

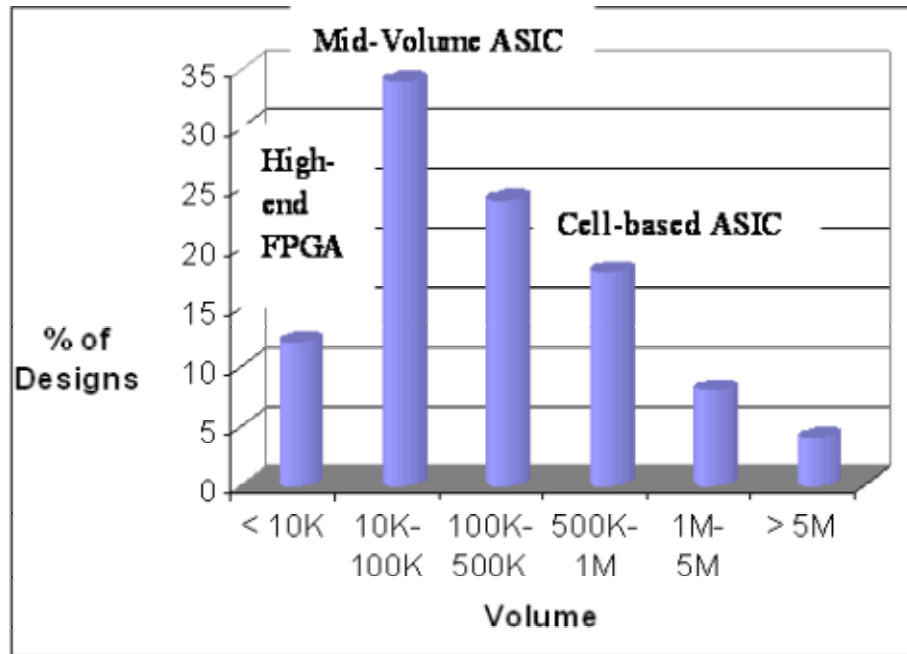
# **Structured ASIC Design: A New Design Paradigm beyond ASIC, FPGA AND SoC**

**Dr. Danny Rittman**  
**August 2004**

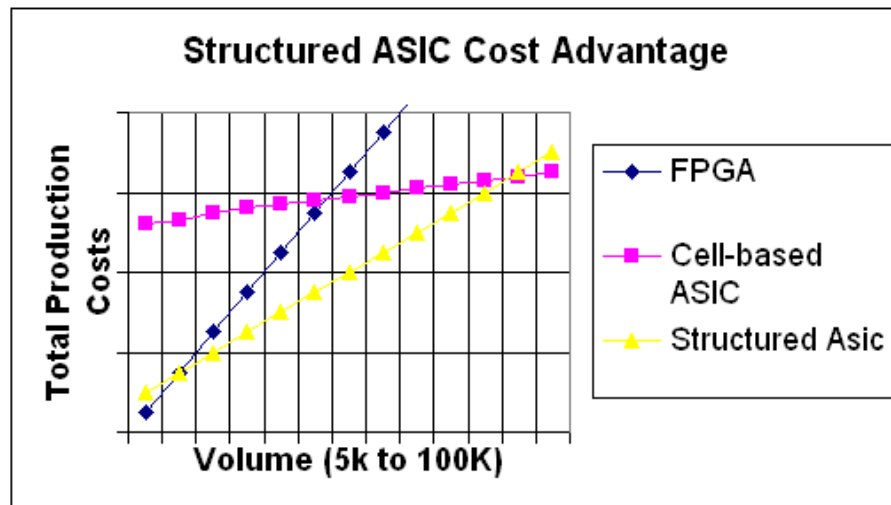
## **Abstract**

Standard Cell ASICs are well known in the IC industry and have been successfully used over the past decade. During recent years there is a significant reduction of traditional ASIC design according to Gartner/Dataquest. (About 1/3rd the number of ASIC designs today compared to 3 or more years ago – See Figure 1) Trying to obey Moore's law and the fact that mask costs passing \$1M, design cycles expanding past a year and reliability issues arising from VDSM physical effects, it's now require a massive production run to justify cell based ASIC. Yet, the demand for custom ASIC performance still remains for a wide variety of applications. In some cases, FPGA's are successfully replacing ASIC's despite their high unit costs, limited performance, high power consumption and power dissipation. Systems manufacturers may still achieve cost reduction in case the product is successful but this can be done in a later stage. The high cost of the development, masks, and tools can justify the production of very high-volume cell-based ASICs only. For mid-volume ASIC designs, the costs are simply becoming out of reach. The industry had to come with an effective and affordable solution. A different concept that will provide custom ASIC's performance yet, with a short design cycle and low cost manufacturing. A new direction was taken, Structured ASICs.

Structured ASICs offer a cost-effective solution for the mid-volume ASIC design with 75% less development costs than cell-based ASICs and unit costs up-to 90% less than complex FPGAs. (See Figure #2) In this article I will present a new design paradigm that has come of age, the structured ASIC.



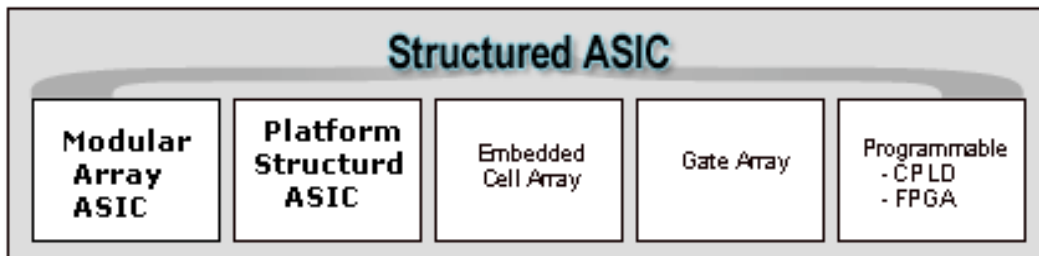
**Figure 1 – ASIC's Volume**  
**Image Source: Bay Area Chip Design**



**Figure 2 – Structured ASIC's Advantage**  
**Image Source: Bay Area Chip Design**

## Introduction

In custom logic design the logic cells are physically organized to implement a system, and further divided into device configuration to implement the desired functions. A Standard Cell ASIC is created by assembling a collection of standard cells, each with an optimal implementation of a particular logic component. The design task becomes mapping the design into a completely 'blank line up' of silicon by choosing from a large library of building block cells, and interconnecting them as necessary. A relatively large number of different collections of standard cells can be implemented to form the same function. Structured ASICs simplify the complexity of custom silicon design by providing a set of identical building block cells that are prearranged in a series of sizes and complexities. This means the design task is mapping the circuit into a fixed arrangement of known cells, rather than mapping standard cells to the design. Using a given set of cells in a variety of combinations is significantly reducing the design time due to the fact that all necessary functions are available immediately. Structured ASIC have fewer layers to customize which implies much shorter manufacturing time. They have fixed layers to incorporate Memories, I/Os, Power lines, PLLs, and other logic components. While similar structured ASIC's are not gate arrays. While gate arrays address the manufacturing cycle time issue for custom ASICs, structured ASICs address the design implementation issues such as time to handoff, cost of design tools, engineering resources, number of design functions, NRE charges, IP integration, and layout turnaround time. A wide variety of technologies fall into the Structured ASIC domain such as Modular Array ASIC's, Embedded cell array and more as illustrated in Figure 3.




**Figure 3 – Structured ASIC's Domain**  
**Image Source - LightSpeed**

## Structured ASIC – The Concept

Structured ASICs provide a new ASIC capability that offers a promising alternative to cell-based ASICs for the large and mid-volume market. Structured ASIC technology uses pre-diffused base metal layers to implement functions that would be common to many designs. For example: Memories, power grids, I/O's, clocks, IP, PLL and other advanced functions. The custom logic is then implemented in a few metal layers, typically between two to five metal layers, thus requiring far fewer mask layers to be created for each design (see Figure 4). The entire design cost significantly drops due to the greatly reduced mask, fabrication, and engineering costs. In addition, the time to market factor is extensively reduced enabling faster product delivery to the end customers. Due to their fast design time Structured ASIC's are very attractive for prototyping and system development.

ASIC vendors are taking a new approach to plan their structured ASIC products, eliminating design flow steps, such as SI (Signal Integrity) analysis, power grid, IP integration, Memory insertion and other difficult or tedious tasks. In addition, structured ASIC vendors reduced the engineering team and design tools needed in order to shorten the design cycle time and lower the overall cost.



**Figure 4 – Structured ASIC's Silicon Cross-section**  
Image Source - EETimes

## **Structured ASIC Revolution leads to new design tools and flow**

An important aspect of structured ASIC design is the design flow and tools. The design tools have to match this new type of ASIC design and more important, the vendor's methodology. Using the current industry standard tools would be non-efficient and expensive. The current design tools are designed for massive designs and consumes significant amount of memory and resources. This new direction leads to a structured ASIC customized design tools demand from EDA vendors. With fully customized EDA tools the design time is much faster and there is major improvement in the timing results. This means that the decreased performance and density inherent in a structured ASIC can be mostly recovered by using tools the ASIC vendor has jointly engineered with the EDA vendor. Effectively, this means performance and density have been taken off the table with respect to whether a system design might use a cell-based ASIC or a structured ASIC to meet their system goals. New generation of EDA customized tools has been introduced to the market.

Since generic synthesis tool will almost always overuse elements for structured ASIC's design and under-use others, there is no balancing of the design resources to maximize logic density. Customized synthesis tools can perform automatic resource balancing, particularly in developing arithmetic and datapath operators achieving simultaneously better area and timing performance. A good example is the physical synthesis area. Here it becomes very important to be vendor-specific in the tools used prior to final routing and layout by the ASIC vendor. EDA vendors are working together with structured ASIC vendors to incorporate into physical synthesis tools features like Vendor-specific LVS/DRC check, Clock distribution constraints, Pre-defined placement for diffused hard macros, Pre-defined floor-planning and die size, Custom routing models and more. With these pieces automated in the user's synthesis tool, instead of separated across multiple tools, parties, and geographies, the design handoff process can be done more efficiently and automatically. With VDSM cell-based ASIC designs the issue is no more timing only but also signal integrity, power grids and other considerations that have to be implemented within the flow. With physical synthesis tools customized to the specific structured ASIC architecture all of these problems can be greatly reduced or completely eliminated automatically by the tools. For example structured ASIC devices have a power grid already developed and pre-diffused in the base layers. With a cell based ASIC design the physical synthesis tool is only working off of approximation provided by the designer in their floorplan, not on the actual power grid itself.

As expected, when the actual power routing is done by the cell-based ASIC vendor, many factors will change such as the space required by the power routes, the IR drop effect of the power grid, and how critical cells are placed with respect to the power grid. This leads to design closure problem. Comparing structured ASIC with a predefined power grid that allows the physical synthesis flow to actively use the power grid directly in physical synthesis. With a power grid, instantaneous IR drop calculations can be performed directly in physical synthesis; that way the tool can understand the true voltage received by each cell, and therefore make a more accurate assessment of timing across that cell. With this detailed knowledge of the routing structures, the physical synthesis tool can automatically perform vendor-specific design rules checking on the complete design, ensuring that when the placement and netlist data are handed off to the structured ASIC vendor, it is done only once and without iterations.

### **Structured ASIC manufacturers and Products**

There is a constant grow in structured-ASIC products and vendors. 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. The next are major structured ASIC manufacturers and products.

### **Altera – Together with Synopsys go Structured**

ASIC design tool leader Synopsys teamed with leading FPGA vendor Altera to develop together solutions for the design and production of Structured ASICs. Altera has long touted their HardCopy structured ASIC as a clean cost-reduction path from an FPGA-based development, prototype, and early production platform to a cost-reduced, performance-optimized mask-programmed equivalent. Altera acknowledged that the advantages of programmable logic for early development will compel design teams to consider their structured ASIC offering. Altera came out with its own structured ASIC entry, dubbed the HardCopy program, in 2001. Through this program, customers that used Altera's FPGA devices can convert them to hard-coded versions, gaining the benefits of higher performance and lower power on a design that had already been proven in the market.

## **Faraday – UMC Spin-off**

Faraday introduced its first structured ASIC, the 3MPCA (three-mask programmable cell array), in 2003. The UMC spin-off has historically focused on IP and ASIC design, with 250 to 300 design projects per year. It is currently offering a metal programmable cell array and will offer a metal programmable I/O later this year, according to Martice Chen, vice president of marketing in the United States.

## **AMI Semiconductor – Replacing FPGAs**

AMI Semiconductor explicitly targets the FPGA-conversion market, continuing the business it has for several years been conducting in that area. AMI introduced second generation of structured ASIC products under the names XpressArray and XpressArray HD (higher density). 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, 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. This type of geometry enables 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  $\mu$ W/MHz/gate. The new design requires less power and ground pads than an FPGA, retargets it to a smaller package. Because the largest devices are denser than the largest currently available FPGAs, designers 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 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.

### **LSI Logic - Embedded IP targets communications**

LSI Logic offers structured-ASIC called RapidChip. RapidChip is structured around an ASIC-design flow, and it aims to deliver a reduced-NRE ASIC product. Devices are stocked prediffused with IP cores from LSI's CoreWare program; LSI describes the prediffused parts at the uncommitted stage as "slices."

The hard-IP blocks are comprise processor cores and other functional blocks that suit a device to a target market; designers then add their own, LSI's, or a third party's IP, plus custom logic. Designers can perform a complete configuration with the last few layers of metalmask programming. LSI was first to fabricate the parts in 0.18-micron and then in 0.11-micron technology.

LSI Logic also has customized tool set for RapidChip: RapidWorx targets low-level issues—particularly those that arise with 0.11-micron processes. The tool is working in the background to automatically resolve design detail issues, concealing them as far as possible. A five-button tool chain gets you from RTL input to a placed netlist, the company says. The tool chain includes customized 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 designers to control the placement of major blocks, such as memories; and the Physical RTL Optimizer, 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.

Designers can use the tool set in a highly automated "default-setting" mode, or they can intervene to fine-tune its processes. Using the TeraForm engine allows the tool set to generate a floorplan from physical-synthesis principles. 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. According to LSI, if designers are using RapidWorx they carry out much more of the process before handoff, which exploits the features of the architecture to reach the best cost point.

Although the design flow is ASIC like, 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. Designers will be able to add soft IP comprising all the commonly used interface standards, plus logic unique to your own design.

### **NEC ISSP Solution Platform**

NEC's offers ISSP (Instant Silicon Solution Platform). 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 pre ISSP family include a 3.125-Mbps SERDES (serialiser/deserialiser) 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 designers may still encounter signal-integrity challenges 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

synthesized 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 optimized 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. Designers 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 – Straight to Nanometer**

AccelArray from Fujitsu is aimed straight to a 0.11/0.09-micron (90-nm) process. 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.

Designers 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. Designers 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 the design must be done within given clock constraints. As many as 16 banks of I/O can be configured, and each bank can use a different signaling 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 flip-flops. Memory is also configurable on a block-by-block basis. According to Ellins, designers can 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.

### **Lightspeed - Clocks to 700 MHz**

Luminance, a modular array is a structured ASIC product from Lightspeed Semiconductor. The product uses TSMC's 0.13-micron/eight-layer copper process, to provide high speed, with system clocks reaching 700 MHz. This product is aimed for military and wireless-infrastructure designs 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 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. Designers can source IP from a number of third-party suppliers and import it directly into the array using 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. 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.

### **ChipExpress – Advanced Gate Array?**

Chip Express calls its structured ASIC products Advanced Gate Arrays and positions them as a standard-cell alternative. The CX5000 series use 0.18-micron technology. System Slice parts target general-purpose SOC (system-on-chip) designs, and the '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 (500MHz in constrained clock domains), 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. Designers 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. 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). It takes about three week for 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.

## **Conclusion**

It's clear to most observers that Standard Cell ASIC is becoming far too costly for all but the largest organizations and huge volume production opportunities. The entry-cost of the mask NRE (\$1.5M+ at .13 $\mu$ ), the large suite of necessary EDA tools, and the design teams required to successfully complete a design create a significant barrier for access to the attractive elements of the standard cell: lowest unit cost and highest performance. When a proper total cost analysis is done, projecting all expenses and opportunity costs across a product lifetime, structured ASIC in general, and Modular Array ASIC in particular, emerge as the ideal choice for the majority of custom silicon applications. The number of announced ASIC vendors offering structured ASIC devices continues to grow, NEC, LSI Logic, Fujitsu, Lightspeed, and others. The commitment from these ASIC vendors to offer design flows that offer substantial improvements in lower risk, lower cost, higher automation, and higher performance is clear. All major EDA vendors already implanted structured ASIC support within their tools. Using Structured ASIC technology dramatically shortens the design cycle by simplifying the design flow which cuts down engineering time by months and reducing manufacturing time to weeks compared to months for cell-based ASICs.

## Reference

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

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

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

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

Flextronix Semiconductor - ASIC Design Practice

Qlogic - Simplified Hardware Design

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

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

Design and Reuse Magazine

EE Times