

This white paper describes the emergence of SoC FPGAs, the drivers behind that emergence, and strategic considerations for executive management and system designers when choosing these devices.

Introduction

Semiconductor devices that integrate FPGA fabric, hardened CPU subsystems, and other hardened IP—SoC FPGAs—have reached a tipping point that will lead to their broad proliferation in the next decade, therefore offering many options for system designers. These SoC FPGAs complement the decade-long availability of soft-core CPUs and other soft IP for building systems on FPGAs. A mix of technical, business, and market forces underpin this tipping point, and vendors such as Altera, Cypress[®] Semiconductor, Intel[®], and Xilinx[®] have announced or shipped SoC FPGA devices. ⁽¹⁾

The primary drivers of this tipping point are:

- The move to parallel and multicore processing for power efficiency
- The shift of FPGAs to the leading edge of new semiconductor process technology
- Increasing use of FPGAs in embedded systems
- The economic realities of Moore's law
- The consolidation of CPU architectures

As the SoC FPGA era emerges, systems designers must consider the following key strategic questions when choosing these devices:

- Which devices will experience a “platform effect,” that leads to a self-reinforcing cycle of vendor, ecosystem, and customer development?
- Which devices will support IP reuse across the broadest array of options?
- Which FPGA technology allows for the maximum cost and performance effectiveness?

Tipping Point for SoC FPGAs

The industry's first decade of integrating FPGAs and CPU systems was marked by both success and failure. The initial SoC FPGAs met with limited commercial success ⁽²⁾ while the use of soft CPUs on FPGAs was broadly accepted, ⁽³⁾ showing the fundamental market desire for the integration of FPGA and CPU technology. A variety of new factors is reshaping the environment, leading to a tipping point where SoC FPGAs will broadly proliferate in the market.



101 Innovation Drive
San Jose, CA 95134
www.altera.com

© 2011 Altera Corporation. All rights reserved. ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS and STRATIX are Reg. U.S. Pat. & Tm. Off. and/or trademarks of Altera Corporation in the U.S. and other countries. All other trademarks and service marks are the property of their respective holders as described at www.altera.com/common/legal.html. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.



The key factors driving the industry to the tipping point are:

- [Computing Power Efficiency](#)
- [FPGAs Move to Leading Edge Process Technologies](#)
- [FPGA Adoption in Embedded Systems](#)
- [Moore's Law Economics](#)
- [CPU Architecture Consolidation](#)

Computing Power Efficiency


The evolution of computing is toward parallelism, with the near-term focus on processors shifting from higher, single-core processing power toward multicore implementations.⁽⁴⁾ As part of this quest for higher computing performance at lower power, attention is moving toward using FPGA logic as hardware accelerators for CPUs.⁽⁵⁾

A monolithic SoC FPGA system maximizes power efficiency and software partitioning flexibility. SoC FPGAs allow hundreds of data signals to connect different functional areas, thus enabling 100-gigabits per second (Gbps) bandwidth, or greater, with nanosecond-level latencies, representing orders of magnitude better performance and latency than discrete implementations. Furthermore, monolithic integration permits sharing of memory controllers, allowing high-bandwidth memory access for hardware accelerators.

This increased performance and memory access enables the use of the FPGA for finer-grained accelerators to address a broad class of computing requirements. Because hardware accelerators can have more than 1,000 times the power efficiency than CPUs, designing for SoC FPGAs creates significant potential for an improved approach to power efficient computing, beyond what simple multicore parallelism can achieve.

FPGAs Move to Leading Edge Process Technologies

In the year 2000, state-of-the-art FPGAs were built on 130-nm process technology, while state of the art CPUs were built on 90-nm process technology. Because more advanced CPUs were available, the appeal of the first generation of SoC FPGAs was somewhat dampened. Today, however, leading-edge FPGAs are targeting 28-nm process technology, which relatively few commercial CPUs or ASSPs use, or are likely to use, in the near future. The process technology advantage of FPGAs significantly increases the market potential of—and the inclination of vendors to invest in—these integrated devices, as designers do not need to compromise on the CPU's capabilities, as illustrated in [Figure 1](#).

Figure 1. FPGAs Move to Leading Edge of Process Technology

FPGA Adoption in Embedded Systems

In the year 2000 FPGAs were still relatively expensive for most embedded systems applications, and as a result were used less frequently than their CPLD or PAL counterparts. In the last decade, however, SRAM-based FPGAs enjoyed riding the CMOS cost reduction curve, such that EE Times annual embedded survey states that nearly 50% of embedded systems also contain FPGAs. ⁽⁶⁾ Because one of the major appeals of SoC FPGAs is their reduced cost compared to discrete implementations, there is a large native market for silicon vendors to achieve a return on investment.

Moore's Law Economics

Moore's law is becoming more expensive. Fabrication facilities that build advanced CMOS semiconductors cost approximately \$6B-\$10B to build. ⁽⁷⁾ Because it costs \$40M to build a new semiconductor, ⁽⁸⁾ a semiconductor must recover \$100M of gross profit to operate in a typical profitable model where 20% of revenue is spent on R&D. At typical gross margins of 50% ⁽⁹⁾, a firm must capture a market of \$200M or more. Outside of consumer electronics, mobile handsets, and PCs, there are few application markets of this size, making single-purpose, or fixed function, device investments difficult to justify. As advanced semiconductor costs grow even more in coming process technologies, this cost structure will make it even more difficult to economically justify building fixed-function semiconductors, suggesting that programmable technologies will see increasing investment, while fixed function devices, including specialized ASSPs and CPU derivatives, will see less. Because SoC FPGAs have the potential to serve many markets, they will be the target of growing investment levels.

CPU Architecture Consolidation

The term *embedded processing* covers a broad spectrum of applications, from extremely cost-sensitive 4-bit processors, to extremely sophisticated multicore 64-bit machines. Similarly, this application breadth has long supported a widely diverse and fragmented base of processor types, operating systems, and software vendors. Yet that diversity is markedly different in the year 2011 as compared to the year 2000. For

all its size and variety, the overall embedded market is shifting en masse to faster and more capable processors; for example, 16-bit microcontrollers give way to 32-bit CPUs. At the same time, support for 32-bit CPU families is consolidating around four widespread architectures: ARM®, MIPS®, PowerPC™, and x86. The reason for this consolidation is due mostly to software commonality and reuse.⁽¹⁰⁾ As a result, SoC FPGAs that incorporate one of these CPU architectures can address an intrinsically larger market, thereby adding more incentive for vendors to invest in this class of semiconductor.

The Platform Effect

Products that have an adoption interrelationship between producer, user, and ecosystem tend to exhibit what is called a *network effect*,⁽¹¹⁾ or a *platform effect*. The basic tenet of the platform effect is that the more use a particular product, or standard, attracts, the more valuable it becomes to members of the user base and ecosystem. As a result, the members of the user base and ecosystem then invest more in the technology, thus attracting more use and creating a self-reinforcing cycle. Familiar examples include PCs,⁽¹²⁾ video recording formats,⁽¹³⁾ and social networking sites.

In general, product categories that have the innate possibility to operate in a self-reinforcing cycle have a strong tendency to evolve in the cycle, as each actor in an emerging product category looks to maximize ROI. As soon as a potential platform starts to gain momentum, it begins to attract investment, and thus the market dynamic can quickly shift toward that standard.

SoC FPGAs have a strong likelihood to see this platform effect. As SoC FPGAs proliferate, there will be a strong desire for users to reuse their FPGA IP and design software knowhow across a wider variety of systems. For example, members of CPU ecosystems will prefer to learn as few FPGA development tools as possible and CPU vendors will prefer to minimize the amount of FPGA development tools and technology they need to know. As a result, an SoC FPGA platform that supports multiple vendors and CPU architectures will be best positioned to trigger this platform effect investment, thereby creating the most advantages to customers and ecosystem members who join in its adoption.

Altera's Approach

After a decade of innovations in embedded systems, Altera has now launched its “Embedded Initiative” with the intent to create a multivendor, multi-CPU architecture SoC FPGA platform based on a single FPGA design flow methodology. The FPGA design flow methodology serves as the basis for a wide variety of SoC FPGAs over time, as well as SoC solutions using soft CPUs and other soft IP. ARM (hardened), MIPS (soft), and Nios® II (soft) CPU options are or will be available from Altera, while the Atom E6X5C Configurable Processor devices will be available from Intel. This integrated approach unites three of the leading CPU architectures and the most popular FPGA-based soft CPU around a single FPGA architecture and design flow.

Fostering Innovation

The integrated approach to FPGA design flow methodology is intended to stimulate the ecosystems from leading processor architectures to invest in a single FPGA platform and tool flow, thus creating an enriched set of tools, application software, operating system software, and professional expertise. With the investment of hundreds of collective, global ecosystem members of CPU architectures, this FPGA platform and its growing list of tools, software, and IP becomes increasingly useful and important to system designers, presenting them with the value proposition that will trigger their adoption, and thus fuel the virtuous platform cycle.

Providing Powerful Tools and IP


A key component of this multivendor platform is the Quartus® II software flow for programming FPGA logic. In addition to its other merits,⁽¹⁴⁾ the Quartus II software includes the Qsys system integration tool, which features Altera's second-generation switch fabric technology used to accelerate soft IP development, reuse, and integration. The GUI-based Quartus II software, available as a free Web Edition or fully licensed version, addresses productivity and performance needs with a design flow that includes system design and timing closure methodologies, in-system verification, and third-party EDA tool support.

In addition to Altera's traditional Avalon® Memory-Mapped (Avalon-MM) interface and data path bus interface specification, Qsys will also support the ARM AXI™ standard, which will allow an automated “mix and match” approach to integrating Avalon-based IP with AXI-based IP. Qsys allows a fast, intuitive design experience that supports easy design reuse and in-system verification across the entire common platform for ARM- and Intel-based SoC FPGAs, as well as MIPS and Nios II soft CPU SoC implementations.

Tailoring a 28-nm Device Portfolio

Altera's 28-nm FPGA portfolio delivers the industry's most comprehensive set of device options tailored to customers' diverse design requirements.⁽¹⁵⁾ Altera provides a substantially superior set of FPGA architectures and process technologies for a variety of end application requirements—high-density Stratix® V devices that optimize for performance, high-volume Cyclone® V devices that optimize for cost, and mid-range Arria® V devices that provide a balanced mix. The entire set of SoC FPGAs benefits from this custom-tailored approach.

Altera's newest SoC FPGA will contain an advanced processor block based on the ARM Cortex-A9MP core, as shown in [Figure 2](#):

Figure 2. Altera SoC FPGA Architecture

The Altera SoC FPGA architecture will feature a broad array of hardened IP within the ARM-Cortex A9 subsystem, as well as high-performance multiport memory controllers for maximum memory bandwidth. High bandwidth, low latency interconnects between the FPGA and the CPU subsystem will allow for high-performance applications and efficient FPGA hardware acceleration. Advanced internal switch fabrics will support efficient data throughput, as well as highly productive in-system observability and debugging. When combined with Qsys, the Quartus II software, and ARM Connected Community software tools, this device will represent a highly cost-effective system design option that leverages standard tool flows for maximum productivity, and offers significant new development and verification possibilities.

Conclusion

The era of SoC FPGAs has begun. These devices have reached a tipping point driven by key economic, technical, and market forces, and numerous vendors have already announced or are shipping. Executive management and system architects should carefully consider platform effect, IP reuse, and FPGA process technology benefits when evaluating system solutions.

Altera has partnered with leading CPU vendors ARM, Intel, and MIPS, to provide a common FPGA platform for SoC FPGA devices and soft CPU solutions. This partnership will allow the industry's most widespread CPU architectures, and their ecosystems, to inherit the same advanced FPGA design flow, thereby maximizing the IP reuse and flexibility within this platform. This integrated approach creates a platform effect that will foster growth and development on this platform, as well as the supporting CPUs and their ecosystems.

Acknowledgements

- Chris Balough, Sr. Director, Software, Embedded, and DSP marketing, Altera Corporation

Notes

1. Altera announced its “Embedded Initiative” on Oct. 12, 2010, including its intention to produce devices that integrate hardened ARM subsystems and 28-nm FPGA technology; Cypress Semiconductor ships the PSoC 5 device family that integrates programmable PLD logic, ARM Cortex-M3 CPUs, and A/D functionality; Intel announced the Atom E6X5C Configurable Processor, combining an Atom SoC with an Altera FPGA into a single package, on Nov. 23, 2010; MicroSemi (formerly Actel) shipped the SmartFusion device family that integrates FPGA logic, ARM Cortex-M3 CPUs, and programmable analog; Xilinx announced its Embedded Product Platform on April 27, 2010 which integrates ARM Cortex-A9 CPUs with 28-nm FPGA logic.
2. In 2004, Xilinx acquired Triscend, and later discontinued all its 8051-based and ARM7-based Configurable SoC devices. Atmel's® FPSLIC family, which combined an AVR processor with FPGA logic, is no longer actively promoted. Altera's Excalibur™ devices that combined the ARM926 with FPGA fabric was never extended beyond the original family. Xilinx included hardened PowerPC processors in some members of its Virtex® FPGA line starting in 2002, but stopped including that hardened processor beginning with their Virtex-6 generation.
3. Remarks from Altera Corporation during their October 2010 “Embedded Initiative” press announcement indicated that 30% of their shipping devices are ultimately programmed with a Nios II processor or other soft CPU.
4. “Issues of power consumption and heat generation have led the foreseeable future of processor engineering away from increases in processor speed, in favor of increases in parallelism, in the form of increasing numbers of execution cores per chip.”, *Transitioning Software to Future Generations of Multi-Core*, Intel Corporation, June 1, 2009.
5. Examples include 1) IBM's Liquid Metal initiative to create a common software approach for CPUs with FPGA accelerators; 2) XtremeData sells FPGA accelerator plugins for Intel and AMD platforms; 3) the FPGA 2011 Conference in Monterey features a panel on Feb. 27, 2010 entitled “The Role of FPGAs in a Converged Future with Heterogeneous Programmable Processors” that includes presentations from AMD, IBM, Intel, and NVIDIA.
6. EE Times Embedded Survey, 2010.
7. “...the cost of a leading-edge fab exceeds \$6 billion, including all the preparatory work...the Taiwanese Semiconductor Manufacturing Company (TSMC) has built two “GigaFabs” for between \$8 billion and \$10 billion each.” From “Under New Management,” *The Economist*, April 2, 2009.
8. “Firms may spend up to \$40m before the first chip is made.” From “Kneading Chips,” *The Economist*, Nov. 18, 2010.
9. Ibid.
10. *Embedded Systems Design* surveys, 2005–2010.

11. "Network effect" Wikipedia entry
http://en.wikipedia.org/wiki/Network_effect
12. The evolution of the x86/Microsoft de facto PC standard is examined at
<http://www-siepr.stanford.edu/papers/pdf/00-51.pdf>
13. VHS vs. Beta "format war" is described at
http://en.wikipedia.org/wiki/Videotape_format_war
14. For more details on the Quartus II software, including compile time and feature set advantages, go to <http://www.altera.com/software>.
15. For additional details about Altera's 28-nm portfolio, visit:
http://www.altera.com/corporate/news_room/releases/2011/products/nr-28nm-portfolio.html