



User-Customizable ARM-Based SoCs for Next-Generation Embedded Systems

Author Abstract

Todd Koelling

Senior Manager, Embedded Products
Intel® Corporation

This white paper discusses Intel's programmable system-on-chip (SoC) approach to ARM-based embedded system implementation. The single-chip approach can be of particular value to embedded systems developers facing stringent time to market, cost, performance, design reuse, and longevity requirements.

Introduction

Today's embedded system developers face unprecedented challenges in their efforts to rapidly deliver competitive products to market. Until recently, most system implementation options have been limited to software-intensive, power-hungry multichip systems or costly SoC ASICs. However, market forces and resource constraints are compounding to make these approaches less viable for many design teams. For ARM-based embedded systems, though, advances in FPGA technology, intellectual property (IP), and design tools have witnessed the emergence of user- customizable SoCs. These devices not only overcome the shortcomings of traditional approaches, but offer unique and significant advantages for embedded system implementation.

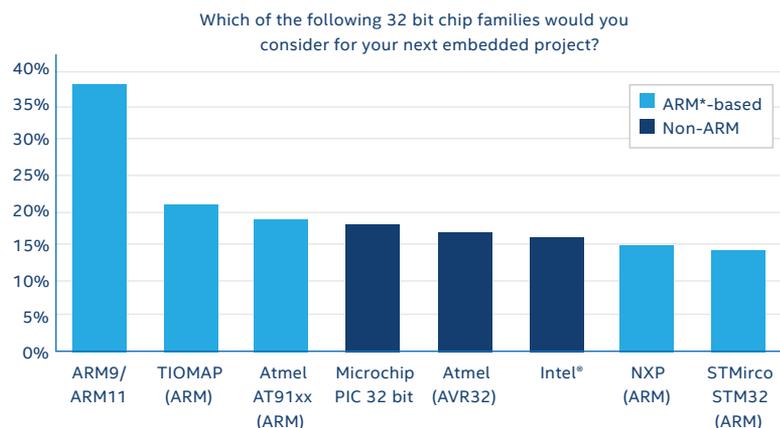
The Ubiquitous ARM* Processor

Only a few years ago, the processor market was heavily fragmented. PowerPC, RISC, MIPS, SPARC, and a number of other platforms competed for market dominance.

However, as the market matured and became more specialized, certain platforms gained dominance in specific application spaces. Nowhere is this more the case than in that of the ARM* processor in embedded system applications (Figure 1).

Table of Contents

- Abstract 1
- Introduction 1
- The Ubiquitous ARM Processor... 1
- Embedded System Design Challenges 2
- A New Type of SoC..... 2
- The SoC..... 2
- Getting to Market Faster 3
- Flexibility 3
- Scalability and Design Reuse..... 3
- Product Longevity..... 4
- The Intel SoC..... 4
- Rapid System Design Tools..... 5
- Software Development 5
- Conclusion..... 5
- References..... 5



Source: EE Times Group, Copyright 2010 by UBM/EE Times Group

Figure 1. Popular Platforms in Embedded System Applications

The ubiquity of the ARM processor in the burgeoning embedded systems market bodes well for designers. First, a well-established and growing ecosystem of software, development tools, and ARM-compatible devices gives them a toolbox of solutions to work with. Second, thanks to the economies of scale of the growing ARM market, new and more advanced system implementation options are becoming available.

Embedded System Design Challenges

More than ever, embedded systems developers must create cost-effective systems. The rapidly expanding global marketplace is demanding much more of designers. A growing number of highly viable opportunities—from new markets such as China, India, and Latin America—have become too significant to ignore. Addressing this broadening customer base requires support of a varied array of standards and price, performance, and feature platforms. In addition, competition is proliferating globally, placing even more pressure on design teams to deliver richly featured products within shrinking market windows.

Unfortunately, at the very time so much is being demanded of embedded products, design teams are shrinking. Economic pressure is forcing many companies to scale back, and design resources are part of the fallout. As a result, development groups are being downsized despite increased workloads.

Cost-conscious ARM-based embedded system designers are becoming more aware of the shortcomings of conventional implementation approaches. Multichip solutions are relatively easy to implement, but are costly and often lack the flexibility and the performance/power profile demanded of today's applications. Single-chip solutions that employ soft processor cores are also relatively easy to implement, but can fall short of power and performance objectives. ASIC SoCs with on-board hard ARM cores offer excellent power, performance, and optimization, but are slow to market, inflexible, and too expensive for the vast majority of applications. In order to be competitive, embedded systems developers need a solution that enables them to develop highly differentiated products with much greater flexibility and efficiency.

A New Type of SoC

The use of embedded processors in FPGAs has steadily increased over the past decade (Figure 2). Thanks to advances in FPGA technology from Altera, a new type of SoC device has emerged that offers the breadth of capabilities desired by today's embedded system applications. ARM-based SoCs combine a hard ARM processor, memory controllers, and peripherals with customizable FPGA fabric in a single SoC. These SoCs resolve many of designers' challenges, while yielding unprecedented product differentiation, price/performance optimization, time-to-market efficiency, and product longevity.

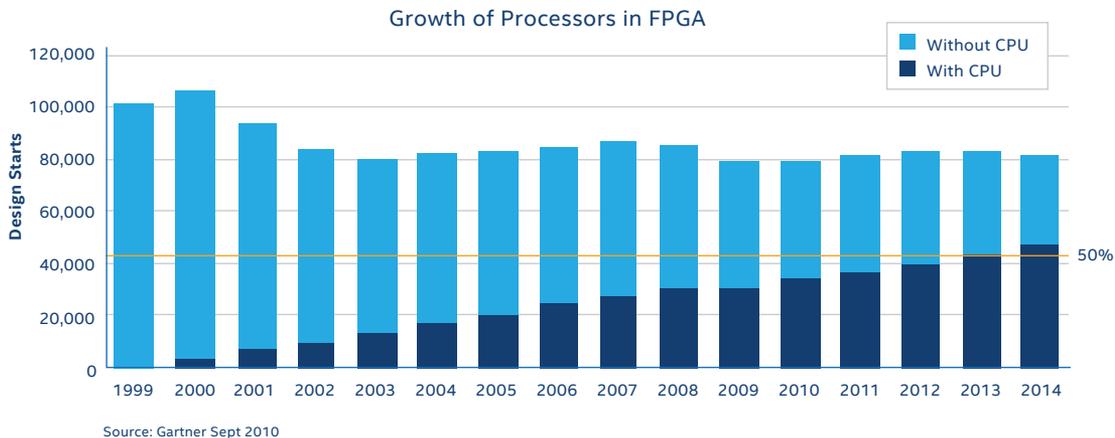


Figure 2. Growth of Processors in FPGAs

The SoC

The ARM-based SoC (shown in Figure 3) tightly couples a highly optimized “hard” processor system (HPS) with an on-chip FPGA. The HPS, which includes the dual-core ARM processor, multiport memory controller, and multiple peripheral elements, offers up to 4,000 MIPS (Dhrystone 2.1 benchmark) of processing performance for under 1.8 W. These hard IP blocks offer high performance while lowering power and cost, and free up logic resources for product differentiation. On-chip FPGA fabric can be customized by the user to create application-specific logic. Programmability allows for ready adaptation to new or changing communications standards, networking protocols, and performance tuning.

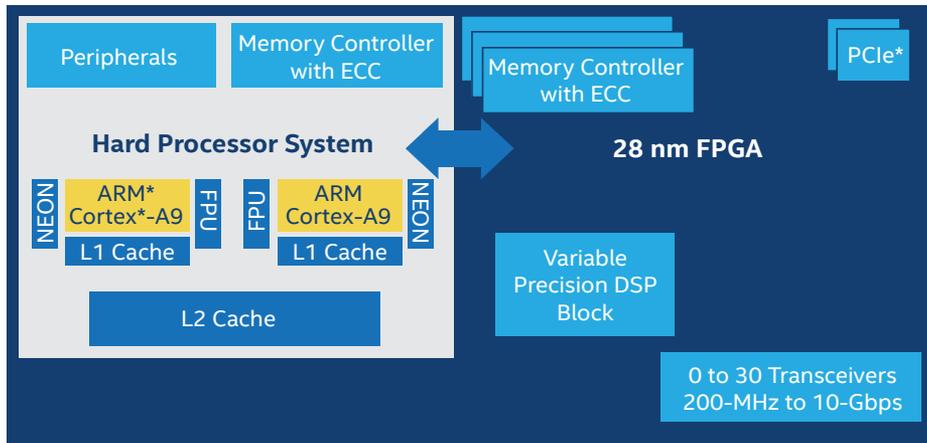


Figure 3. Processor Tightly Integrated with the FPGA Fabric

The SoCs offer significant performance advantages compared to traditional solutions. Hard elements are highly optimized compared to their soft IP counterparts, and thus deliver the best performance, lowest power, and highest density possible for the process node. Since FPGAs are often the first devices produced on new process nodes, designers utilizing SoCs gain access to the latest and best semiconductor technology available. The tight coupling of elements made possible by on-chip busing offers performance and power efficiencies over board-based solutions, as well. From an overall system standpoint, the SoC implementation significantly reduces overall system size, power, and cost.

Getting to Market Faster

Using a field-programmable platform, a custom ARM-based SoCs can be built using off-the-shelf devices in a fraction of the time and cost required of other custom devices. The combination of proven FPGA design tools, intuitive system integration tools, and the well-established ARM ecosystem help assure a speedy and low-risk development cycle. Even those new to FPGA implementation find that tools commonly available today support interface formats and standards such that they can leverage and reuse legacy software, IP, and other design content with relative ease. The SoC platform and supporting development infrastructure enables developers rapidly respond to embedded system market opportunities.

Flexibility

Competing in the embedded systems market requires agility. Networking and communications applications must adapt to new or changing standards. Industrial and medical providers may not necessarily require large production volumes, but still demand a high degree of product specialization. Regardless of the specific application space, embedded systems developers must have the flexibility to contend with intense competition, to take advantage of the latest technology advancements, and to quickly respond to new market opportunities.

SoCs are uniquely suited for the demands of the embedded market. Many changes and adaptations can easily be accommodated within the logic portion of an ARM-based SoC. In a SoC, this logic is field programmable, meaning changes to differentiate, update, or spin a product variant can be made quickly and easily. In cases that warrant, reconfiguration can even be made after a device has been deployed in the field.

Scalability and Design Reuse

Design reuse is a key means to manage costs, shrinking market windows, and limited design resources. The ability to scale and reuse design content across multiple devices and families can be invaluable. This technique minimizes the need to redesign from scratch when implementing an existing design or IP block in a device with different size, power, or performance requirements. For the growing number of geographically dispersed design teams, reuse not only boosts productivity, but also provides a method to share and leverage content.

The SoC architecture supports multiple IP cores and scalability to different device families. A broader range of applications with different price, performance, power, or other constraints can be more easily served with reusable and scalable content. As process technology advances, migrating to the next node is also greatly simplified within such FPGA families through porting IP to the next generation.

Product Longevity

For a large number of applications, product longevity is an important consideration. The product life cycle, particularly for medical, industrial, and military devices, can extend 10 to 15 years or longer. However, it is very common for IC providers to phase out products over a much shorter time frame. Designers of such products often must incur significant costs and dedicate precious design resources to migrate or respin when the product life cycle exceeds that of the IC components.

To avoid the negative consequences of product obsolescence or failure, developers of embedded systems with the potential for longer term deployment should seek implementation choices that have been rigorously tested for reliability and that are vendor supported over the expected product life cycle. FPGAs have a long history of being deployed in industrial, military, aerospace, automotive, and medical applications because of their longevity and ease of migration. Device longevity translates into less product maintenance in the field, thereby reducing maintenance costs and allowing the application of design resources to new products.

The Intel SoC

As part of its “Embedded Initiative” launched in 2011, Intel offers the 28-nm ARM- based SoCs for the embedded systems market. These Intel® products feature:

- Advanced hard ARM processor, peripherals, and a high-speed interconnect
- On-chip Intel FPGA fabric, secondary memory controller, and a PCI Express* (PCIe*) interface
- Design tools that support integration of legacy content and IP
- Simulation environment for accelerated software development with or without silicon

Silicon

At the 28 nm process node, Intel uses a tailored approach to address the spectrum of embedded applications by offering the ARM processor in the Arria® V and Cyclone® V FPGA families. Both families, produced on the same TSMC 28 nm low- power (28LP) process, offer the low system power and system cost required by embedded systems. The Cyclone V FPGA family is well suited for applications for which power or size are most important, while the Arria V FPGA family is tailored architecturally to address more performance-driven embedded applications.

The Intel ARM HPS (Figure 4) combines a dual-core ARM Cortex*-A9 MPCore* processor, memory controller, and peripheral IP in a full-featured HPS. The high- performance dual-core ARM Cortex-A9 MPCore processor operates at up to 800 MHz at the 28 nm process node. The dual-core configuration provides scalability that corresponds with the SoC product offering, as well as performance headroom for future needs. The built-in NEON* media processing engine and double-precision floating point unit provide standardized acceleration for multimedia and signal- processing applications. Two 32 KB Level-1 caches per core, backed up by a 512 KB shared Level-2 cache, help boost performance by minimizing latency and memory access time.

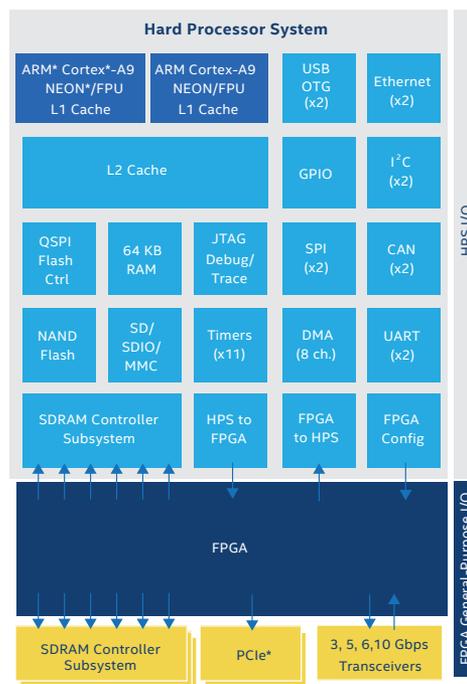


Figure 4. ARM-Based Hard Processor System Overview

In addition to the ARM cores themselves, the HPS includes a SDRAM controller subsystem, an array of general-purpose peripherals, and a high-speed, on-chip interconnect. The peripheral set includes a hard flash controller, MMC, DMA, USB 2.0, Ethernet, UART, SPI, and GPIO interfaces. Finally, Intel's unique on-chip busing architecture connects the HPS and FPGA with very high-speed interconnect with >125 Gbps total bandwidth. Application-specific logic is implemented in an on-chip Intel FPGA. Because the Cyclone V and Arria V FPGA families are well supported over the long term, Intel products can achieve greater than 20-year useful life under nominal operating conditions.

Rapid System Design Tools

The Intel Quartus® Prime design software provides an efficient design environment that enables developers to rapidly implement an ARM-based SoC. The included Platform Designer (formerly Qsys) saves significant time and effort in the design process by automatically generating interconnect logic between IP functions and subsystems. For IP written using common interface standards and protocols such as AMBA*, the Platform Designer (formerly Qsys) automatically recognizes the IP and connects it to the SoC. This tool makes it possible to easily reuse or mix and match legacy or third-party IP elements with different standard interfaces in a single SoC. In addition, developers can easily leverage existing content for rapid FPGA implementation.

Software Development

Software development can be a significant task when developing embedded systems. For this reason, embedded software developers can use a "virtual target" simulation environment to write, simulate, and debug software before silicon is available. The Intel SoC Virtual Target simulation environment supports register- and binary- compatible software development, even before first silicon is available. By completing most of the software development work in advance using the virtual target, designers can reduce risk and deliver products to market more rapidly.

Conclusion

Today more than ever, embedded systems developers must contend with intense competition, rapidly respond to changing standards, protocols, and requirements, and address an increasingly diverse marketplace with reduced resources. SoCs that include hard processor cores have emerged to empower designers to not only surmount these challenges, but to gain significant time-to-market, price/performance, differentiation, and product longevity advantages. A tipping point has been reached at which FPGA-based SoCs are now a viable and preferred approach over conventional solutions, and are staged to proliferate broadly in the market.

References

- SoC Overview: www.altera.com/devices/processor/soc-fpga/proc-soc-fpga.html
- Intel User-Customizable SoC FPGAs brochure: www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/br/br-soc-fpga.pdf
- Cyclone V FPGAs: Lowest System Cost and Power: www.altera.com/products/fpga/cyclone-series/cyclone-v/overview.html
- Arria V FPGAs: Balance of Cost, Performance, and Power: www.altera.com/devices/fpga/arria-fpgas/arria-v/arrv-index.jsp
- Dual-Core ARM Cortex-A9 MPCore Processor: www.altera.com/products/soc/tools_and_software/getting-started.html
- Platform Designer (formerly Qsys): www.altera.com/products/design-software/fpga-design/quartus-prime/features/qts-qsys.html
- Webcast: "Introducing the User-Customizable ARM-Based SoC FPGA": www.altera.com/education/webcasts/all/wc-2011-arm-based-soc-fpga.html
- Video: "Jump-Start Software Development with the SoC FPGA Virtual Target": www.accelacomm.com/acc/socfpga-wp/1/51451958/

