, you can alternatively use them as ROMs.

One of the four ispXPLD I/O banks does double duty as the system-configuration port, and each device, independently of logic and memory density, includes two system-clock PLLs. IspXPLD chips run at 1.8, 2.5, or 3.3V with no speed penalty at lower voltages. In addition to the programmable analog devices, Lattice offers a series of SPLDs and a logic-deficient—compared with CPLDs—but routing-rich programmable-interconnect architecture called ispGDX. For you nonbinary fans, the company also sells ispPAC programmable analog chips.

Lattice's PC design software, ispLever, comes in numerous variants. For workstations, you need to obtain ispExpert Compiler for Lattice-developed architectures and DesignDirect Summit for the Mach4A and SPLDs. In line with its heritage as the in-system-programming pioneer, Lattice also supplies a variety of programming software, source code, and hardware to help you get your design up and running.

AT A GLANCE: LATTICE

- In-system programming remains Lattice's key corporate-marketing motto.
- Mach 4A and ispMACH 4000 follow-ons overlap SuperFast and SuperWide CPLDs with a timing twist.
- SPLDs, programmable switch devices, and programmable analog devices complete the product-line plethora.

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 coupled an 8-bit microcontroller to a conventional memory. From those humble beginnings, the product line has expanded in numerous directions. For example, the company now offers in-system programmable and reprogrammable flash memory, including a separate boot array, instead of EPROM.

AT A GLANCE: STMICROELECTRONICS

- Memory-plus-logic hybrid chips began with the vision of preserving microcontroller ports and expanding code memory.
- Latest variants combine multiple memory types and increase programmable-logic versatility.
- Development tools and support for 8- and 16-bit processors speed your time to market.

PSDs also now combine nonvolatile code, nonvolatile data (EEPROM), and volatile data (SRAM) partitions within a single chip. The integrated logic has become more flexible and useful for general-purpose functions, not just port re-creation; the same part, after design-specific configuration, can bolt to numerous 8- and 16-bit embedded controller and microprocessor buses. You can program PSDs either offboard using a PROM programmer, onboard under system-processor control, or via a JTAG interface.

Lithography reductions have boosted read and write speeds and 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. For all its PSD devices, STMicroelectronics supplies logic-design tools, device programmers, and evaluation boards.

XILINX

JUST AS ALTERA is a CPLD vendor that later added FPGAs to its product portfolio, Xilinx is an FPGA supplier that has subsequently bought and developed several CPLD families. Available in 2.5, 3.3, and 5V versions, XC9500 devices have 36 to 288 macrocells and incorporate a fairly mainstream global-routing architecture. One of the parts' most notable characteristics is their flash-memory architecture, which the vendor claims results in lower cost and higher reprogramming-cycle capabilities than does the more common EEPROM approach. (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. The latest iteration of the CoolRunner family, Coolrunner II, migrates previous-generation XLPA3 to a 1.8V, 0.18-micron lithography. In the process, it tackles one key limitation of XPLA3: the chips' slower performance—albeit with much lower power consumption—than sense-amp-based alternatives from Xilinx and other suppliers.

CoolRunner II propagation delays are as low as 3.5 nsec for the 32-macrocell version and 6 nsec for the 512-macrocell part, coupled with 300-MHz clock frequencies. All family members offer 2× clock multipliers at each macrocell; 128-macrocell and larger CoolRunner-II devices also include a 2 to 16× on-chip clock divider. CoolRunner-II I/O-buffer-voltage options are 1.5, 1.8, 2.5, and 3.3V; 32- and 64-macrocell parts contain one I/O bank. The 128- and 256-macrocell versions offer dual I/O banks that can run at independent voltages, and the 384- and 512-macrocell version supplies four I/O banks.

The company has two interesting—and free—software products. The Internet-based WebFitter tool accepts designs in VHDL, Verilog, ABEL, EDIF, and XNF

AT A GLANCE: XILINX

- Xilinx's architectures slowly but surely make inroads into CPLD applications.
- Acquired, then enhanced, CPLDs keep their cool and strive for a perfect fit.
- Hit the Web for free software and other stuff.

(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 the CoolRunner CPLD product lines as well as several FPGA families. Xilinx's Alliance and Foundation Integrated Software Environment tool sets also support XC9500 and CoolRunner-II devices.

EMBEDDED-PROGRAMMABLE-LOGIC CORES

ACTEL

Actel has chosen a conventional logic-cell approach, comprising dual three-input LUTs (look-up tables) feeding a register, for the VariCore embedded-FPGA architecture it acquired when it purchased Prosys Technology. 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: no hardware multipliers or other arithmetic-optimized structures, no three-state buffers or bidirectional buses, and no 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 FG (functional group) contains four pairs of three-input LUTs and four registers. The registers share common control inputs, including enable, preset, reset, and clock.

The next level of hierarchy is the PEG (programmable-embedded-gate) array, an 8×8 matrix of FGs 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 PEG structures on 0.18-micron process technologies; The **table** on embedded-programmable-logic cores in the Web version of this article at www.edn.com also discloses the company's plans for 0.13-micron processors. You can orient these multi-PEG arrangements in square, rectangular, and even L-shaped structures.

AT A GLANCE: ACTEL

- 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.

According to Actel, the ideal VariCore shape is square, which 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 multiple 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, whereas most of the design resides in ASIC gates.

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. In fact, 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 FGs. Each multi-PEG array also contains JTAG circuitry and a BIST (built-in-self-test) 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-intellectual-property form.

ATMEL

With the sale of Agere Systems' product line to Lattice, Atmel is now unique as a provider of both FPGAs and ASICs. The company not only produces FPGAs and

AT A GLANCE: ATMEL

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

FPGA-plus-ASIC hybrid chips of its own, but also supplies embedded programmable-logic cores for ASIC designs. The architecture allows for a range of logic-cell counts and associated RAM-array sizes, a reflection of the fine-grained AT40K FPGA architecture for which the company originally envisioned embedding.

LEOPARD LOGIC

Since the publication of last year's directory, Adaptive Silicon and partner LSI Logic terminated their embedded-FPGA program. At least part of the reason for the demise of the MSA product line centered on the companies' choice of an ALU as the fundamental logic-cell building block, versus the more design-tool-friendly LUT (look-up table). Unwilling to make that same mistake, Leopard Logic bases its HyperLink CCL core on a traditional combination of four-input LUT and dual registers.

Leopard Logic focuses its innovation efforts on the cores' routing resources. It bases point-to-point wiring and hierarchical interconnect on multiplexers versus the conventional pass-gate approach. According to the company, user benefits include extremely high efficiency at full speed, low power consumption, and low configuration-memory requirements. Standard array sizes employ a 1-to-1 aspect ratio; other shapes are possible.

System-on-chip integration hooks include a command-driven host interface for an embedded CPU or JTAG, built-in self-test with a single command, and debugger access to all registers in the core. The design and simulation tools simplify integration with an ASIC flow; LEF (library-exchange-format) models for floorplanning, GDSII files for tape-out, and both Verilog and VHDL synthesis and simulation libraries with timing back-annotation in standard SDF.

AT A GLANCE: LEOPARD LOGIC

- Embedded-FPGA newcomer learns from competitors' mistakes.
- Multiplexer-based routing promises higher efficiency and low power consumption.
- Test chips and design kits enable evaluation before implementation.

To try before you buy, Leopard Logic offers 256-, 1024- and 4096-LUT test chips built on TSMC's 0.18-micron process and mounted on evaluation boards. The company's design kits support TSMC's 0.13-micron technology.

References

1. Dipert, Brian, "*EDN*'s second annual programmable-logic directory," *EDN*, Aug 30, 2001, pg 36.
2. Dipert, Brian, "*EDN*'s first annual programmable-logic directory," *EDN*, Aug 17, 2000, pg 54.
3. Dipert, Brian, "Tools help you lose the hardware blues," *EDN*, April 18, 2002, pg 36.
4. Dipert, Brian, "Do combo chips compute (or even compile)?" *EDN*, Feb 15, 2001, pg 101.
5. Dipert, Brian, "Cunning circuits confound crooks," *EDN*, Oct 12, 2000, pg 103.
6. Dipert, Brian, "Low-cost programmable logic: How low should you go?" *EDN*, March 16, 2000, pg 123.
7. Dipert, Brian, "The best (or worst?) of both worlds," *EDN*, Nov 11, 1999, pg 139.
8. Dipert, Brian, "Figuring out reconfigurable logic," *EDN*, Aug 5, 1999, pg 103.
9. Dipert, Brian, "Lies, damn lies, and benchmarks: The race for the truth is on," *EDN*, May 27, 1999, pg 54.
10. Dipert, Brian, "Synthesis shoot-out at the *EDN* corral," *EDN*, Sept 11, 1998, pg 95.
11. Dipert, Brian, "Counting on gate counts? Don't count on it," *EDN*, Aug 3, 1998, pg 52.
12. Dipert, Brian, "Moving beyond programmable logic: if, when, how?" *EDN*, Nov 20, 1997, pg 77.
13. Dipert, Brian, "Programmable logic: Beat the heat on power consumption," *EDN*, Aug 1, 1997, pg 57.
14. Dipert, Brian, "Shattering the programmable-logic speed barrier," *EDN*, May 22, 1997, pg 36.

Author's biography

You can reach Technical Editor Brian Dipert at 1-916-454-5242, fax 1-916-454-5101, e-mail bdipert@edn.com.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to www.edn.com/info and enter the reader service number. 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. 301

Altera
1-408-544-7000
www.altera.com
Enter No. 302

Anachip
1-408-434-0678
www.anachip.com
Enter No. 303

Atmel
1-408-441-0311
www.atmel.com
Enter No. 304

Cypress Semiconductor
1-408-943-2600
www.cypress.com
Enter No. 305

Lattice Semiconductor
1-503-268-8000
www.latticesemi.com
Enter No. 306

Leopard Logic
1-408-777-0905
www.leopardlogic.com
Enter No. 307

QuickLogic
1-408-990-4000
www.quicklogic.com
Enter No. 308

STMicroelectronics
1-781-861-2650
www.st.com
Enter No. 309

Triscend
1-650-968-8668
www.triscend.com
Enter No. 310

Xilinx
1-408-559-7778
www.xilinx.com
Enter No. 311

Other companies mentioned in this article

Agere Systems
www.agere.com

ARM
www.arm.com

AT&T
www.att.com

Cadence
www.cadence.com

Chartered Semiconductor
www.csminc.com

IBM
www.ibm.com

LSI Logic
www.lsil.com

Lucent Technologies
www.lucent.com

MathWorks
www.mathworks.com

MIPS Technologies
www.mips.com

Philips
www.philips.com

TSMC
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, enter no. 312 at www.edn.com/info.