



Structured ASICs: More gain, less pain?

**STRUCTURED ASICs GIVE
YOU LEADING-EDGE
PERFORMANCE AT A
MANAGEABLE COST.**

At a glance21
A swing of the pendulum21
What's next?22
For more information28

IT MATTERS NOT WHICH ANALYST'S FIGURES you look at. Dataquest, i-Supply, and others all agree: Design starts in complex ASICs are down. Complex reasons drive this situation, only a few of which relate to the current state of the semiconductor industry. Today's leading-edge semiconductor processes offer specifications

that are attractive in performance (speed), density, and—if you are in a position to take advantage of true high-volume production—price. However, these gains come at a cost that includes high upfront expenses and large minimum-order quantities.

If you are involved in a design for, say, an industrial product, a medical project, or even a consumer product that will not be in the very-high-volume-production category, you may be looking enviously at the performance you could achieve with an ASIC in a leading-edge technology. But, for a variety of reasons, this technology may increasingly become out of reach.

However, the recently introduced category of structured ASICs may offer a way forward, giving you access to advanced technologies, in moderate volumes, with affordable design-cycle costs—and without compromising performance or versatility. The suppliers entering this market use a variety of terms—“structured ASICs,” “platform ASICs,” and others—but this article refers to them as structured ASICs. In some quarters, structured ASICs have been dismissively described as the “return of the gate array,” but they are much more. Before looking at them, it is worthwhile to review the scale of the developing

problem with “conventional” standard-cell products.

0.13-MICRON COT IS TOUGH

Today’s leading-edge technologies offer performance with clock speeds into the hundreds of megahertz, low-voltage (1.8/1.5V) and low-power operation, and very high densities (number of logic gates per unit area)—hence, small dies. But designing in a 0.13-micron (and smaller) technology is a challenging and expensive business. A look at the past several years’ topics at the EDA industry’s annual DAC (Design Automation Conference) shows the technological barriers that the tool vendors have been struggling to surmount. Aside from the sheer challenge of getting the logic design right and achieving the timing “closure” in an environment in which the interconnect dominates the on-chip delays, placement and routing of the cells must take into account a number of “new” phenomena that become significant only at these geometries. You must analyse power grids to see whether, at the grid’s extremities, the current drawn through fine metal lines is such that the IR drop in actual voltage supplied to a cell reduces the operating voltage below specification. You also need to analyse current densities and electric fields to determine

AT A GLANCE

▷ Structured ASICs occupy the space between FPGAs and standard-cell ASICs, though most structured-ASIC designs offer close-to-standard-cell speeds.

▷ Structured ASICs provide much higher logic density than FPGAs but somewhat less logic than standard-cell designs.

▷ Mask charges and other NRE costs for structured-ASIC designs are dramatically less than those for standard-cell designs.

▷ The structured-ASIC market is still being defined; no two offerings are directly competitive, so look for one with the features that best match your application.

▷ Structured ASICs return the design problems of deep-submicron silicon to the domain of the chip vendor.

▷ The most demanding and highest density designs will still have to wrestle with the complexities of leading-edge standard-cell technology.

whether they are likely to impact reliability by causing conductors to act as very slow fuses. And, you must subject paral-

lel signal paths to a field analysis to evaluate signal-integrity issues. All of these steps and checks are the domain of state-of-the-art EDA tools, which add to the cost of the exercise.

The headline figure that tool suppliers and others often quote for building a leading-edge standard-cell device is the cost of the mask set—the 20-plus precision patterns needed to fabricate all layers of the device. Quoted figures for 0.13-micron or soon-to-come 90-micron technologies are \$750,000 to \$1 million. And if a respin (design revision) is necessary, you again incur that cost. EDA-tool vendors frequently cite mask costs to persuade you that using their products will obviate the need for a respin. In fact, if you are in the business of handling the complete design process for a product in that deep-submicron league, the mask set is the least of your problems; at least you know who is going to make it and that it will work. You also need to put together a functional design flow, patching in leading-edge point tools to solve the special problems that your design reveals. You will need to pay for the licences for all those tools, and you will need a highly skilled design team to operate them (a scarce commodity in itself). Then, there is the question of minimum-order quantities. Suppose you end up with a mod-

A SWING OF THE PENDULUM

There is a historical perspective to the development of structured ASICs, and it may show a cyclical process at work. Early in the life of the ASIC, it was very much a bespoke business that evolved from the design of the first custom-silicon products. In the 1980s, a move began to wrest control of the design process away from the silicon vendors; if you could carry the design right through to geometry yourself, you could take it to any silicon house running a compatible process and be free of the restrictions that captive customers face. The silicon foundry emerged, and the catch phrase became “foundry independence.”

With the move to submicron

design and the shift from gate-delay-dominated timing to interconnect-dominated timing, the centre of gravity of the design process shifted back to the silicon vendors, as they wrestled with the increasingly complex layout problems of that era. In time, the tools caught up, and vendors once again offered the appetising prospect of letting you execute all but the final stages of your design in your own facility, passing only the later stages of the exercise to the silicon house. This time around, it was called COT (customer-owned-tooling), but the principle was the same.

You can argue at length about whether the independence you could secure from dependence

on a particular supplier was ever more than illusory, and ASIC-industry commentators do.

Today’s structured-ASIC suppliers will claim that, in practice, you have as much freedom as you ever did, because you are completing a design to a recognised format (RTL, HDL, or netlist) with the objective of handing it off to a formulaic design-completion process that itself has a short turnaround time. If, therefore, you fall out with one supplier, submitting the design to another’s process should require no more time than it takes to port a standard-cell design from one process to another.

Synplicity’s director of ASIC marketing, John Gallagher, is an enthusiastic proponent of the

technology and not, he says, only because his company’s tool appears in the design flow of several vendors: “Gate arrays were about manufacturing time; structured ASICs are all about design time.” You can cut the time to carry out a design by more than 50%, he says, and the cost by an order of magnitude. Gallagher notes that the two key software tools specific to the structured-ASIC-design path are affordable on a six-month licence, even budgeted against a single project. He claims that focussed synthesis can also restore the performance margins you lose by using the array architecture and get your device’s performance nearly back to that of a standard-cell architecture.



est 1-cm-square die; not many wafer fabs are yet building this type of product in 300-mm facilities, but such capability will soon become the norm. These production facilities offer more than 500-plus dies from each wafer; also, they don't really want to discuss processing less than half a "boat" (12 wafers). And they would rather deal in complete boatloads. One batch could be the lifetime requirement of many ASIC designs.

You have, of course, an alternative: FPGAs. But the step from the largest FPGA to a standard-cell ASIC remains a significant and well-documented one. The FPGA is less area-efficient by more than an order of magnitude, and the largest

devices are very expensive. They're suited to prototyping and short production runs but prohibitive for volume production runs. Of course, the economic crossover point between the technologies is a moving target as the FPGAs get larger and their immediate predecessors progress down the cost curve. A differential factor in performance also exists, but it, too, has narrowed as FPGA vendors have switched production to the latest and smallest geometry silicon processes.

RETURN TO A SEA OF LOGIC

Structured ASICs aim to take away some of the pain. They are somewhat

similar to older gate-array products but offer significant differences. Like the older products, they are prediffused arrays of logic elements that are finally interconnected and configured in the top two or three layers of metal (**Figure 1**). And as with older gate arrays, wafers—in a range of sizes—are diffused by the silicon manufacturer and held in stock. You create your logic design, which is mapped to the structures on the array and then create masks for only the two or the three layers of metal that it takes to define your logic configuration. In this way, vendors can more quickly return prototypes than they can for standard-cell devices in which all layers are custom and must be fabricated

WHAT'S NEXT?

With few structured-ASIC designs yet in volume production, it may seem a little premature to ask, "What's next?" But there is an answer in development in more than one ASIC house, and first announcements will likely come at the end of this year: a hybrid device offering both a standard cell and a structured ASIC on a single platform. You might use such a product to produce a range of silicon designs with a common set of core functions, but in a number of distinct variants—for example, a communications product that has to interface to a number of interface standards but also provide the same basic processing function to the data it handles.

You might also use such a platform to get a product into a market whose standards are still evolving. Your strategy would be to get the first product out as quickly as possible, and then produce a new variant, or even variants, as the standards evolve and mature. Therefore, you would place the core functions, which would not change after the first definition, into standard cells, and use blocks of structured ASICs to implement the parts of your circuit that would be subject to change at short notice. You

could look at this design in a number of ways: as a structured ASIC chip with your own, dedicated, embedded IP (intellectual property) or as an ASSP (application-specific standard product) with a degree of flexibility. In either view, you would have a design with which you could produce a new IC with good volume-manufacturing economics for the cost of the mask set for just the two or three metal layers you need to redefine the structured-ASIC part of the chip (plus the costs you incur to redesign that portion).

This argument should sound familiar; it is almost exactly the one that silicon vendors, including LSI Logic, put forward a year or two ago for a hybrid product that sought to build blocks of FPGA into a standard-cell ASIC platform. The concept did not succeed. So, is there any reason to think that a standard-cell/structured-ASIC hybrid would succeed where a standard cell/FPGA would not? There are some plausible explanations why it would. The standard-cell and structured-ASICs processes are much more closely matched in area efficiency, performance, and design style than was the case for standard cells and FPGAs. If you wanted to

build enough FPGAs into a standard cell/FPGA to be able to do a useful degree of reprogramming or adapting to new standards, you soon reached the point at which overall die area would grow rapidly. The relatively small FPGA gate count—compared with the bulk of your design, which is executed in standard cell—would then consume amounts of silicon comparable with the "main" part of the chip, eroding the area-efficiency argument that took you to standard cell in the first place. Then, the FPGA might not match up to the standard cell in speed. Also a potential difficulty, you had to adopt a complex design flow that incorporated both ASIC- and FPGA-design styles. Adding FPGA also complicated the silicon processing and added to the cost. With small and even medium-sized FPGAs descending the cost curve, keeping the programmable fraction of the design in a second chip looked like a more attractive option.

With a standard-cell/structured-ASIC hybrid, these factors are more closely matched. Area efficiency is closer; you can mix in a useful amount of structured ASICs without increasing die area too much. The same is true for

performance. And, you can handle both types of circuit block in essentially the same design flow. Such a hybrid product also allows the option of introducing mixed-signal functions via the placement of analogue cells—a feature that does not figure prominently into the first round of structured-ASIC products. The single-chip, "flexible ASIC" might be a real option this time around.

Interestingly, IBM has chosen not to build a structured-ASIC product. Its CCP (customisable-control-processor) series devices embed a "hard" PowerPC 405 core but use standard cell for the surrounding logic. IBM does, however, intend to go forward with the hybrid embedded-IP-plus-FPGA concept. European ASIC tactical marketing manager, Gordon Fairley, is unperturbed by the lack of success such concepts have so far demonstrated. IBM believes, he says, that geometry is the solution to this conundrum. Using 90-nm process geometries, the area penalties for introducing useful amounts of configurable logic via the FPGA route diminish enough to make the whole concept viable. Fairley anticipates a product announcement along these lines by the end of this year.

from a bare wafer. Because the base product is the same, the silicon vendor will undertake smaller production runs.

Structured ASICs differ from older gate arrays in several important respects. The granularity of the base logic cells from which the array is constructed is different. Structured-ASIC base-logic units are typically somewhat more complex than older gate arrays (which would have been a “sea” of simple logic gates or combinations of one or two gates). But they are much less complex than typical FPGA cells. Simple combinations of a handful of gates, multiplexers, and flip-flops seem typical. The density is less than that of standard-cell processes, by a factor of perhaps two or three, so the ultimate cost per device is higher—somewhere along the axis from standard cell to FPGA. Every design has a different cost point, and you need to factor in everything from design cost to ultimate volume to determine exactly where it lies. As you might expect, structured ASICs achieve lower performance than standard cells, but, with typical clock speeds of 200 MHz or better, they are sufficient for many applications.

The differences go much further. As the silicon vendors use latest generation technologies to fabricate the base wafers, they will have at their disposal multiple layers of metal interconnect. Of these layers, you get to configure only the top two or three. (In some devices, you configure as many as five layers.) Unlike older gate arrays, the “stock” structured-ASIC wafer will already contain a lot of interconnect. Power and ground, global clock routing, even (though not in every vendor’s solution) test access will already be in place. One effect of this interconnect is that, although you will be using a leading-edge technology, the base layers handle many of the attendant problems (power-grid provision, IR drop, electromigration, and the rest), and they will not impact the final interconnect design.

ARCHITECTURE CUTS DESIGN HASSLE

In most cases, the structured-ASIC vendor will deal with design issues arising from the complexities of 0.13-micron (or smaller) process geometries. In a typical design flow, you will carry out your logic design in the front-end environment of your choice, through to a verified netlist, or HDL description at the register-transfer level. The silicon vendor




Figure 1

Following the general principle in fabricating all structured ASICs, all of the functional circuitry is predefined and fabricated in the stock versions of the chips; you finalise the design to your specification with the top three mask layers (courtesy Chip Express).




Figure 2

The outline design flow for LSI Logic’s RapidWorx tool set is typical of the steps involved in structured-ASIC design. As a customer, you deal only with steps 2 to 4. Unlike full standard-cell design, the base array handles many of the complexities of deep-submicron design, or vendors deal with them in the final design stages.

will then undertake placement and routing, selecting a suitable base-die size from its range. The predefined architecture constrains your design freedom some-

what, perhaps limiting performance to some extent. However, the back-end placement-and-routing process can employ a high degree of automation, leading to a short turnaround time to first samples. The process flow will also generate test data and perform full physical-layout checks.

Most of the products make a virtue of the fact that you can design them with standard-ASIC tools that you likely already use. For most cases calling for a special tool variant, the design stage is synthesis, and the tool of choice appears to be Synplicity’s Synplify. According to John Gallagher, director of ASIC marketing at Synplicity, Synplify’s architecture is suited to being tuned to the structured array; written to be adapted to different FPGA logic cells, it directly constructs logic to make optimum use of the basic modules on the arrays.

Of the six or so structured-ASIC products now available, no two are directly comparable. Some target conversion of designs from high-end FPGAs, and others aim to capture business that would normally use standard-cell ASIC technology. Some look relatively similar to older gate arrays, in that the logic array consumes most of their area; others incorporate significant blocks of IP (intellectual property) that suit them to a particular application domain.

SWAP OUT FPGAS WHEN VOLUMES RISE

Now introducing its second generation of products under the names

XpressArray and XpressArray HD (higher density), AMI Semiconductor explicitly targets the FPGA-conversion market, continuing the business it has for several years been conducting in that area. The products aim to take a high-density FPGA design into the better volume-production economics of ASIC technology without incurring standard-cell NRE charges. They are positioned as drop-in replacements for 1.8 and 1.5V, high-end FPGAs.

XpressArray-HD is built using what AMI terms a hybrid production process. The company buys wafers from TSMC (www.tsmc.com), using that foundry's 0.18-micron process. TSMC builds the wafers up to the second level of metallisation, and then AMI adds as many as five more layers of metal at its own facility, using a more relaxed geometry of 0.35 or 0.25 microns. You can have as many as 2.5 million "ASIC" gates of logic and 1.4 Mbits of RAM, which is distributed throughout the logic array, in eight base device sizes. Phase-lock- and delay-lock-loop timing structures, with a range of commonly used I/O-interface types, ease conversion from the most popular programmable devices. XpressArray parts embed test structures, but their base layers do not include a fully connected test structure; AMI's design-completion process provides scan chains, BIST (built-in self-test), and JTAG access as part of the layout process. The company quotes maximum system-clock frequency at 220 MHz, and a range of soft IP includes blocks such as an Ethernet MAC (media-access controller) and a 64-bit PCI interface. AMI quotes power at 0.06 μ W/MHz/gate. You can mimic the pad layout of an existing design or, because the new design requires fewer power and ground pads than an FPGA, retarget it to a smaller package. Because the largest devices are denser than the largest currently available FPGAs, you can consider combining more than one programmable device into one structured-ASIC part on conversion. Not all of the designs are conversions; AMI says it also gets "pure" ASIC-style projects in these families. AMI application/system-architecture director, Bob Kirk, says that most of the company's production volumes are in the 10s-of-thousands




Figure 3 In its ISSP parts, NEC embeds IP that is dedicated to specific functional areas, such as high-speed serial interfaces for generic, fast-data applications.

area, but it will accept business down to a few thousand devices per year. Design input is from standard tool chains with synthesis by Synopsys or Synplicity. NRE charges are \$80,000 to \$200,000, and you can have samples 10 days after handing off a design.

EMBEDDED IP TARGETS COMMUNICATIONS

Also currently announcing an update to its structured-ASIC offering is LSI Logic, with its RapidChip products. RapidChip is structured around an ASIC-design flow, and it aims to deliver a reduced-NRE ASIC product. Devices will be stocked prediffused with IP cores from LSI's CoreWare programme; LSI describes the prediffused parts at the uncommitted stage as "slices." The hard-IP blocks will comprise processor cores and other functional blocks that suit a device to a target market; you then add your own, LSI's, or a third party's IP, plus custom logic. Again, you complete configuration with the last few layers of metal-mask programming. LSI will first fabricate the parts in 0.18-micron and then in 0.11-micron technology.




Figure 4 The CX5000 family from Chip Express includes an option that provides memory-hungry applications with as much as 4.5 Mbits of fast SRAM.

LSI has just announced a customised tool set for RapidChip: RapidWorx targets low-level issues—particularly those that arise with 0.11-micron processes. If you follow the default path through the tool chain, it will work in the background to automatically resolve detail design issues, concealing them as far from you as possible. A five-button tool chain gets you from RTL input to a placed netlist, the company says. The tool chain includes customised versions of Synplify and, ahead of that, Tera Systems' TeraForm RTL design-planning tool. Major components in the tool chain include RapidBuilder to configure the device at a high level, construct test strategies, and configure memories; RapidView, which lets you control the placement of major blocks, such as memories; and the Physical RTL Optimiser, which employs the TeraForm engine and the TeraGates format to generate and verify a physical view at the RTL. Synplify, which becomes Amplify in this form, maps layout directly to the RapidChip primitive cells. Other tool elements handle matters such as clock generation.

You can use the tool set in a highly automated "default-setting" mode, or you can intervene to fine-tune its processes. Using the TeraForm engine allows the tool set to generate a floorplan from physical-synthesis principles on your desktop. This process goes beyond what is possible with other RTL-linting software and goes a step further than most structured-ASIC-vendors' tools before handoff. RapidWorx, LSI says, avoids problems such as congested routing invoked by poor RTL. It also follows good design-reuse practices throughout. Likewise, Amplify embodies numerous rules to avoid crosstalk problems in placed designs. If you use RapidWorx, LSI says, you are carrying out much more of the process before handoff, which exploits the features of the architecture to reach the best cost point (Figure 2).

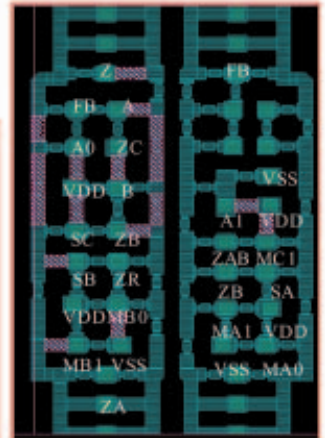

Although the design flow is ASIClike, LSI switches to a comparison with a high-end programmable device to illustrate a per-unit cost that LSI claims is as low as 10% of that of a high-end FPGA. Overall estimated NRE is 25% of a cell-based design, and, LSI adds, the outcome is more predictable. The company bases the initial selection of "slices" (now 11)

on the functions necessary for communications, storage, and consumer applications; their complexity ranges from around 3 million to 6 million gates. For example, a slice for storage applications contains an ARM 7 or 9 processor core, which can run as fast as 333 MHz, with several megabits of memory, key interfaces, and configurable I/O. You will be able to add soft IP comprising all the commonly used interface standards, plus logic unique to your own design.

Another structured-ASIC product already entering its second generation is NEC's ISSP (Instant Silicon Solution Platform, **Figure 3**). NEC builds the family, which now includes base arrays offering as many as 1.5 million usable gates and 3.7 Mbits of embedded configurable memory, in a 0.13/0.15-micron technology. The newly announced ISSP2 family will take the series to 90-nm technology. Embedded cores in the preISSP family include a 3.125-Mbps SERDES (serializer/deserializer) core supporting the current round of high-speed serial-interface standards. This variant, known as ISSP-HIS, will operate with system clocks to 250 MHz.

NEC does not believe that ISSP technology directly competes with its continuing cell-based ASIC business. Rather, it views ISSP as a means of widening its offering to those who would like to achieve cell-based levels of performance but are forced to use programmable solutions. ISSP aims to solve most high-speed-design problems in the base array, including signal integrity, testing, and clocking strategies. Christoph Hecker, ASIC product-marketing manager of NEC's European semiconductor and displays business unit, notes that you may still encounter signal-integrity chal-

Figure 5



Chip Express uses one of the simpler fundamental blocks of logic to build customer-configured logic. The company duplicates this cell many thousands of times, and two or three layers of metallisation on the final design interconnect its individual elements.

lenges in the routing of a design but also assures that in the upper metal layers, these problems are very "fixable." Test structures are all embedded, and multiple clocks are globally routed. NEC will accept verified RTL or synthesised netlist as a hand-over point; Hecker says the objective is an early design hand-over. Again, Tera and Synplicity tools figure into the picture. An optimised version of Synplify provides improved results in array usage, but you can also use Synopsys' synthesis. ISSP uses a relatively complex multigate cell structure with inverters, multiplexers, and a single register in each cell. You can select or bypass the individual combinational or sequential elements in the logic-to-array mapping process. Volume targets are medium-sized projects—not high, production-run numbers. Design to production time is 14 days, with NRE charges of less than \$100,000.

Fujitsu has gone straight to a 0.11/

0.09-micron (90-nm) process with its AccelArray. In the words of its European marketing manager Mark Ellins, Fujitsu intends AccelArray to "fill the gap between FPGAs and standard cell," opening up leading-edge process performance to a new market sector. The CS90A process employs six metal layers, of which three are for final programming. Ellins says that turnaround time from design completion to prototypes is typically one-third that of standard cell. The base array handles most signal-integrity and clock issues, and the device has test structures already built in. Fujitsu claims a maximum clock frequency of 333 MHz.

You can choose from two interface types: MegaFrame devices offer high-speed I/O to 600 MHz, and GigaFrame parts offer 1 GHz or more. MegaFrames are available now in 0.13-micron technology, and 90-nm devices will be ready later this year. You can expect a 30% area penalty and a 20% speed penalty relative to full standard cell, but for about one-third the NRE costs. On the AccelArray parts, memory is distributed in the regular logic structures. PLLs and clocks are preconfigured, and you must design within given clock constraints. You can configure as many as 16 banks of I/O, and each bank can use a different signalling standard. High-speed SERDES functions reside in the I/O area of GigaFrame devices. Five base-device sizes span 500,000 to 3.5 million gates, and logic is arranged in blocks of 10,000 gates each with 500

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in *EDN Europe*.

AMI Semiconductor
www.amis.com

IBM Microelectronics
www.ibm.com/chips

LSI Logic
www.lsilogic.com

Synplicity
www.synplicity.com

Chip Express
www.chipexpress.com

Lightspeed Semiconductor
www.lightspeed.com

NEC Electronics
www.ee.nec.de

Tera Systems
www.terasystems.com

Fujitsu
www.fme.fujitsu.com

Synopsys
www.synopsys.com

flip-flops. Memory is also configurable on a block-by-block basis. You can, says Ellins, use any standard ASIC design flow; Fujitsu will take the resulting data and apply a few extra tool steps to map the design onto the AccelArray structure. IP comes from the IPWare portfolio, and “platforms” with embedded high-speed I/O for a number of communications standards will follow. Back-end design takes two to four weeks, and prototypes require two more weeks. Fujitsu expects typical volumes of 5000 to 100,000 units per year.

CLOCKS TO 700 MHz

Lightspeed Semiconductor calls its structured-ASIC product, Luminance, a modular array. It designed its latest announced series, which uses TSMC’s 0.13-micron/eight-layer copper process, to provide high speed, with system clocks reaching 700 MHz. Military and wireless-infrastructure-baseband designs typify target applications. Lightspeed’s vice president of marketing and application engineering, Michael Sydow, notes that in the 0.25-micron family, the company has seen designs that, as COT (customer-owned-tooling) standard-cell exercises, might cost \$10 million to reach silicon; Lightspeed says it can reduce this bill by two-thirds. Sydow also notes that the company is finding that fabless semiconductor houses are considering the arrays as vehicles for ASSP (application-specific-standard-product) designs. Quoted array sizes are as many as 10 million usable ASIC gates with as much as 5 Mbits of embedded memory. Lightspeed has embedded PLLs, SRAM, and configurable I/O and has announced a high-

speed SERDES function. The company is considering introducing a high-performance 12-bit DAC function to satisfy the demands of the wireless-baseband market.

There is currently no specific processor-IP core associated with the modular array, although Sydow acknowledges that making one available in the technology is a priority. You can source IP from a number of third-party suppliers and import it directly into the array. Once again, you use a targeted version of the Synplicity tool that “understands” the Lightspeed logic module. Otherwise, you can use standard-ASIC-design tools. The base-array structure includes testability; an approach termed AutoTest provides 100% stuck-at-fault detection. Together with AutoBIST, AutoTest eliminates the entire design-for-test process from a design. “Test is free,” Sydow claims. You also get full connectivity to view any internal node. Lightspeed positions its offering to compete with the low to middle range of standard-cell designs. Looking forward to widespread use of 90-nm technology, Sydow anticipates that structured ASICs will be able to address as much as 70% of all designs. At 65-nm, Sydow notes, Lightspeed thinks it will address close to 100% of designs.

NOT ENOUGH MEMORY?

Not shrinking from the “gate-array” label, Chip Express calls its offering Advanced Gate Arrays and positions its product squarely as a standard-cell alternative. Its two latest introductions, in the



You can reach Editor Graham Prophet at +44 118 935 1650, fax +44 118 935 1670, e-mail gprophet@reedbusiness.com.

CX5000 series, use 0.18-micron technology (**Figure 4**). System Slice parts target general-purpose SOC (system-on-chip) designs, and the enterprisingly named Memory Pig handles applications with heavy memory demands. Eight System Slice parts span 44,000 gates and 64 kbits of fast SRAM or ROM, configurable to 1.8 million gates and 2.6 Mbits of memory. They

also include PLLs and DLLs. For “memory-voracious” designs, the Memory Pigs come in four sizes that shift the balance of memory to logic to around nine to one.

System clocks run at more than 200 MHz, but Chip Express’ vice president of marketing, Doug Bailey, anticipates that constrained logic areas will run much faster, because 200 MHz is in fact a global power constraint and not determined by gate delay. You can implement high-speed SERDES functions and other IP blocks specific to I/O functions at chip edges, where the power grid can supply ample power. The basic logic module is simple, and Chip Express constructs it around a single flip-flop (**Figure 5**). Chip Express uses a Cadence back-end placement environment, with a maze-router algorithm that targets the architecture. Design NREs are \$35,000 to \$100,000; unit prices span \$2 to \$60 (100,000/year). You can expect a three-week handoff-to-prototype cycle. Chip Express continues in production with 0.35- and 0.25-micron families that offer a range of options, including one- and two-mask programming and a “hard-array” route for higher volume. □