Tier Logic, Inc.

Taking semiconductors to the next level

Tier Logic: Taking Semiconductors to the Next Level

Preface

The drive to reduce the full function of a computer or consumer device to a single integrated circuit (IC) has relied heavily on system-on-a-chip (SoC) design methodologies, which result in very complex million-gate devices to define next-generation computing systems. But the SoC movement has not meant success for the traditional semicustom IC or the ASIC. In fact, companies such as LSI Logic Corp., who defined their business on ASICs, have changed business plans as their core ASIC business has all but disappeared.

The reasons for this are simple; ASIC design requires:
- Expensive EDA tool suites
- Large teams of experienced design engineers
- Endless verification time
- High NRE costs
- Risk of multiple, costly re-spins.

Design, test and verification costs for a 40-nm ASIC can exceed $30 million, with the mask set adding another $2 million. System designers who elect to hand over software-verified designs to an ASIC foundry to get a customized IC still face the same concerns.

One industry answer to the ASIC dilemma has been to "do it yourself" – program a device that can take on several personalities for different applications. Since the late 1980s, the ASIC has increasingly been replaced by user-programmable logic, most commonly implemented in an FPGA. In more than 20 years of product development, FPGAs have matured to a predictable architecture, dominated by the two largest FPGA vendors, Xilinx, Inc. and Altera Corp. With these programmable solutions, the system designer has traded off lower upfront NRE costs and fewer delays against higher end-unit cost, power, and performance.

But are current FPGAs the best possible platforms for user-programmed custom designs? Tier Logic, a new player in the FPGA market, can show that a separation of the functional logic of the device and the configuration circuitry that programs the device, results in a TierFPGA™ device with greater performance, greater efficiency, and lower power dissipation than any device on the market today. At the same time, Tier Logic's solution allows an easy, cost-effective migration to a TierASIC™ solution that is based on the same die, preserving the original TierFPGA design, including its exact-timing performance. The advantages are great enough to break the stranglehold the top two vendors have had on the FPGA market for the last few years.

Introduction

Tier Logic is changing the programmable-logic landscape by introducing a new means of reducing FPGA cost and power dissipation while increasing logic density. This new technology also allows effortless FPGA-to-ASIC conversion through a single mask change. This is all accomplished through a method Tier Logic calls "TierFPGA," which separates static configuration circuitry and active logic circuitry into two layers that are stacked on top of each other.

In more than 20 years of evolution, FPGA architectures have advanced in density and performance largely through improvements in process technology. Despite a variety of experiments with programmable technologies, such as metal-fuse, EPROM, EEPROM, flash, and anti-fuse, mainstream FPGA vendors have reached a consensus on the use of SRAM as the process technology of choice. SRAM lends itself to CMOS process scaling rules quite well.

Consequently, FPGA cost, performance, and power improvements have solely relied on CMOS technology nodes (or transistor size). As a result, improved performance is only achieved through process scaling according to Moore's Law. The FPGA supplier has become a prisoner of silicon process node. If a higher-value FPGA is desired, the supplier must move from 90-nm to 65-nm devices, then to 40-nm, and next even to 28-nm devices. However, the penalty for early technology adoption comes at a price—devices fabricated with the smallest process geometry carry the highest price tag due to high risks and development costs. At the most advanced process nodes, performance and power benefits are diminishing. Static power can be several watts for large FPGA devices at the 40-nm node.

This reliance on smaller and smaller CMOS transistor size carries both obvious and subtle consequences. Fixed IO voltage standards do not scale with process scaling, and supporting these standards adds complexity and cost to advanced process nodes. As the largest and most complex FPGAs rely on cutting-edge 40-nm processes, the cost of the entire manufacturing process, and thus the price per device, can be difficult to reduce. To improve performance and maintain power consumption, the added circuits in the FPGA increase silicon area, adding even more to the cost of usable logic gates. Conversion from FPGA to semi-custom ASIC (such as gate array, structured ASIC, and standard cell) implies a redesign and/or re-optimization, which requires an NRE charge ranging from several hundred thousand dollars up to millions of dollars, depending on the complexity of the design, and a considerable delay in the order of three to six months.

In an environment where the FPGA market appears to be locked in a two-way race between Xilinx and Altera, the promotion of new approaches in programmable logic would appear to be problematic. Both vendor-specific and third-party FPGA design tools have improved significantly over the last decade, and the user base has evolved with the tools. New entrants at the turn of the millennium made few inroads, due to a focus on speeding up design tools and enhancing logic and routing resources. These gains were marginal as the silicon overhead associated with the programmability in traditional FPGAs was never addressed. Thus, with no new significant FPGA technology advancements in the industry, existing Altera and Xilinx users learned to live with the limitations and became proficient designers with those architectures and design flows. As a result, users of these FPGAs hardly see a need to learn new design-entry methods or architectures. Convincing them to move to another vendor would invariably require one or more compelling reasons to consider an advantageous and novel FPGA technology that adheres to existing design methodology and tool flows.

Tier Logic's revolutionary TierFPGA technology provides several compelling reasons to take a look beyond the conventional offerings:

- Significantly lower unit price than traditional FPGAs;
- 70-percent lower device power dissipation in static mode
- Very low NRE cost for conversion to an ASIC (a TierASIC device)
- Superior design security and soft error immunity
- Use of traditional tool flows: design entry, simulation, and physical implementation are all managed with readily available tools from Mentor Graphics

When these advantages are considered together, it is clear that the Tier Logic™ solution provides a superior alternative to traditional FPGAs, as Tier Logic has resolved the main problem of today's FPGA – the overhead penalty of programmability.

Tier Logic has accomplished this breakthrough in part through its pedigree: Every executive at Tier Logic has over 20 years' experience in FPGAs, logic, or EDA tools. The U.S. Patent and Trademark Office has already granted over 50 patents to Tier Logic covering its architecture and process, with more than 20 more pending.

By developing a new architecture capable of leveraging existing design tools and standard CMOS manufacturing techniques, Tier Logic can offer breakthroughs in density, power consumption, and performance that require no additional costs to be borne by systems designers, third-party IP suppliers, or semiconductor foundries. In fact, the TierFPGA device is a better product, and far easier to convert to an ASIC than any standard FPGA in the market today.

Process: Simplicity Through Separation

FPGAs in general contain two types of circuits: usable logic and configuration circuits. Usable logic circuits require performance and carry timing signals of the user's design. Configuration circuits are static and orthogonal to timing signals, and only program the usable circuits. In traditional FPGAs, both of these circuits coexist in the same silicon. The configuration circuit is invisible to the systems designer, but nevertheless consists of over 50 percent of the transistors in the silicon, consuming about 30 percent of the total area, resulting in cost, power, and performance overhead. Configuration bits are prone to soft errors and piracy. Tier Logic has addressed these inefficiencies by simply stacking the configuration circuitry on top of the usable logic circuitry.

The benefits of this rearrangement are many:

- Freed-up silicon is used for packing in more user logic, thus enabling higher levels of system integration within a single IC.
- Routing distances between related usable logic is reduced, thus providing higher performance.
- Shorter routing also reduces wire capacitance, thus lowering power consumption and increasing overall system performance.
- Moving the configuration circuitry to the second layer provides a natural separation for an intuitive migration path to an ASIC. The upper RAM layer is easily removed and replaced with a single metal ROM layer, capturing the identical functionality and timing at a lower system cost, lower unit cost, higher reliability, and better security.

The configuration circuitry is implemented using proven low-cost thin-film transistor (TFT) technology. TFTs have been in mass production for years, specifically in display technology, and follow process scaling rules. Tier Logic uses the low-cost TFT SRAM cells as the non-performance-critical configuration memory. Thus, the two-tier solution is cheaper, faster, and consumes less power than traditional FPGAs. As designs evolve in the R&D stage, the TFT-SRAM-based TierFPGA devices provide a platform for altering designs on the fly, prototyping, and low-volume production. Once end products start shipping and volumes ramp up, the same design can be converted to a Timing-Exact™ solution, using the same base silicon, where the TFT layer is simply replaced with a single metal mask. Tier Logic calls this transformation a TierASIC device. Tier Logic software tools allow users of traditional FPGAs to migrate their existing FPGA design to a TierASIC device within a few days with no extra design effort.

While other vendors attempt to provide a similar migration path, these other methods are not implemented using the same base die, so they require additional design efforts and costs to meet timing and technology limitations. With most vendors, existing FPGA designs must be converted to an ASIC design to get any cost reduction. For most customers, having to deal with altered timing means a complete re-verification of the design and removes the interest in ASIC as a cost-reduction path except in the most compelling cases. This is especially true in an era where engineering resources are becoming particularly scarce, meaning that few design teams are able to afford the cost and time of designing twice–first for the FPGA and then for the ASIC version. Tier Logic's approach means that the engineering team designs once and is finished. In fact, the decision to convert to the TierASIC device can be an operational one based on stability of the design and volume ramp—the design team doesn't even need to be involved.

Other advantages of Tier Logic technology are less obvious but nevertheless significant. Because configuration TFT is not employed in TierASIC devices, they have better immunity to soft errors than conventional FPGAs; this is applicable to everything from radiation-hardened applications to higher-reliability commercial applications. The lower power dissipation of the TFT-based TierFPGA device is further reduced in the TierASIC device because hundreds of millions of configuration TFTs are eliminated. Within the TierASIC methodology, designs can be further optimized to further reduce resistance and wire capacitance for higher performance or lower power consumption with no loss of the design's functional integrity. Because Tier Logic provides a series of product densities in pin-compatible packages, TierASIC users can migrate across densities as their design needs change—at very low NRE costs. Tier Logic has thus extended traditional FPGA migration to ASIC migration for systems designers.

One of the significant reasons for the reduction in new ASIC designs over the last few years has undoubtedly been the high and increasing costs of NREs. Tier Logic has made a significant step towards enabling an increase in the number of designs that can justify using ASIC technology by reducing the NRE down to an extraordinarily low level–the cost for migration to a TierASIC device is in the range of $25 thousand. Furthermore, the turnaround time is only four weeks, far faster than any normal ASIC can be manufactured.

Tier Logic's FPGA design tool flow consists of familiar synthesis, place-and-route, and simulation methodologies; design tools from Mentor Graphics support synthesis targeting Tier Logic architecture. Designers familiar with Mentor Graphics' design-entry, synthesis, and simulation tools can readily use module generators, libraries, and functional simulation models targeted for Tier Logic architecture. Following functional simulation and synthesis, Tier Logic's GUI-based Mobius™ software suite can be used for placement, routing, timing analysis, and generating configuration bit-streams. Any design written and debugged in an existing FPGA device can be moved to Tier Logic devices within days. Expert designers wishing to use script-based flows can ignore the GUI and write scripts to drive the tools through Tool Command Language (TCL).

The FPGA Market Yesterday and Today

FPGA market softness in 2001 did not mean a turn back to gate arrays or cell-based semi-custom ICs. Instead, system designers relied on a mix of processors, DSPs, and application-specific standard products (ASSPs) to implement similar products. While these designs were often far from optimal in footprint and power dissipation, they served the purpose of keeping prototype costs to a minimum. As access to capital is severely restricted today, more end users and suppliers will continue to watch their NRE budgets very closely.

In the last two years, however, FPGAs have made a comeback and ASSPs have lost favor in several applications. In data-path processing for networking, for example, the highly touted network processor has all but disappeared as OEMs discover that their value-add is resident in proprietary protocol processing at layers three and four in the OSI stack. Meanwhile, the option of proprietary ASIC designs has all but disappeared, with vendors such as LSI Corp, Fujitsu Microelectronics, and NEC stepping back from gate arrays and structured cells.

In order to ease the pain of conversion from FPGAs to ASIC-equivalent designs, traditional FPGA leaders devised structured-ASIC conversion paths with low NRE cost, exemplified by Altera's HardCopy program. Without a programmable technology and FPGA-compatible tool flows, LSI's Rapid-Chip structured-ASIC had a much lower-than-expected market appeal, which led to its demise. Similarly, NEC pulled out of their ISSP structured-ASIC offerings. In this realm, however, traditional FPGA vendors are also limited in conversion approaches due to the melding of configuration and logic elements in the device. For example, Altera's HardCopy is a completely new gate-array device, whereas Altera's Stratix FPGA is a traditional LUT-based device. Design conversion from Stratix to HardCopy is still a complete redesign due to the completely different timing of the two implementations, even though FPGA tools are used for both.

Traditional FPGA vendors implement logic architectures that are suitable for the configuration technology they use. The EPROM-based cells adopted by Altera and the EEPROM cells adopted by Lattice lend to product-term-based NAND-NOR logic ICs. The anti-fuse-based cells used by Actel and SRAM cells used by Xilinx lend themselves to truth-table logic ICs. In most cases, the programmable overhead can reduce logic density as much as 20x over similar gate-density ASICs. Furthermore, high-voltage requirements related to EPROM, EEPROM, flash, and anti-fuse do not scale with shrinking gate oxide thicknesses and fine spacing rules when processes scale to smaller geometries. In order to take advantage of manufacturing scaling, FPGA vendors have abandoned configuration architectures based on high-voltage requirements and have converged on SRAMs. Tier Logic has adopted the next-generation configuration technology ideally suited for FPGAs by using TFT-based SRAMs. In addition to using a mature industry process technology, Tier Logic leverages the scaling of TFT that follows identical electric-field scaling laws of the logic transistors. Hence Tier Logic can deliver superior FPGAs with older process technologies compared to traditional FPGA suppliers. By separating the configuration circuitry from the usable circuitry and moving it to a second layer, Tier Logic has solved a previously insoluble problem: that of converting FPGA to ASIC without having to design twice because of different timing results in the two silicon implementations.

The FPGA Market of Tomorrow

Tier Logic offers a single solution wherein the user can keep the programmable overhead or remove it completely without altering the design. The two-tiered programmable-logic solution can simply replace any existing FPGA device to reduce power or save cost; it can be used for design debug, prototyping, and mid- to low-volume production. The same-die TierASIC device is a low-NRE design-freeze technique to move the TierFPGA device into a much lower-cost, high-volume solution with no risks.

Tier Logic offers advantages in manufacturing that benefit all: the system designer who seeks productivity and low unit prices, the silicon supplier who seeks a high-volume, standard-process platform for silicon, and the "soft" intellectual property (IP) provider who seeks their "soft" IP to harden into high-volume ASICs without silicon qualification. By using a single mask change in migrating from a TierFPGA device to a TierASIC device, Tier Logic allows the customer, silicon supplier, and IP provider to take advantage of economies of scale. By using industry-standard design entry and implementation tools and allowing verified netlists for distribution, Tier Logic offers a platform for third-party IP vendors to develop, optimize, and deliver soft-logic cores to system designers.

Competition in future FPGA architectures will revolve around the cost, power, and performance of hard- vs. soft-IP cores and around the integration of multiple system functions onto a single device. Earlier architectures often relied on hard-wired cores for cost and performance, though the limited choices in hard IP had turned vendors to greater use of software equivalents. As Xilinx and Altera have moved to 65-nm and 45-nm processes for newer generations, soft cores such as Altera's Nios and Xilinx's MicroBlaze are proving to be more suitable than the older embedded hard-IP microprocessor cores.

Tier Logic can offer FPGAs in mature feature sizes that can meet higher soft-IP performance in a much smaller footprint – hence lower cost – and at much lower power consumption. The design methodology used in an independent configuration layer allows for straightforward "higher speed" implementation of third-party soft IP, while the change to a single metal layer will offer a "frozen" soft IP in the TierASIC platform. Such a methodology offers IP vendors to develop their soft IP within a software environment and incur no manufacturing or silicon characterization costs to implement the exact soft core in a TierASIC device. The systems designer has many more soft IP options to use.

Conclusion

Some analysts may ask if designers can afford to consider a startup like Tier Logic in the current FPGA environment. We would ask if designers can afford to continue using traditional FPGA architectures, when the cost and performance is worse and the path to cost reduction is much less certain than in our new 3D layered approach.

Contact

Tier Logic, Inc. U.S. Headquarters
2975 Scott Boulevard, Suite 215
Santa Clara, CA 95054

www.tierlogic.com

Phone: 408.970.9500
Fax: 408.970.9599
info@tierlogic.com