



## Report

### Structured ASICs Offer Application Adaptability

Cary Snyder and Murray Disman

## Summary

Structured ASICs (SAs) are a new class of semiconductor device fitting in-between FPGA and cell-based integrated circuits in cost, performance, and function. By virtue of offering fast development times and significantly lower costs over standard sell ASICs, SAs have become an important part of the Application Adaptable Integrated Circuit (AAIC) landscape. All but the highest volume semiconductor applications today require innate application adaptability to maintain market position, whether by performance-, feature- or cost-based criteria.

Rapid low-cost flexibility is the primary attribute that led to the creation of SAs and will be what drives this segment. A level of necessary application adaptability exists such that an OEM can have fewer committed resources in its manufacturing pipeline. A faster turnaround time (TAT) on critical components lowers risk and saves money. SAs offer shorter TAT compared to ASICs. This means a manufacturer can benefit from quickly adapting products to meet changing customer demands.

This report focuses on the major technical and business trends that are most important to the SA market. Forecasts for this segment on SA technology can vary widely by what definition is used for inclusion or exclusion in the SA arena. Further segmentation of the SA market is likely to occur as simple and complex SoCs are created from the same basic technology and meet the basic definition of an SA.

## Introduction

The MOS gate-array market is currently segmented into three major categories: conventional gate arrays, embedded arrays, and structured ASICs. Conventional wisdom defines the SA market to include other device types like modular arrays or platform ASIC devices, although vendors may want to claim otherwise. SemiView believes there is validity to these claims when the comparison of complexity of some platform ASICs or platform FPGAs to more simply structured ASICs is contradicting. However, in many cases there can be a functional or application overlap where both device types may compete for the same socket together with other AAIC devices.

The primary focus of this report will be on the broader definition of the SA market. With this in mind, it is important to draw distinctions between low- and high-complexity devices, with lower and higher performance requirements.

## Economics Drives the Market

The transition to newer silicon technologies, first to 130nm and now to 90nm geometries, has greatly changed the IC resource cost model to where different economics drive the market. Uncertain market adoption of new standards and features creates complicated economic issues while a product is still in development. Customer acceptance can change quickly and long product development cycles increase the risk of missing product-market forecasts. Increasing uncertainty is what primarily drives the SA market-the length of time a device or product can meet market requirements has been greatly reduced. Substantially higher resource costs further exacerbate the problem.

The economic cost model can be complicated. Expenses not only include a device's unit cost, but the tools, engineering time, and up-front vendor non-recurring engineering (NRE) charges that are a part of the system design. Table 1 is a cost comparison between a typical 1M-gate design at 0.13 microns for FPGA, structured ASIC, and cell-based ASIC implementations.

**Figure 1 - FPGA, cell-based ASIC, and structured ASIC development costs**

<b>Table 1 - Comparing Total Costs of doing in 1M gate designs in 0.13um</b>			
<b>Note: Prorated Unit Cost = Per piece cost + (Qty/Total Design Cost)</b>			
	FPGA	Structured ASIC	Cell-based ASIC
Total Design Cost:	~\$165K	~ \$500K	~ \$5.5M (Typical)
Vendor NRE:	None	~ \$100K - \$200K	\$1M to \$3M
# Tools Required:	2 to 3	2 to 3	6 to 10
Cost of Tools:	~ \$30K	~ \$120K to \$250K	> \$300K (# seats?)
# Engineers:	1 to 2	2 to 3	5 to 7
Price per chip:	\$220 to \$1K	~ \$30 to \$150	~ \$30
Total Unit Cost Qty 1K:	~ \$1000 ('03)	\$500 to \$650	\$ 55K
Total Unit Cost Qty 5K:	~ \$220 (4Q'04)	\$100 to \$150	\$1.1K
Total Unit Cost Qty 500K:	~ \$40 (4Q'04)	> \$21	\$11 to \$20

Standard FPGA-based products uniquely address the economic issues for lower volume applications by being easy to design and program with the fewest number of tools and engineers and without up-front NRE costs. However, FPGAs consume more power, are lower in performance, and can have substantially higher unit costs when compared to a standard-cell ASIC in high-volume applications. With this in mind, FPGAs are unsuitable for some higher performance, higher volume applications, or those applications with low-power requirements that are best met with cell-based ASIC technology.

Structured ASICs have been created to target the space in-between cell-based ICs and FPGAs. Unlike standard cell ASICs, SAs can offer the advantages of higher performance, lower power and lower unit cost. This is due to reduced NRE, minimum unit-volume commitment, design-tool cost, and overall engineering resources. Everything about these devices, from short turnaround time, NRE costs, fewer tools, fewer engineers and lower volume commitments creates an economically compelling case for using these devices.

## **Structured ASIC Evolution**

The evolution of a structured ASIC is similar to that of a gate array (GA) in that a pre-built "base-wafer" is stocked until ordered by a customer. The predefined architecture of a GA uses transistor pairs as basic building blocks, and final fabrication of the metal layers connects these transistor pairs to form gates, flip-flops, and other basic structures. This arrangement shortens ASIC manufacturing-cycle times, but does nothing to lower design implementation cost.

Structured ASICs use more complex logic elements than those found in GAs. For example, NEC's ISSP family of devices use a total of five metal layers with three fixed layers implementing standard logic, memory cells, clock-routing, power/ground, and debug circuitry. The last two layers are reserved for netlist routing and extra power routing for higher speed circuit options.

## **Structured ASIC Definition**

The key to reducing structured-ASIC design cost and complexity is reducing the number of custom mask and via layers. Structured ASICs are based on a predefined and pre-built logic fabric. This logic fabric can include multiple input lookup tables (memory elements), flip-flops, and multiplexers arranged as either combinatorial or sequential logic elements. Although these look to be FPGA terms, they are actually used to describe SA attributes.

These blocks can also include special structures and design features to reduce design-flow complexity, simplify the test and validation process, and reduce the number and cost of design tools. The key element in keeping the design task manageable is to limit the number of modifiable layers. Typically, they are limited to two user-modifiable metal layers, but can be as high as five. Other fixed layers can be used to incorporate PLLs, RAMs, power distribution, I/Os, and other IP depending on what the vendor implements.

## **Structured ASICs and Platform ASICs**

The real differences and similarities between a 'structured' and 'platform' ASIC become more of a choice of words than actual definition. Both implement blocks of IP, and the goals of each can be identical. Yet implementation of IP in SA devices can vary significantly in both capability and complexity—a Platform ASIC implies complex IP functions like processor or DSP cores, or complex interfaces such as USB or 1394. Segmentation of the SA arena into a separate platform ASIC is happening, but only to distinguish more complex SoC-type IP functions from simpler IP implementations.

In this regard, most platform ASICs can be considered as SAs but not all SAs can be considered to be Platform ASICs. We include as SAs those devices that limit user-modifiable metal layers to fewer than 5. Devices requiring a user to modify more than 5 layers, including partial modification of layers, are not considered SAs for the purposes of this report since they would then be closer to standard-cell technology. However, we discuss

other fast TAT ASICs based on vendor claims they meet part of the SA criteria-substantially faster development times-than a full custom ASIC.

## **Other Cost-Reduction Approaches**

Substantial reductions in development times are possible using other approaches. IBM's Customizable Control Processor, or CCP, not considered an SA since users are required to modify all the device's layers, may end up competing with SAs based on the CCP's fast development time claims. This time saving may have more to do with IBM using an FPGA together with the CCP test chip to create its CCP programmable development platform. The development platform is what enables a rapid start to system software development. This capability gives silicon designers more time to shake out and test features during functional testing.

## **Time to Market Advantages**

The design and device delivery time for an SA varies from three to nine months; standard-cell ICs take 12 to 18 months. Altera's Structured ASIC strategy is unique in that it includes five Stratix HardCopy SA members that overlap its Stratix FPGA devices. This allows Altera customers to ship and/or test an FPGA-based device prior to even starting a HardCopy conversion. Altera cleverly promotes its HardCopy Stratix devices as "a minimal risk ASIC-alternative" based on having a standard FPGA design-path compatible with a cost-reduced, performance-enhanced HardCopy device.

## **EDA Requirements**

The response from EDA tool vendors to the emerging SA market has a clear leader with Synplicity having announced special versions of its Amplify ASIC tools targeting both LSI's RapidChip and NEC's ISSP. Most SA users would be taking the lower-cost path as non-customer-owned-tooling (non-COT) type customers. With this in mind, most EDA tool vendors targeting this space realize that the primary design handoff will stay at the gates-to-placed-gates level.

Hedging bets that front-end designers, who outnumber back-end designers by a four or five to one ratio, will be attracted to the substantially lower tool cost that is part of the SA design flow. This will make EDA tools a critical requirement for everyone's success. A key attribute of the SA design flow is its lower tool cost and simplified design handoff where a placed-gates netlist is given to a Structured ASIC vendor. This is a very attractive feature.

With this in mind, companies experienced with FPGA tools might have an advantage over use to dealing with cell-based ASIC design flows. For example, Synplicity, a leading FPGA synthesis company, doesn't have to worry about creating IC-implementation tools and can, instead, be more focused on providing standalone physical-synthesis tools such as its new Amplify ASIC tool.

## Representative Tool Providers

### Magma Design Automation

Magma acquired its ArchEvaluator and PALACE tools from APlus as announced in June of this year. In addition to targeting FPGAs, the tools are also said to be suitable for Structured ASICs.

### Synplicity

Synplicity and NEC Electronics jointly announced that they will work together to craft Amplify ISSP, a specialized physical-synthesis tool for NEC's Instant Silicon Solution Platform (ISSP) structured ASICs. Delivery is expected in the first quarter of 2004. Synplicity's Amplify ASIC RTL-synthesis tool already offers mapping technology for NEC's ISSP family. The new Amplify tool will place gates to improve performance and help timing closure. Synplicity also worked with LSI Logic to develop Amplify RapidChip, targeting LSI's RapidChip structured ASICs.

The Amplify ISSP software is expected to perform automatic memory-block placement, concurrent full-chip placement, customized physical synthesis, timing analysis, and ISSP-specific datapath and arithmetic-operator generation. The tool is expected to handle a flattened ISSP of any size without a need to partition for synthesis. NEC has established a third-party ISSP EDA Vendor Alliance program with Synplicity as a charter member.

### ViASIC

ViASIC has created ViaPath, its EDA tool targeting optimized placement and routing for single-mask-programmable structured ASICs. ViaPath works as part of the proprietary ViaMask architecture and it's unclear if the company will use this tool to target other structured ASIC technologies.

### ASIC Design Starts

Obvious to many observers is the decline in cell-based ASIC design starts-worldwide design starts peaked at a little over 11,000 per year in the late 1990s and is expected to continue to drop. According to recently released American Technology Research information, ASIC design starts will drop to 2500 by 2005. Bryan Lewis, chief analyst at Gartner Dataquest, similarly forecasts a drop in ASIC design starts, but only to below 4000 by 2006. While it is unclear if SA design starts are included in these estimates, the 70-to-80 SA design starts SemiView estimates for this year are expected to grow by a healthy 28% CGAR through 2005, and reach over 225 SA design starts per year by 2007.

### The Structured ASIC Market Forecast

The SA market forecast is optimistic based on the many advantages this technology offers. The primary reason for limiting ASIC design starts or to look at alternatives is for cost. Other complexity issues and problems created by rapidly changing standards are more manageable when using faster TAT SAs. Fine tuning a particular hardware feature closer to a product's introduction helps insure its market acceptance. Shorter development time increases the chance of hitting the market at the right time with a product having the right features.

Contrary to predictions of several industry analysts, SemiView feels that FPGAs will not overtake mask-programmable chips and become the industry's dominant silicon architecture. FPGA technology will play a key role, but the increased die size and cost requirements of an FPGA, along with the device's extra power consumption are physical characteristics that will limit its wholesale use. Even the accelerated trend towards FPGAs displacing a large number of printed-circuit-board sockets now occupied by ASICs will slow,

as inherent FPGA technology limitations provide a ripe and healthy market for alternative AAIC devices such as structured ASICs. FPGA vendors seem to also recognize this trend since they, too, are gearing up to be major players in the SA market. Many long-existing perceptions, such as FPGAs are best suited for prototyping ASICs, or FPGA design tools and IP are ultra-low margin products, are just not true.

Our forecast for the SA market is positive based on simple economics. We expect the number of SA design starts to reach 80 by the end of this year and show a healthy growth rate over the next five years, with at least 225 design starts per year by 2007. There is considerable upside potential as other device types that fit the SA definition join this market segment. SemiView believes in the SA market segment as part of the trend towards AAIC-type devices that are neither ASICs or ASSPs, nor other types of fixed-function silicon devices.

### **Chip Cost and Technology Trends**

Chip-cost trends are well established in that dramatic increases in silicon fabrication costs have to be amortized over increasingly high volumes. These high volumes can be reached in different ways, but the key will be "application-adaptability," hence SemiView's focus on the AAIC market segment. Devices such as SAs are an important part of this segment. Advances in silicon-fabrication technology shift the primary cost factor from die size and package type to mask charges and development cost. Structured ASICs provide the process and means to minimize the cost impact of using more advanced silicon technology.

There are some very important technology trends that affect the acceptance of SA-based solutions. As senior managers better understand the complex economic and technical trends as they more effectively adopt newer silicon technology, the advantage of using alternative devices such as SAs will become more compelling. For example, the effect of line/interconnect delay caused the demise of older gate array technologies that made sense at 350nm, but did not scale well to 180nm or smaller silicon geometries.

The basic premise of SA technology is to simplify the design process. With SAs, the back-end design process is greatly simplified through the use of embedded IP along with specialized design and analysis techniques. These techniques cover DDR macros (in the case of Fujitsu's AccelArray), test structures, IR drop, crosstalk, timing closure, memories, hard and soft macros for high-speed I/Os, and others. The biggest design challenges are solved at the wafer level once. Users thus share the benefits derived from the substantial development costs that go into using more advanced silicon-fabrication technologies.

From FPGA and ASIC design viewpoints, SAs use the best of both worlds. SA EDA tools can deal with higher levels of design abstraction and avoid some of the toughest problems associated with the design process.

Structured ASICs have found a clear following with a number of committed design wins at multiple vendors. These designs will succeed based on the benefits they deliver to silicon design teams who are under immense pressure to lower costs. The biggest challenge facing SA vendors isn't a technical or marketing one, but is with the design teams who are accustomed to having total flexibility in a higher-end silicon-design process. The capital investment in high-end computer systems and EDA tools, and the skills required to make them all work, used to mean a lot more than they do today. In reality, a design team adopting SA technology means the group will then have the means to do more designs and take advantage of the short SA TAT compared to an ASIC design.

Several companies have already committed to developing and marketing SA and SA-like devices. Following is a summary of some of these companies, along with descriptions of their AAIC products.

## Company Profiles

### Altera

101 Innovation Dr.  
San Jose, CA 95134  
408-544-7000  
www.altera.com

### Company Background

Altera was founded in 1983 to develop and produce complex PLDs. The company added FPGAs to its product line and is now the second largest producer of programmable-logic devices. Altera leads the CPLD segment of the market and is second to Xilinx in terms of FPGA deliveries.

The company had CY 2002 revenues of \$711 million and has reported revenues of \$205.3 million for 2Q03. Altera claims some 14,000 customers and employs 1880 people in 14 countries.

The company's primary products are the FLEX 10, APEX II, APEX 20K, Mercury, Stratix, Cyclone, and Stratix GX FPGAs. Its flagship CPLD is the venerable MAX 7000 family in its various configurations. The MAX 3000A is a cost-reduced version of the MAX 7000 family.

APEX device shipments were initially seriously impacted by problems with the Quartus design system. These problems have since been eliminated and Altera is now concentrating its marketing and sales activities on the Stratix, Stratix GX, and the low-cost Cyclone families of FPGAs.

The company recently released a roadmap for product introductions during 2003 and 2004. These products are centered on the use of TSMC's 90nm process. First 90nm silicon for Stratix II is due to be shipped during 1H04, for Cyclone II by mid-2004, Stratix II GX by mid 2005, Stratix II HardCopy by mid-2005, a new version of the Nios embedded processor by 2Q04, and MAX II by the end of 1H04.

Altera is being much more conservative with their 90nm products than Xilinx. Xilinx is already sampling Spartan 3 devices from UMC's and IBM's 90nm processes and planned for production to start before the end of 2003-this date has now slipped to 1Q04.

### Structured ASIC Products

Altera introduced its HardCopy program in October 2001 for conversion of its APEX 20KE and 20KC families. The four members of the HardCopy series introduced at that time were the HC20K1500, with up to 52,000 logic elements, along with the HC20K1000, HC20K600, and HC20K400. The company claims that typical lead times for design migration and prototype manufacturing is about eight weeks.

At the same time as this announcement, Altera stated that it would produce HardCopy versions of the Excalibur family, which incorporates an embedded ARM9 processor core. The company never introduced HardCopy versions of its Excalibur devices, but did announce an extension of the HardCopy program to include APEX II devices in February 2002. The HardCopy devices for APEX II replacement were to be produced using an all copper 130nm process at TSMC.

Altera also stated that the APEX II products supported by the HardCopy program included the EP2A70, EP2A40, and EP2A25 devices. Stratix was introduced soon after the release of



APEX II. As a result, APEX II and the HardCopy option did not receive much exposure or emphasis from the company.

HardCopy qualifies as a structured ASIC since the majority of the device is fabricated to the point where only two metals layers are needed to implement a specific customer design. Actually, Altera applies two more layers of metal after the customer-specific metal is applied. These are for power and clock distribution and are common to all conversions. The customer does not have to pay for the masks for these layers.

HardCopy parts are basically the same as for initial FPGA, but without the elements that allow programming of the logic and routing and the SRAM that stores the configuration data. According to Altera, this results in a reduction of 60% to 70% in chip size that leads to a price of the HardCopy parts that is significantly below that of the equivalent FPGA.

Altera stated that it had completed 10 to 20 HardCopy conversions of APEX devices at the time it announced the Stratix HardCopy program in June 2003. Most of these conversions were for the largest APEX parts and typically involved production quantities in the 5K range. Three reasons were identified for the HardCopy customers' decision to convert an FPGA design: for production runs, as a bridge product to use while an ASIC was being developed, and as insurance against slippage in the ASIC development program.

The Stratix HardCopy program represents a subtle, but extremely important change, in emphasis compared to the earlier HardCopy programs. In the past, HardCopy was more of a defensive approach that was used to steer customers away from cell-based ASICs. The 10-to-20 conversions of APEX designs are a reflection of an early sales and marketing effort that was applied to HardCopy.

Feedback from Stratix customers about HardCopy indicated a desire for lower cost, higher performance, and lower power, with higher performance being the most important. HardCopy devices are inherently faster than their FPGA counterparts and Altera decided to take advantage of this feature with the Stratix HardCopy program. An indication of the company's seriousness in this area is that a HardCopy product group was established under a Vice President to help insure the program's success.

Altera, in order to verify performance capabilities, ran about 100 customer Stratix push-button designs that targeted a HardCopy device. The results ranged from almost no gain to an increase of 110% in performance. Minimal gains occurred for those designs that made extensive use of the hard IP in the device - the MACs and embedded memory. The best gains were seen for those designs that were dominated by routing. The average performance increase was 50%. It should be noted that these performance gain figures are from fixed FPGA designs without benefit of increasing performance or redesigning the FPGA to avoid specific performance bottlenecks.

The key to Altera's Stratix HardCopy strategy is the unified design environment contained in the latest release of the company's Quartus II tools. Altera was able to preserve the traditional HardCopy migration route for those customers that were satisfied with the FPGA's performance and were only interested in reducing costs. The customer could now do an FPGA design and automatically produce a HardCopy design. Either or both designs could be altered independently to reach the desired performance levels.

The other approach is to do and optimize the HardCopy design first and then automatically produce the FPGA design. The resulting FPGA could then be used as a prototype of the design. This alternative placed the company squarely in the structured ASIC business. A unique advantage for Altera is that both design flows can be done with the \$2000 Quartus design package.

The HardCopy Stratix family members range in density from 25,660 to 79,040 Logic Elements (LEs) and are available in FineLine BGA packages. The actual parts are the HC1S25F672, HC1S30F780, HC1S40F780, HC1S60F1020, and HC1S80F1020. The number following the S in the part designator is approximately equal to the number of logic elements in the device. Altera started accepting customer designs during 3Q03 and expects to deliver the HC1S80 starting in 4Q03 and the HC1S25 in 1Q04. The company claims that volume pricing will range from \$25 to \$120.

The delivery timetable for Stratix HardCopy devices is three weeks for the actual design work at Altera. It then takes about five weeks for prototype delivery after the customer approves the timing results. Production units can then be delivered within eight weeks after prototype approval. Typical NRE charges have been reduced to \$150K and vary with the size of the device.

Altera claims that it does not require customer-supplied test vectors. Test circuitry embedded in the HardCopy arrays includes BIST for memory and the PLLs, and boundary scan logic. Automatic test-pattern generation (ATPG) vectors are created by Altera which, according to the company, results in fault coverage of 99%.

All IP cores offered for Stratix devices can be migrated to HardCopy parts. An additional license fee may be required for the cores from Altera's third-party IP partners.

The company seems to be very satisfied with its HardCopy program. It claims that the number of companies in production with HardCopy is in the low double-digits and that they are doing two to four conversions per quarter. The backlog of conversions is growing and Altera has some 20 designs "in the hopper." One Stratix HardCopy design has been converted, but the company expects major growth in this area as it has hundreds of Stratix HardCopy customer "engagements."

The company is staffed to do 50 to 100 HardCopy designs per year. A major portion of its customer base comprises ASIC engineers who have migrated to FPGA design. Designers are fairly evenly distributed among those that first design and implement a design in an FPGA and those that target a HardCopy from the start.

## **AMI Semiconductor**

2300 Buckskin Rd.  
Pocatello, ID 83201  
208-234-6795  
www.amis.com

### **Company Background**

AMI Semiconductor (AMIS) traces its roots back to the very early days of the IC industry. It was founded in 1966 and been transformed several times through acquisitions and divestments. The company recently completed an IPO.

AMI Semiconductor reported sales of \$211.2 for the first six months of 2003. Approximately \$120 million of these sales were for mixed-signal devices. Around 50% of these sales resulted from the 2002 acquisition of Alcatel's mixed-signal business.

The company also reported \$40.9 million for the sales of "structured digital products" for the first half of 2003. Sales of these products exceeded \$200 million in 2000. The sale of structured digital products seems to be recovering and grew 18% from 2Q03 to 3Q03.

According to AMI, structured digital products are those that result from the conversion of FPGA designs to lower-cost gate-array or cell-based ASICs. The company does not consider these products structured ASICs.

AMIS' structured digital products include gate arrays, cell-based ASICs and structured ASICs. Gate Arrays have accounted for the majority of ASIC products shipped by the company to date. In gate arrays, AMI offers the 0.35-micron XLArray XL3 family targeting 2.5V and 3.3V designs and the 0.5-micron XLArray XL5 family, which the company describes as a low-cost 5V solution. The 0.35-micron family contains as much as 330k bits of dual-port memory and runs at speeds of 150MHz. AMI claims it has performed 1600 FPGA-to-gate-array conversions since 1987.

### **Structured ASIC Products**

The company introduced what it considers to be its first structured ASIC, the 1.8V XpressArray (XPA) family in January 2002. The 0.18-micron XPA devices can operate at 1.5V, but with a 30% decrease in performance. The family contains eight members with capacities ranging from 49K to 1742K ASIC gates. AMIS defines ASIC gates as equivalent to usable 2-input NAND logic gates. The embedded memory, which can be configured as single or dual port memory, ranges from 38K bits for the smallest device to 1362K bits for the largest.

In July 2003, AMIS introduced a higher density version of the XPA family, XpressArray-HD (XPA-HD), with 40% more capacity than the XPA series. Like the XPA family, XPA-HD comprises eight devices with capacities ranging from 64K to 2664K ASIC gates. The embedded memory ranges from 38K to 1403K bits.

The XPA-HD family was introduced primarily to increase the density of on-chip circuitry by 30% to 40%, thereby lowering chip cost. This was accomplished by migrating the back-end processing from 0.35 to 0.25 microns. These are the two metal layers that configure the device to a customer's design and are applied at AMIS. The capacity of the XPA family was increased with XPA-HD to match the increasing capacity of the newer FPGAs.

Both the XPA and XPA-HD families contain DLL and PLL circuitry to match the requirements for Altera and Xilinx FPGA conversions. In addition, both families contain mask-

programmable I/Os that can meet a wide range of signaling standards including those for PCI-X, HSTL, SSTL, GTL, and 622 Mbps LVDS. This capability is again needed for Altera and/or Xilinx FPGA conversions.

The basic structure of AMIS' XPA and XPA-HD families is a sea of modules surrounded by the PLLs and DLLs. Next are eight banks of I/O cells and the bonding pads. The company would not reveal the detailed architecture of the logic modules, but did say that they contained RAM, MUXs, test circuitry, NAND, and NOR gates. AMIS would not divulge whether the modules contained configured registers such as those available in other structured ASICs.

The distributed RAM approach is different from the block RAM technique used several of the other structured ASIC vendors. This approach was chosen by AMIS to facilitate the conversion of Xilinx FPGA designs that make use of this distributed feature. The company has found that many designers are using block RAM in FPGAs and intends to introduce this feature in a future product release. AMIS will introduce a new family of structured ASICs in early 2004.

The two XPA families are being produced on TSMC's 0.18-micron process. Parts are processed to the point where only two metal layers are required to configure the device to a specific customer design. The final two layers, applied at AMIS' facility, are relatively wide to minimize reticle cost, according to the company. AMIS also claims that in-house metallization provides them with a significant advantage in terms of control and turnaround times. Other structured ASIC suppliers use the foundry that produced the part to apply the final metal layers.

AMI is the historical and volume leader in the FPGA-to-ASIC conversion business. It has grown in the ASIC segment and today only slightly more than half of its design wins are for FPGA conversions. The company tends to focus on the larger customer, but claims it will work with the smaller company as well. AMI has a minimum order requirement of 2K units per year, but is really looking for a minimum of \$0.5 million of business from each design encounter.

Typical NRE charges are \$50K for a design that fits in a 44K ASIC-gate XP164E-PQ240 and \$150K for designs that require the 1.2M ASIC-gate XP560E-BG560. Typical per-unit production prices for these parts in 50K per year quantities are \$8.50 for the XP164E-PQ240 and \$38.00 for the XP560E-BG560.

According to AMIS, the timing problems that once plagued the conversion business are not a problem for most customers and that the majority of its designs are at frequencies of 133MHz and below.

## **Chip Express**

2323 Owen St.  
Santa Clara, CA 95054  
408-988-2445  
www.chipexpress.com

### **Company Background**

Chip Express was formed in Israel in 1985 as an activity within Elron Electronic Industries, Ltd. It started US operation in 1989 as an ASIC prototyping company. The technique used involved a laser to cut connections in a pre-fabricated gate-array structure to implement the customer's logic. The company completed a number of designs, but found that the technique was inefficient compared to mask-programmed devices. Chip Express discontinued the laser approach and emerged as an ASIC supplier in 1995. The company claims that it has completed more than 1000 structured ASIC designs.

Chip Express, which is private, had sales of about \$29 million in 2001 - a peak sales year for many in the electronics industry. Chip Express is experiencing a substantial pickup in sales and has had 10% quarterly sequential revenue growth for this year. It expects that 4Q03 revenues will be the highest since 1Q01. Chip Express has raised a total of \$44 million in venture backing.

### **Structured ASIC Products**

Chip Express is by far the leading supplier of structured ASICs, having shipped substantial volumes of these products for at least the past four years. The company actually claims that all of its products, even its 0.6-micron CX2000 family, are structured ASICs. The company's more recent products are its 0.35-micron CX3000, 0.25-micron CX4000, and the 0.18-micron CX5000. The CX5000 was introduced in April 2003 and actually consists of two families - the CX5000 System Slice and the CX5000 Memory Pig. The parts in both families contain both analog PLLs and DLLs.

The CX5000 System Slice family contains eight members ranging in capacity from 44K to 1755K ASIC gates. The corresponding block RAM ranges from 64K bits to 2568K bits. The CX5000 Memory Pig series consists of four members with capacities ranging from 117K to 546K ASIC gates and block RAM ranging from 1104K bits to 4548K bits - a much higher ratio of memory bits to ASIC gates than in most other structured ASICs.

The logic cells are relatively simple compared to other structured ASIC architectures. While it would not divulge the precise details of the cell, the company did say that each cell contained a multiplexer, one NAND gate, and two drivers. It appears that a configured register is not included in the logic cell.

The I/Os are made to match those of the FPGA design through the use of the metal layers. The company claims that it can configure some 1400 different I/O cells in this manner. The maximum speed attainable on the differential I/Os is 622 Mbits/sec, a little short of the 740 Mbits/sec needed for some communication systems.

Chip Express claims that its ASIC gate ratings are based on 70% utilization of logic available on the chip. It has seen designs where 75% and even 80% utilization has been achieved.

None of the CX 5000 parts are yet in production, but designs are underway and the company expects that production orders will start to be filled before the end of 2003. The company claims that its design capability for the 0.18-micron CX5000 is completely utilized.

Chip Express can currently execute three designs at one time and can do a maximum of 60 designs per year. It plans on hiring additional staff during 4Q03 and 1Q04.

The company is finding that the size of its contracts has been increasing over time. Typical production orders for its 0.35-micron parts were \$150K, \$250K for its 0.25-micron devices, and \$400/\$500K for its 0.18-micron CX5000. It is now rejecting smaller jobs. Typical NRE costs are \$54K for a one-million-ASIC-gate design, with per unit production parts priced at \$41 for 10K unit quantities.

Even though the company does not target the conversion business, some 30% of its revenue is derived from the conversion of FPGAs - mostly for small chips. Chip Express does not do the design conversion, but farms the work out to one of several third-party design companies. These firms convert the design to a Chip Express netlist that is free of timing problems and changes the embedded circuitry in the FPGA to match the logic in the structured ASIC.

Chip Express has not yet decided whether its next family will be produced using a 130nm or 90nm process. It did state that a new family will be introduced next year and that it will be very different from its current structured ASIC architectures.

## **eASIC**

3555 Woodford Drive  
San Jose, CA 95124  
408-264-7128  
www.easic.com

### **Company Background**

eASIC was founded in 1999 with a goal of simplifying ASIC design and manufacture and to shorten turnaround time. The company is pursuing the licensing of its eASIC fabric for use as a core in ASICs and ASSPs. eASIC is also considering offering structured ASICs on its own or in conjunction with a partner.

ST Microelectronics will be using the fabric as an embedded core and Flextronics Semiconductor will be a partner to co-develop and market the structured ASIC. Initial test devices are now being produced by Flextronics.

Flextronics, through its acquisition of Orbit, has been active in the FPGA-to-ASIC conversion market for about 10 years. The first structured ASIC products will be produced using a 0.13-micron process and should be available by the middle of 2004.

### **Structured ASIC Products**

eASIC does not have a Structured ASIC product that it can deliver, but has designed and evaluated the basic programmable fabric for such a device.

The eASIC logic cell comprises a pair of three-input SRAM-based LUTs and a register, very much like the basic building block used in SRAM-based FPGAs. The cells are interconnected using the upper metal layers. The company's approach is unique among structured ASIC suppliers in that the interconnection can be accomplished using a single via mask. A segmented wiring grid allows for the customization to be done by via changes between Metal 6 and Metal 7.

The logic programming of the cell is done through the SRAM and, like an FPGA, requires that a bitstream be loaded into the device during power-up. The ability to change the bitstream offers a number of advantages as it is possible to slightly modify the design after the device has been fabricated. In addition, the bitstream can be modified for testing and debugging operations. The cell logic can be permanently set by using the via mask, thereby eliminating the need to download a bitstream.

eASIC has suggested that a direct-write e-beam approach can be used in place of the via mask. This technique, while not cheap, eliminates the need for any masks and can significantly reduce delivery time. Fujitsu, for one, estimates that the e-beam approach can reduce the prototype delivery cycle from four weeks to one week.

The company makes the point that it is much faster and less expensive to implement the via than a metal layer when using e-beam direct-write. This is because the vias occupy only 1% of the chip's area versus 30% for the metal in the layer. eASIC claims that customization through the via results in e-beam throughputs that are 10 times faster than for a metal mask.

Another advantage of the e-beam approach is that multiple designs from a variety of different customers that can be written on the same wafer. This clearly reduces the prototype's cost along with the cost of low-volume production.

The agreement with Flextronics gives eASIC the opportunity to develop its fabric into a family of structured ASIC products. Optimizing the tool flow and embedding the test and other IP structures will keep the two companies busy until the middle of 2004. Faraday Technology



## **Faraday Technology**

Hsinchu, Taiwan

886.3.578.7888

[www.faraday.com.tw](http://www.faraday.com.tw)

### **Company Background**

Taiwan-based Faraday Technology was founded in 1993 and had a public offering on the Taiwan Stock Exchange in 1999. The company had 460 employees and revenues of \$96.2 million in 2002. Its US operations are difficult to assess.

Faraday describes itself as a fabless ASIC company and an ASIC/SoC design services and IP provider. It claims that it has developed 800 silicon-proven IP cores and has completed more than 1500 ASIC projects.

IP offered by Faraday includes RISC CPUs, DSPs, USB 2.0, Gigabit Ethernet, and Serial ATA. The major part of the company's revenue is derived from the delivery of ASICs to its customers. Almost all of its work is done with UMC as a silicon foundry.

### **Structured ASIC Products**

Faraday announced its 3MPCA (Three-Mask Programmable Cell Array) in June 2003. The company also calls this product a Flexible ASIC. Like most of the other structured ASICs, 3MPCA is customized using three masks for two metal layers and a via layer. Faraday announced in a news release that the 3MPCA family is available today for UMC's 0.35-, 0.25-, 0.18-, 0.15- and 0.13-micron CMOS processes for ASIC and IP customers.

Faraday has not defined a family of products and it is highly doubtful that products are available in all these different processes. It is difficult to say whether the company is supplying structured ASIC products or a core for insertion into a customer's ASIC.

One surprise is the relative complexity of the basic logic cell, which the company claims contains several LUTs, drivers, and registers. The family supposedly covers the range from a very low 5K to 5M ASIC gates, has a built-in scan chain, and supports unlimited clock domains and gated clocks. Another claim is that the devices can implement any type of combinatorial and sequential logic.

## **Fujitsu Microelectronics America**

1250 E. Arques Avenue  
Sunnyvale, CA 94088-3470  
408-737-5600  
www.fma.fujitsu.com

### **Company Background**

Fujitsu Microelectronics America (FMA) was established in 1979 as a subsidiary of Fujitsu Ltd., a \$38 billion company. FMA produces a wide variety of electronic components that include CMOS ASICs, networking and wireless ASSPs, microcontrollers, memory devices, and displays. The company is among the top three suppliers of both gate-array and cell-based ASICs.

### **Structured ASIC Products**

FMA introduced its first structured ASIC product line, the AccelArray, in April 2003 and claims to have registered more than 20 design wins since then. The Mega Frame family is being produced using Fujitsu's 1.2V 0.11-micron process and comprises five devices with capacities ranging from 512K to 3.8M available gates and 860K to 4.55 Mbits of embedded memory.

The company claims a maximum operating frequency of 333MHz, 800MHz PLLs, and LVDS I/Os that will run at 311MHz and support 622 Mbps. An array of two-input NAND gates is the customizable logic fabric. Three-to-four metal layers are required to implement a customer's design.

The replicated block in the AccelArray architecture is an array of 43K usable two-input NAND gates surrounded by embedded RAM blocks. There are a number of register files, primarily to aid in the I/O functions. The array also contains an embedded DDR macro that can pass data at a rate of 400 Mbps and metal-programmable I/Os.

Future releases will probably contain configured registers embedded within the NAND array blocks. The use of simple two-input NAND gates for the user-configured logic leads to a requirement for more metal layers than those needed for architectures with more complex logic cells. LSI Logic's RapidChip is another example of a structured ASIC that needs more than two metal layers to interconnect its array of NAND gates.

FMA is also delivering its Giga Frame version of the AccelArray family. This series, to be announced during 1Q04, is similar to the Mega Frame series, with the main difference being the inclusion of metal-programmable SERDES structures that can operate at rates up to 3.125 Gbps. Fujitsu is finding that PCI Express is a very popular I/O standard with this family.

A good deal of the circuitry for the different high-speed I/O standards is embedded in the part. The company supplies other pieces of the I/O protocols as soft macros. One member of the Giga Frame series will contain an embedded ARM 9 processor.

The company expects to migrate a redesigned version of AccelArray to 90nm in mid-2004. Devices from this new family could support speeds as high as 500 MHz.

FMA is quoting four-to-eight weeks for the time required from netlist acceptance to tape out. Most designs can be done in four weeks and it is only the more difficult ones that might require the full 8 weeks. First samples are delivered four weeks after tape out. NREs of \$200K are typical for AccelArrays at the 130nm. The company would not discuss minimum

quantities, but stated that AccelArray has been very well accepted and that they are now being very selective in adding new customers.

Fujitsu also plans on offering direct an e-beam writing option as a means of speeding the sample delivery process next year. The company estimates that sample delivery time can be reduced from four weeks to one week. This is not a high-volume process and it would still be necessary to produce masks for production runs.

## **Lightspeed Semiconductor**

209 N. Fair Oaks Ave.  
Sunnyvale, CA 94085  
408-616-3200  
www.lightspeed.com

### **Company Background**

Lightspeed was founded in 1996 with the specific idea of offering an array of complex cells that could be customized with a few metal layers. The company had several false starts in pursuing the FPGA-conversion market, but eventually developed a strategy based on the structured ASIC business. A total of \$85 million has been invested in the company. The latest round of \$35.5 million was closed in 1Q02.

The company developed and introduced three generations of its modular-array architecture. Its latest family, Luminance, was to be fabricated using a 0.13-micron process at TSMC. Lightspeed has 0.35-micron and 0.25-micron (Lightning) modular-array families in production.

Lightspeed, in November 2003, again changed its business strategy and moved from supplying silicon products to being a vendor of its intellectual property. The company announced that it would not be accepting new structured ASIC customers, but would endeavor to license its mask-programmable cores to ASSP and ASIC companies. In addition, Lightspeed will offer versions of the fabric directly to customers for inclusion in their silicon designs.

Lightspeed claims that it has a number of ASSP companies interested in using the 0.13-micron fabric. These same companies were not at all interested in the company's 0.25-micron or 0.35-micron technologies. Using a Structured ASIC approach begins to make more sense for creating variations of a standard product as the cost of developing completely new devices at 0.13-micron skyrockets.

The company discovered that many potential customers did not want to deal with a small startup company and had difficulties developing a production base for its 0.25-micron and 0.35-micron families. While it generated substantial interest in the 0.13-micron Luminance series, Lightspeed found that cost of marketing and developing the family and the NRE for implementing all of the devices was more than it could bear at this time.

### **Structured Array Products**

Lightning, Lightspeed's 0.25-micron modular-array family, contains some 11 members with usable gates ranging from 30K to 1 million, assuming an estimated usage of 60-70% of the available gates. This range of gates is accomplished with the use of 4K to 127.7K modules.

The corresponding embedded SRAM in 4-Kbit blocks ranges from 81.9 to 851 Kbits. Two members of the family have embedded 8-Kbit blocks, while one member has 14 embedded 64-Kbit blocks of embedded SRAM. Additional smaller blocks of memory can be implemented using the modules in the array.

Lightning targeted FPGA-to-ASIC design conversions from Xilinx's Virtex E FPGAs. The company claims that some 20 conversions were done, but found this to be a very difficult process. FPGA-to-ASIC conversions never became of major part of Lightspeed's focus.

Luminance, Lightspeed's 0.13-micron family, consisted of eight arrays with capacities ranging from 200K to 10M usable ASIC gates. Embedded RAM ranged from 720K to 5.7

Mbits. All of the family members contain 18-Kbit blocks of embedded RAM. The three largest members of the family contain 512-Kbit RAM blocks and fewer 18-Kbit blocks.

Other features of the new family included metal-programmed I/Os that could meet a wide variety of single-ended and differential-signal standards. Lightspeed planned to include a 3.125Gbps SERDES in a future release of the family. The devices also included AutoBIST and AutoTEST circuitry that provides 100% stuck-at-fault coverage and 100% visibility of internal nodes.

The basic logic cells in both the Lightning and Luminance families are relatively complex and, in addition to the test circuitry, contain buffers, sequential elements, and two multiplexers plus an AND gate. It is interesting that Lightspeed, in the data sheet for the Lightning series, rates each logic cell as having an equivalent capacity of 10 available and 8 usable gates - a relatively conservative rating for the amount of logic contained in the cell.

In spite of Lightspeed's supposed lack of interest in the FGPA-to-ASIC conversion business, a major part of the only application note on Luminance is devoted on how to convert Virtex-II designs to a Luminance device. It is questionable as to whether Lightspeed ever did produce a production qualified Luminance part.

## **LSI Logic Corporation**

1621 Barber Lane  
Milpitas, CA 95035  
866-574-5741  
www.lsilogic.com

### **Company Background**

LSI Logic, founded in 1981, is one of the major ASIC companies. It reported total 2002 revenues of \$1.81 billion and employs about 5000 people worldwide. The company tends to focus on high-end ASICs, ASSPs, host bus adapters, software, and storage systems.

### **Structured ASIC Products**

LSI prefers to refer to its RapidChip families of devices as a Platform ASICs, rather than Structured ASICs. Like Fujitsu's AccelArray, RapidChip uses an array of gates to implement the user's logic with several metal layers. The customizable gate-array structures used by LSI and Fujitsu require more metal layers, five for LSI and three for Fujitsu, more than the two layers typical of other structured-array technologies with more complex logic cells. It is by far the most complex device being offered in this market, but it is likely that the other ASIC suppliers will migrate to this level of complexity to match their customer's needs.

There are four basic families within the RapidChip series-the Foundation Slices and StreamSlice, produced using LSI's 0.18-micron process, and the 0.11-micron Integrator and Xtreme families. LSI claims that it has recorded more than 13 design wins at 0.18-micron and more than 15 wins at 0.11-micron.

All of these parts are being produced using the Black Diamond low-k dielectric material. At the time of RapidChip's introduction, the company said that it expected to have 90nm versions available late in 2003.

RapidChip was introduced in September 2002 with the expectation that it would cut entire design cycles to six months, the same as required for FPGAs and about one-half that needed for a cell-based ASIC design. LSI, in January 2003, introduced the first slice, StreamSlice, for Foundation.

StreamSlice targets high-end switches, routers, and other communications-system applications and offered 3.3 Mbits of embedded memory and 3M usable gates for customer logic. The platform contained a number of high-speed embedded interfaces including 12 GigaBlaze (1.0625 - 3.2 Gbps), 36 HyperPHY (622-832 Mbps), and an 80-bit DDR SDRAM interface. Cores to work with these interfaces for SPI4.2, 10/100/1G/10G MACs, XGXS, and Fibre Channel are also available.

LSI then introduced a series of seven Foundation Slices in March 2003. The customizable capacity ranged from 0.5M to 2.5M gates and the corresponding amount of embedded memory ranged from 0.6 to 2.0 Mbits. The embedded memory was composed of 144-Kbit single-port RAM, 36-Kbit 2-port RAM, 9-Kbit dual-port RAM, and 9-Kbit 2-port RAM blocks. Three of the slices contained embedded 200 MHz ARM926EJ-S processors. The parts also contained 4 to 12 GigaBlaze 3.18Gbps SERDES blocks supporting interfaces such as 1- and 10-Gbit Ethernet, Fibre Channel, Serial ATA, Serial Attached SCSI, and PCI.

The flip-chip-packaged Xtreme family targets high-bandwidth I/O applications. The family supports slices with 8, 16 or 32 GigaBlaze channels of up to 4.25Gbps SERDES and is compliant with 1- and 10-Gbit Ethernet, Fibre Channel, Serial ATA, Serial Attached SCSI, and PCI Express standards. In addition, the Xtreme family also includes support for LSI

Logic's HyperPHY transceivers that enable multiple OIF SPI4.2 and SFI4.1 link interfaces at 155 to 3200 Mbps per channel.

While specific platforms within the Xtreme series have not yet been defined, LSI states that high-performance 333 MHz processors are pre-built into several of the slices. This includes processors such as ARM 1136, ARM 1026, MIPS 5Kf, and ARM 926. According to LSI, all slices in the Xtreme family support the ARM 7 and 966 in soft or firm formats. The processor IP is delivered as a complete subsystem with the AMBA bus and peripherals provided to the customer.

The Integrator family is designed for mainstream logic applications and is available with capacities ranging from 2.9M to 9.8M ASIC gates. Embedded RAM can reach 5.3 Mbits. DDR PHY support for 400 Mbps is also provided. Designs with up to 25 logic levels can run at speeds up to 200 MHz.

Configurable I/Os will support SSTL, HSTL, CMOS, LVTTTL, PCI, and LVDS. Some of the slices will support multi-standard-compliant SERDES with speeds up to 4.2 Gbps. LSI will support soft or firm ARM processors.

LSI claimed that two slices from the Xtreme family and four slices from the Integrator family were available for design when the family was introduced in July 2003. Pricing for the entry-level slice starts at \$25.00 for 100,000 units/year quantities in 2004. NRE for the 0.18-micron families ranges from \$100K to \$200K and from \$200K to \$400K for 0.11-micron parts.

## **NEC Electronics America**

2880 Scott Blvd.  
Santa Clara, CA 95050  
www.necel.com

### **Company Background**

NEC Electronics, established as a separate entity in November 2002, is a wholly owned subsidiary of NEC Corp., one of the world's largest electronics companies. NEC Electronics reported sales of \$1,414B for 2Q03.

The company describes itself as a firm specializing in non-DRAM semiconductor solutions. Its products include a variety of processors and peripheral components, SRAM and flash memory, a number of ASSP devices, linear components, audio-video ICs, interface devices, and gate array, cell-based, and Structured ASICs. NEC claims to be the leading supplier of gate-array ASICs.

### **Structured ASIC Products**

NEC Electronics introduced SoCLite two years ago. This product included an embedded ARM7TDMI processor, a number of peripheral components, a relatively small amount of memory, and 190K gates of user-configurable logic. Targeted applications for the part included factory automation, industrial-bus systems, card readers, business telephones, terminals, and home-communication equipment. NEC described SoCLite as ideal for emerging applications or applications requiring rapid time-to-market because of the device's low unit cost, low NRE, and short turnaround time.

SoCLite meets the requirements for being included as a structured ASIC. It does, however, use a sea-of-gates for the user-defined logic, not the more complex logic cells found in other structured ASICs. The issue would be how well did this technology scale. NEC still supports the part, but is not actively promoting its use.

NEC's primary structured ASIC products are its series of ISSP (Instant Silicon Solution Platform) devices. Its first family, ISSP 1, was announced in March 2002, with availability set for 3Q02. This series is being produced using NEC's 0.13-micron (drawn) UX4 ASIC 1.5V process- UX4 is really a 0.15-micron process. The family, at the time of the announcement, consisted of three devices with usable logic capacities of 227K, 530K, and 1109K gates. The corresponding numbers of embedded 16-Kbit SRAM blocks were 16, 48, and 64.

ISSP 1 contains embedded analog PLLs and DLLs. The DLLs are mainly used for DDR interface applications. NEC would not divulge details of the logic cell except to say that it contains both combinatorial and sequential elements. The family allows the implementation of a wide variety of I/O signaling standards, including 3.3V LVTTTL and 2.5V LVCMOS, PCI, SSTL and HSTL memory interfaces, LVPECL, and LVDS I/Os

The company, in March 2003, expanded the ISSP 1 offering with two devices. One part offers up to 1.5M usable gates with 2.5 Mbits of embedded memory. The other device supports 1M usable gates and contains 3.7 Mbits of memory. The two new parts contains 32 4-Kbit SRAM blocks each, in addition to the 16-Kbit blocks in the earlier family members. The new additions also contain four eight-phase high-speed analog PLLs.

Another addition to the new devices is the use of a common metal layer that is applied after the two customer-specific metal layers. The primary purpose of the final layer is to reinforce power line distribution, especially to the center of the chip where IR voltage drop can cause problems. One more mask is required to redistribute I/Os for flip-chip packages.



The ISSP devices use a relatively large logic cell. While NEC would not reveal the details of its logic cell, it does list the number of registers contained in each part. The number of configured flip-flops ranges from 9151 for the 214K-gates part to 65,076 for the 1,693 million-gate device. This amounts to a claimed usable gate rating of about 25 for each logic cell - about twice the 12 gates normally claimed by FPGA manufacturers to rate their four-input LUT/register logic cells.

Clock circuitry is embedded in the base device. Customers do not have to provide test vectors, since the test technologies, including SCAN, BSCAN, BIST, and TestBus, are also embedded in the base part.

NEC, at the same time that it announced the larger devices in the ISSP 1 family, announced the ISSP1-HSI (high-speed interface) family with embedded high-speed SERDES circuitry. The 0.15-micron ISSP-HSI family consists of three configurations with usable densities of 500K, 700K, and 1M gates. The corresponding amounts of embedded SRAM are 780 Kbits, 1.3 Mbits, and 2 Mbits. Four 3.125 Gbps SERDES blocks are included in the smallest part in the smallest family. The two larger parts each contains 16 3.125Gbps SERDES blocks.

The SERDES can be set for 3.125Gbps, 2.5Gbps, 1.25Gbps, and 622 Mbps and can be implemented for XAUI, GB Ethernet, InfiniBand, PCI Express, and Fibre Channel compliance. 8B/10B encoding/decoding is accomplished via a soft core. The 3.125Gbps interface dissipates 220 mW per channel. NEC claims to have achieved 5 Gbps operation when using a 0.13-micron process. The company expects to be in volume production with this family before the end of 2003.

In June 2003, NEC announced its intention to migrate the ISSP families to a 90nm process. Device libraries for the new 1.0V ISSP 2 series are expected to become available during 2Q04, with engineering samples and production-device availability starting during 2Q04. The company expects to achieve clock speeds of up to 500MHz.

The precise lineup of parts for ISSP2 has not been set, but will include devices with up to 4M usable gates and 10 Mbits of embedded SRAM. NEC plans to release the ISSP2-HSI series during 2H04. These devices will contain 10 Gbps SERDES as well as a next-generation 3 Gbps ATA interface, up to 3M usable gates, and 10 Mbits of SRAM.

NEC claims that it has recorded 30 design wins for ISSP 1 and that 15 of these are already in production. The company says that these wins are from external customers and from in-house design groups. When announced, NEC stated that the turnaround time from design to production would be 14 business days and the NRE was expected to be below \$100K.

The company says it will accept orders for as few as 400 pieces, but, like other structured ASIC vendors, it is really looking for opportunities of at least \$0.5 million per year. It is considering implementing a "shuttle-like" program for its IP partners and to handle small orders.

NEC has not done many FPGA conversions, but seems to be concentrating on its traditional ASIC customers. The FGPA conversions done are typically to increase performance, to add features to the design, or to combine multiple FPGAs into a single chip. Performance is the key parameter that leads designers away from FPGA approaches.

A wide variety of interface/communication IP cores have been re-synthesized for the ISSP devices. These NEC-developed cores include Ethernet, POS Phy Level 3, SPI4.2, UTOPIA,

and PCI/PCI-X. NEC intends to expand its core offerings through the use of third-party suppliers. It will accelerate core development by providing ISSP platforms.

NEC has recognized a potential roadblock in the buildup of its structured ASIC business - a shortage of in-house design engineers to work with customers for adapting designs for the production of GDSII tapes for the ISSP masks. NEC will, in order to circumvent this potential problem, certify a number of design houses to aid in this effort. It has already identified one U.S., one Canadian, and 3 to 4 Japanese companies for this program.

The company is also working to optimize the design flow for ISSP. It is currently using Synopsys and Synplicity for synthesis to the ISSP structure. Synopsys is using its traditional design compiler, but Synplicity has developed ISSP-specific synthesis tools.

Synplicity has been working with NEC since June 2002. Its latest effort is the Amplify ISSP Physical Optimizer. One key aspect of the design tool being developed is that it will handle any ISSP as a flat design. Not having to partition the design for synthesis leads to substantial improvements in performance and area. Synplicity expects that Amplify software for ISSP1 and ISSP2 devices will become available during 1Q 2004.

[www.semiview.com](http://www.semiview.com)