PROGRAMMABLE-LOGIC DIRECTORY

coverstory *By Brian Dipert, Technical Editor*

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 that both customer needs and competitive pressures dictate.

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

Compared with ASICs, programmable-logic devices have shorter leadtimes, lower upfront NRE (nonrecurring-engineering) 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 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 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 ar-

EDN's FOURTH ANNUAL PROGRAMMABLE-LOGIC DIRECTORY HIGHLIGHTS THE ARCHITECTURES AVAILABLE FOR YOUR NEXT DESIGN. FIND OUT WHAT'S NEW, WHAT'S OBSOLETE, AND WHAT'S EVOLVED IN PALs, PLDs, FPGAs, AND ASIC/FPGA HYBRIDS.

chitecture 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 "hard" analog and digital circuitry (ranging from multipliers, through high-speed transceivers, and all the way up to CPU cores with associated peripherals), 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, most designers have turned from traditional low-level state-machine and schematic-entry synthesis to high-level languages, such as VHDL and Verilog, and some are even exploring 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. Third-party IP start-ups have achieved at best mixed success; most of the IP you purchase or freely obtain probably comes from your silicon and design-software vendors.

The months subsequent to the publication of *EDN*'s last programmable-logic directory, in Sept 2002, were filled with industry turmoil. Several vendors exited the market, because they were either acquired or ran out of money, and at least one company reinvented itself after its original business plan failed to bear fruit. For the vendors that remained, their product lines in the best case stagnated or in the worst—and more common—case were drastically whittled back. Late 2003 and early 2004, however, have witnessed a resurgence of demand for programmable-logic devices and, as a result, a resurrection of product and technology roll-outs led by market leaders Altera and Xilinx. Commensurate with this resurgence, *EDN* is publishing its fourth edition of the programmable-logic directory, and I welcome your feedback on how it can better address your needs in future iterations.

For more information on the FPGAs, PALs and PLDs, and ASIC/FPGA hybrids mentioned here, please see **tables 1**, **2**, and **3**, respectively, at the Web version of this article at www.edn.com.

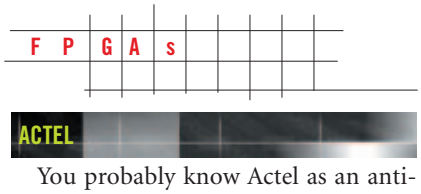

You probably know Actel as an antifuse-based-FPGA supplier, but the company is broadening its focus to include another programmable-logic technology. Five main device families form the foundation of today's Actel antifuse-product line: the 5V MX; 3.3V SX; and closely related 2.5V SX-A, 2.5V eX, and 1.5V Axcelerator. Actel also offers the RTSX-S and RTAX-S antifuse-based product lines, tailored for aerospace and military applications and respectively derived from the SX-A and Axcelerator architectures, and the flash-memory-based ProASIC and ProASIC Plus FPGA families.

Compared with other configuration technologies, antifuse delivers lower impedance and, therefore, lower power and higher speed signal interconnection and more robust immunity to high-radiation operating environments. The diminutive size of an antifuse configuration element, compared with a flash- or SRAM-based alternative, makes efficient use of silicon real estate. 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 lesser 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.

- 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 is on its second product generation.

The company's first-generation PLICE (programmable low-impedance circuit element) antifuse-technology approach, used in 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- 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 experienced with ASICs and their design tools. ProASIC, like Axcelerator, brings embedded-SRAM and FIFO capability to Actel's arsenal. The Plus parts, which Actel built on more advanced processes, also offer a higher memory-to-logic-resource proportion.

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 and FS2 CLAM (configurable-logic-analyzer module) verification and logic-analysis tools enable you to observe and analyze internal device nodes. The Silicon Sculptor device programmer supports programming 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 from Actel and strategic partners.

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.

From the Flex 8000 and follow-on Flex 10K architectural foundations, Altera has taken its FPGA product line in multiple directions, all of which it until recently based on a common, essentially unchanged LAB (logic-array-block) structure. The Acex 1K family is a cost-optimized Flex 10KE variant, which Altera 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 bigger siblings.

Historically, Altera's FPGAs have employed a long-line routing-dominated architecture that provides timing predictability. And, 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 Stratix family adopts a three-length mix more significantly shifted toward short routes with no lines spanning the die as in the past. Stratix also provides three onboard-memory structures, each with application-optimized array numbers, densities, and bus widths. Combine Stratix with as many as 20 channels of high-speed transceiver circuitry and functionally enhance and boost the speed of the source-synchronous I/O buffers, and you have the Stratix GX family, supporting serial data rates as high as 3.125 Gbps.

Whereas Acex 1K was a comparatively minor tweaking of the Flex 10K architecture, Cyclone is a clean-slate FPGA revamp with low cost as its dominant objective. Cyclone contains only short global-routing lines, and Altera builds it on a detuned and one-metal-layer-reduced variant of Stratix's 0.13-micron copper process. To synthesis tools, the Cyclone logic-cell structure appears similar to the cell of other Altera FPGAs, but internal optimizations have made it 30% smaller than the Stratix logic cell. Cyclone comes with a lower memory-to-logic ratio than Stratix, the embedded memory runs at lower speeds, and Altera has removed support for several specialty-memory features. Cyclone FPGAs also come with a maximum of two PLLs, again with fewer features than their Stratix counterparts. The I/O buffers focus on single-ended protocols and lack Stratix's integrated termination resistors, though they do comprehend some differential-signaling schemes and support programmable slew rate and drive strength.

Stratix II, on the other hand, takes Stratix features to even higher levels, commensurate with its high-density eminence and based on the ALM (adaptive-logic module), a revamped fundamental logic element. Peer inside a Stratix II ALM and you'll find a collection of logic functions: two four-input LUTs (look-up tables), a three-input LUT, two registers, and two three-input adders. The triple-LUT structure can implement a diversity of functions, and Altera has doubled the number of inputs to and outputs from each ALM compared with the LE predecessor. All Stratix II chips, built on a 90-nm, low-K copper process, also support nonvolatile 128-bit AES (Advanced Encryption Standard) security for their links to external configuration memories, along with 1-Gbps differential source-synchronous signaling and dynamic phase alignment.

Altera's "hard"-CPU-core-inclusive Excalibur XA-hybrid chips are on indeterminate hiatus after receiving underwhelming market acceptance, and Altera is instead placing its bets on the optimized 8-, 16-, and 32-bit Nios and Nios II RISC processors, which aim to reside in FPGA logic as a soft core instead of in ASIC gates. Over time, Altera plans to port its Nios family, along with other internally developed and partner-developed cores, to all of its FPGA architectures. The company's cost-reduction vehicle for high-volume designs, the

HardCopy program, involves the company's creation of mask-programmable FPGA variants; Altera has recently expanded its focus to address the burgeoning structured-ASIC market. The same Max+ software you'd use to design with Altera's CPLDs also supports Acex 1K and Flex FPGAs. For Apex, Cyclone, and Stratix devices, you'll want to fire up Altera's more advanced Quartus II development-tool environment, along with the SignalTap logic analyzer.

ATMEL

- Second-generation FPGAs have multiplication in mind.
- An 8-bit AVR controller makes a system on chip a reality.

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, though, and the 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. A dedicated two-input AND gate resides ahead of the LUT. Atmel combines this diagonal routing with this logic improvement to enable AT40K FPGAs to better support the matrix-multiplication operation that is 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, a 128-bit dedicated RAM array appears at the point at which four logic-block clusters, each containing 16 logic blocks, intersect. 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.

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. The Secure FPSLIC contains two data buses—one going to the system and the other running only within the package between the memory and the AT49K. Its 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 bit-stream trace between memory and the 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 is good enough to foil the efforts of would-be thieves.

LATTICE SEMICONDUCTOR

- Orca logic-structure innovations require synthesis support or core-generator intercession.
- The internally developed ispXPGA offers multimemory integration and synthesis-friendly logic.

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 in the same way that FPGA pioneer Xilinx has expanded into CPLDs. This broadening has come in the form of acquisition, 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.

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 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 a decoder as large as 10 bits 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. Each Series 3 PFU is more than two times bigger than that of Series 2 devices, now including eight four-input LUTs and nine registers.

Series 4 devices focus attention on signal routing, based upon the 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/data-recovery) macro. The ORT42G5 and ORT82G5 are 1.25-, 2.5-, or 3.125-Gbps backplane-interface FPSCs (field-programmable system chips), and the ORSO42G5 and ORSO82G5 are intended for high-speed serial SONET data transmission. The 10-Gbps ORLI10G line-interface device and the ORSP14 embed two SPI4.2 cores, a 3.7-Gbps SERDES (serializer/deserializer), and a high-speed memory controller.

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 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 256×18 bits, support FIFO and single- and dual-port-memory structures with 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 SERDES 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 IP (intellectual-property) cores.

QUICKLOGIC

Antifuse advocate QuickLogic has, since its earliest pASIC 1 architecture, employed the ViaLink, a metal-to-metal antifuse technology, above the logic grid. Actel only later began to match ViaLink's capabilities. QuickLogic's chips have all of the inherent antifuse advantages as Actel's chips, 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, some of which, unfortunately, have fallen by the wayside.

- Logic cells promote design flexibility.
- ESPs (embedded standard products) integrate RAM and system-interface modules.
- ASIC-housed DSP and PCI functions are now available, along with a MIPS processor and peripheral suite.

The company builds its pASIC 3 product family on a smaller lithography than the original pASIC 1 and follow-on pASIC 2 technologies, and the pASIC 3 devices have correspondingly higher logic counts, lower operating voltage, 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 2.5V 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 and QuickPCI 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.

QuickLogic turned to the Eclipse architecture to implement 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 multiplication, 16-bit addition, or accumulate with carry), and multiple passes through the ECU support the common multiply-accumulate function. Eclipse II is a 1.8V, 0.18-micron descendent of EclipsePlus, which includes ECUs only in the largest family members. And QuickMIPS fills the hole that Altera created when it suspended its MIPS-based Excalibur program. A MIPS32 4Kc processor combines with a separate arithmetic unit and 16 kbytes each of instruction and data caches and 32-bit Advanced High Performance and Advanced Peripheral Buses. Peripherals include dual 10/100 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.

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.

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. ASIC-replacement 3.3V Spartan-XL parts trace their lineage back

to the XC4000E. Spartan forgoes the XC4000E's parallel-interface-configuration option and dedicated, on-chip, wired-AND decoding, and the devices come in only low-cost plastic packaging.

Compared with the baby steps that preceded it, the XC4000-to-Virtex transition was a giant architecture leap forward for the company's high-density-product line. Xilinx found that synthesis tools rarely took advantage of XC4000's 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 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 the more-than-300-Mbps-per-pin, chip-to-chip, buffered, DDR SelectLink communication protocol. 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 for 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.

- Increased LUT (look-up-table) 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.

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 external reference resistors you supply (two resistors for each of the device's eight I/O blocks).

Xilinx embraces the single-chip, ASIC-plus-FPGA concept, which competitor Altera has turned its back on, with Virtex-II Pro, a product family containing high-speed serial transceivers (supporting bandwidth as high as 10.3125 Gbps in the ProX proliferation) and zero to four integrated IBM PowerPC CPU cores. Xilinx's early-2004 acquisition of Triscend meant the end of that company's 8051- and ARM-based hybrid-chip evangelization but promises to accelerate the development of future PowerPC-based products. Spartan-3, the company's latest cost-optimized product family, is based on Virtex-II but makes several silicon-slimming alterations: It's the industry's first FPGA manufactured on a 90-nm-lithography process, a staggered-I/O-pad arrangement minimizes die size for a given package pin count, and only every *other* logic slice contains LUTs you can use as distributed RAM.

In 2003, Xilinx provided a sneak preview of its next-generation, high-density ASMBL(application-specific-modular-block) architecture, with corresponding Virtex-4 products scheduled to appear in autumn 2004. (No, I don't know what happened to Virtex-3, either.) ASMBL arranges multiple stripes' worth of logic, memory, and other circuitry across the die and will ease the development of future ASIC-plus-FPGA devices. 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 Easy-Path, which evaluates parts with a customer-design-specific subset of the normal testing flow. 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 Wind-River-based software products that enable chip programming and reprogramming at the design bench, on the manufacturing line, and in the field.

PALs AND PLDs

ALTERA

- Over the years, Altera has performed only minor fine-tuning on its Max 7000 flagship CPLD-product line.
- Relaxed and restricted testing, along with plastic packaging, lead to low-cost Altera devices.

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 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 counts. Along the way, Altera has in numerous ways fine-tuned the Max 7000. Perhaps the most significant change 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 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.

The company tests its cost-optimized Max 3000A parts to more relaxed specifications than those for Max 7000 devices, allowing the Max 3000A devices to have higher power consumption. Altera offers the parts with fewer packaging options and doesn't support all of the Max 7000 features. 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 reminiscent of a 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, though, is that pin-to-pin timing and logic-block-to-block timing depend on placement and are, therefore, unpredictable.

Altera's Max II family takes FPGA similarity to the next level; it *is* an FPGA, although Altera goes to great lengths to promote it as a CPLD-product line, explaining why it's in the PLD section of this directory. Granted, Altera has tweaked the device's routing and made other enhancements intended to improve Max II chips' timing predictability compared with that of conventional FPGAs. But, it's a stretch to equate Max II with the highly predictable, product-term-based and global-routing-matrix-based CPLD that most of us think of when we see the abbreviation. Based on the flash-memory-inclusive variant of TSMC's 0.18-micron process, Max II chips embed their configuration memories and are fully functional less than 200 μsec after the application of stable system power. You can run Max II parts at 1.8 to 3.3V supply voltages; an internal regulator downconverts the externally supplied voltage to 1.8V, and 1.8V-only device variants are also available that bypass and turn off the regulator, thereby reducing standby-power consumption. Other intriguing Max II features include a separate dual-block, 8-kbit, user-accessible flash-memory partition and the ability to in-system reprogram the nonvolatile configuration memory while the logic in the SRAM-based portions of the chip is running.

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.

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 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 (Advanced Boolean Equation Language)-, 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 penalty for using the full 16 product

terms per macrocell and no delays for fan-outs, expanders, or I/O pins versus dedicated pins or product-term steering or sharing. Also, 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 for device-configuration memory already exists, 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.

Outside each logic-block cluster and closely connected to multiple sets of

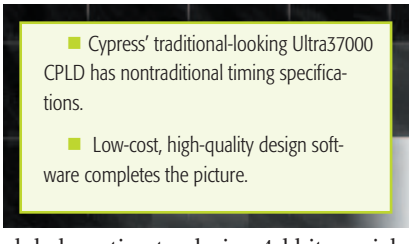

- Cypress' traditional-looking Ultra37000 CPLD has nontraditional timing specifications.
- Low-cost, high-quality design software completes the picture.

global-routing tracks is a 4-kbit specialty SRAM. 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 as fast as 266 MHz) or divide (down to as little as one-sixteenth) 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.

The company is proud of not only its silicon, but also 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 and dynamic-timing simulator, fitter software, and reference documentation. You can order Warp from 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.

LATTICE SEMICONDUCTOR

Lattice first differentiated itself from the competition by bringing in-system programming to PALs. Before that, 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. Lattice's ispMach 4000 SuperFast parts, architecturally based on the ispMach 4A device family that Lattice adopted when it took over Vantis, range from 32 to 512 macrocells in size and come in 1.8V (with a "zero-power," 20-µA standby-power spin), 2.5, and 3.3V variants. They offer propagation delays as fast as 2.5 nsec, even if your design uses as many as 20 product terms per macrocell—a feature the company calls Speed-Locking.

The 3.3V ispLSI and 2.5V ispMach 5000 SuperWide families overlap Mach 4A and ispMach 4000 at the latter families' higher macrocell counts. SuperWide devices, as their name implies, 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 lower than that of some other CPLD architectures. The ispMACH 5000VG devices come in 768- and 1024-macrocell variants, courtesy of a hierarchical-routing architecture, and span the SuperWide and SuperBig categories.

The newest SuperBig architecture, ispXPLD, is SRAM-based, a trait it shares with Altera's Max II and Cypress' Delta39K 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 in-

puts and 32 macrocells. Alternatively, the multifunction block can construct an 8-kbit, dual-port RAM; 16-kbit, single-port or pseudo-dual-port RAM, including one read/write port and one read-only port; a 16-kbit FIFO with built-in control logic; or 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 programmable-analog devices, Lattice offers SPLDs and a logic-deficient (compared with CPLDs) but routing-rich programmable-interconnect architecture called ispGDX. For you nonbinary fans, the company also sells programmable analog chips called ispPAC.

- In-system programming remains Lattice's key corporate-marketing motto.
- The company's product line subdivides into SuperFast, SuperWide, and SuperBig groupings, and some product lines span multiple categories.

Lattice's PC design software, ispLever, comes in numerous variants, supports PCs and workstations, and, in its latest iterations, encompasses the entire programmable-logic line. For ispPAC programmable analog devices, you'll want to fire up the company's PAC Designer software. 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 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 you lose when you couple an 8-bit microcontroller with 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.

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

Lithography shrinkages have both boosted read and write speeds and reduced operating voltages and currents. From its PSD foundation, STMicroelectronics has moved the product line in several vectors. DSM (DSP-system-memory) devices tailor their functions and external interfaces for DSPs. And, instead of connecting to a microcontroller, the µPSD and Turbo µPDS chips bring the microcontroller core inside, along with a full-featured peripheral set sufficient to create a low-end system on chip. For all its devices, STMicroelectronics supplies logic-design tools, device programmers, and evaluation boards.

- Memory-plus-logic hybrid chips began with the vision of preserving microcontroller ports and expanding code memory.
- Subsequent variants combined multiple memory types and increased programmable-logic versatility.

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-product 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 that Xilinx based the devices on flash memory, not EEPROM, which, the vendor claims, results in lower cost and higher reprogramming-cycle capabilities. (Most applications have little need for this reprogram-

ming-cycle feature, however.)

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's

- Xilinx's architectures slowly but surely make inroads into CPLD applications.
- Hit the Web for free software and other stuff.

Delta39K devices. The latest iteration of the CoolRunner family, CoolRunner II, migrates earlier 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. 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 versions supply 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 (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 (integrated software environment) supports both the XC9500 and the CoolRunner CPLD-product lines, as well as several FPGA families.

ASIC/FPGA HYBRIDS

LEOPARD LOGIC

- Silicon newcomer learns from both competitors' and its own past mistakes.
- A common design-and-simulation flow simplifies the development of ASIC- and FPGA-housed circuits.

Past editions of this directory labeled this section "cores." However, Actel terminated its embedded-FPGA program after the publication of the 2002 directory. Adaptive Silicon had gone out of business roughly one year earlier. Atmel,

it turns out, offers no embedded-FPGA cores in its ASIC program; it offers only standard products with embedded-FPGA arrays. And, although in its initial incarnation, Leopard Logic provided embedded FPGA-IP (intellectual-property) building blocks to ASIC and foundry suppliers, the company obtained insufficient industry interest. So, Leopard Logic reinvented itself as a silicon supplier, thereby competing with its former potential customers.

This directory so far rather loosely uses the term "ASIC plus programmable logic" to describe FPGAs and CPLDs with elaborate and fixed "hard" functions, such as CPU and associated peripheral cores, diffused on the die—that is, not located in the user-programmable fabric. In Leopard Logic's case, however, both the ASIC and the FPGA portions of the chip are generic and therefore can work in tandem to implement your design. Leopard Logic's Gladiator devices combine HyperBlox MP (mask-programmable) and FP (field-programmable) logic modules, DLLs and PLLs, dual-port SRAM arrays, and MAC (multiply-accumulate) circuits on a single chip, encircled by a configurable-I/O ring.

FP and MP logic blocks are functionally identical, simplifying EDA tools' efforts, but, whereas SRAM elements configure the FPs when the chips are in your hands, a single via layer customizes the approximately 20-times smaller MP blocks during the final stages of device fabrication. Robust, hierarchical interconnect structures tie together the FP and MP subsystems and tie them to the remaining circuits on the device. Leopard Logic estimates that it can turn around product samples in less than four weeks after you ship the company $50,000 in NRE charges and a compiled bit stream representing your design.□

References

You can find the references to this article on the Web version at www.edn.com.

Author's info

Reach Technical Editor Brian Dipert at 1-916-454-5242; bdipert@edn.com.

FOR MORE INFORMATION...

Actel, www.actel.com
Altera, www.altera.com
Atmel, www.atmel.com
Cypress Semiconductor, www.cypress.com
Lattice Semiconductor, www.latticesemi.com
Leopard Logic, www.leopardlogic.com
QuickLogic, www.quicklogic.com
STMicroelectronics, www.st.com
Xilinx, www.xilinx.com

Other companies mentioned in this article

Agere Systems, www.agere.com
ARM, www.arm.com
AT&T, www.att.com
IBM, www.ibm.com
Lucent Technologies, www.lucent.com
MIPS Technologies, www.mips.com
Philips, www.philips.com
TSMC, www.tsmc.com