



Power Matters.™

Advantages of the 32-Bit Cortex-M1 Processor in Microsemi FPGAs

WP0073 White Paper

August 2015

Table of Contents

Introduction	3
An Industry-Standard Solution Designed for FPGA Implementation	4
A Unique Business Model—FREE	4
Cortex-M1 Features	4
Improved Code Density with Performance and Power Efficiency	4
Fast Memory Access Improves Performance	5
Efficient Interrupts Reduce Latency	6
Simplified Programmer's Model Eases Coding	6
Implementation in Microsemi M1-Enabled FPGAs	7
Cortex-M1 Tools	8
Conclusion	9

List of Figures

Cortex-M1 Block Diagram	3
Thumb-2 Performance and Code Density	5
Tightly Coupled Memory Signal Timings	6
Cortex-M1 Register File	7
Processor Tools	8

Introduction

The embedded market continues to move toward 32-bit processing. At the same time, the market has seen a significant increase in the use of FPGAs as flexible, cost-effective platforms for the rapid design of high-performance embedded systems. In combination, these trends are driving demand for 32-bit processors in programmable logic.

Certainly, there is no end to the number of new 32-bit processor architectures. However, most processor intellectual property (IP) is developed for ASIC implementation. As a result, when implemented in the coarse-grained architecture of FPGAs, the processors are often large and slow—a fate suffered by many widely used processors when ported to programmable logic. Of course, a few soft proprietary IP processor solutions are available for FPGA implementation. However, only limited tools, support, and designer experience exist for these proprietary solutions, making them harder and riskier to use.

What has been lacking in the market is an FPGA-optimized, 32-bit processor based on an industry-standard architecture. To address this need, Microsemi and ARM® developed the 32-bit Cortex-M1 processor (Figure 1), the first ARM core designed specifically for FPGA implementation. When combined with Microsemi nonvolatile M1-enabled ProASIC®3 FPGAs and Fusion Programmable System Chips (PSCs), the small, fast, and highly configurable Cortex-M1 processor offers a number of benefits. Because the Cortex-M1 processor is available from Microsemi free of license fees and royalties, the combination of a low-cost FPGA and the Cortex-M1 processor extends the ARM architecture to lower-volume applications, provides a lower-cost entry into system-on-chip (SoC) design, and shortens time-to-market for ARM users. For designs that scale to ultra-high volumes, the 32-bit Cortex-M1 processor runs the Thumb®-2 instruction set and is upward-compatible with the Cortex-M3 processor, providing an easy migration path to ASIC implementation. The industry-standard Cortex-M1 processor also enables economical reuse of tools, code, and knowledge, which reduces risk and gets products to market sooner.

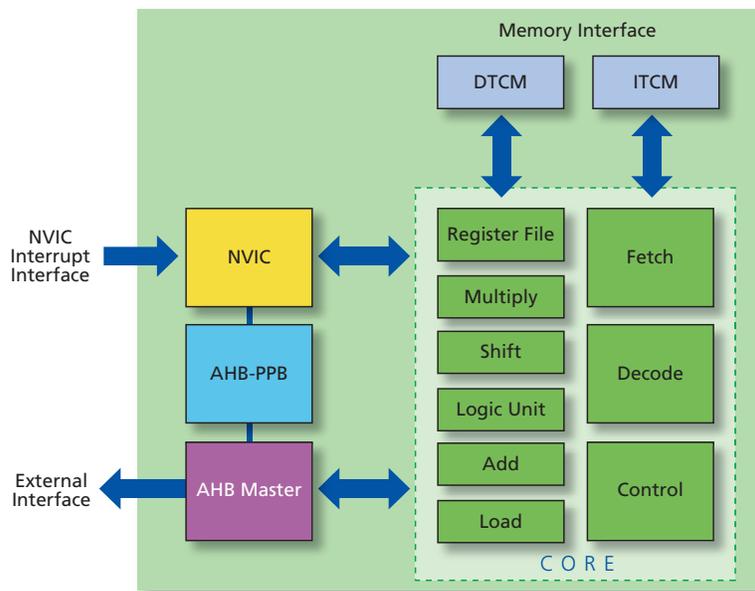


Figure 1: Cortex-M1 Block Diagram

An Industry-Standard Solution Designed for FPGA Implementation

The challenge associated with proprietary architectures is making them efficient in targeted applications and putting tools in place to support them. Experienced engineers know that there is a learning curve when using anything new; it takes time to climb the learning curve and gain the experience to effectively deal with the product's unique characteristics. This is in direct conflict with ever-shortening development schedules and increases design risk. For these reasons, designers tend to reuse what they are familiar with or what they have used before. Over time, this causes a few architectures to become widely used industry standards, while most are only used in narrow vertical niches.

When investigating which 32-bit processor would best suit customer needs in its flash-based FPGAs, Microsemi realized early on that an industry-standard architecture offers significant benefits over a proprietary architecture. Industry-standard processors have a broad selection of development tools, a significant volume of available program code, and a large following of design engineers who have knowledge and experience using them. These benefits enable users to get their designs developed faster and to market sooner while reducing risk, and as a result, offer customers a better solution and value.

A Unique Business Model—FREE

When Microsemi and ARM developed an optimized version of the ARM7™ processor for use in Microsemi FPGAs, it was offered with an innovative business model that dramatically increased industry access to the ARM architecture. By removing the license, royalty fees, and contracts typically associated with licensing models for industry-leading processor cores, Microsemi offers free access to advanced ARM processor technology for the broad marketplace. The same free delivery model is offered with Cortex-M1 for use in Microsemi flash-based, M1-enabled Fusion and ProASIC3 FPGAs. This provides all embedded designers with access to programmable flexibility and system-level integration with the ARM architecture, enabling the development of low-cost, high-performance systems.

Cortex-M1 Features

Derived from the ARM 3-stage Cortex-M3 processor pipeline, the highly configurable Cortex-M1 processor provides a good balance between size and speed for embedded applications. The core in its smallest configuration is less than 5 percent larger than the Microsemi Core8051, an industry-standard 8-bit controller. It runs at over 70 MHz in Microsemi M1-enabled ProASIC3 and Fusion devices. The processor runs a subset of the new Thumb-2 instruction set and features support for tightly coupled memory and a sophisticated low-latency interrupt controller to improve embedded performance and capabilities.

Improved Code Density with Performance and Power Efficiency

Thumb-2, a new instruction set for the ARM architecture, provides enhanced levels of performance, energy efficiency, and code density for a wide range of embedded applications. Thumb-2 technology builds on the success of Thumb, the innovative high-code-density instruction set for ARM microprocessors, to increase the power of the ARM microprocessor core available to developers of low-cost, high-performance systems.

The technology is backward-compatible with existing ARM and Thumb solutions, while significantly extending the features available to the Thumb instruction set. This allows more applications to benefit from the best-in-class code density of Thumb.

For performance-optimized code, Thumb-2 technology uses 31 percent less memory to reduce system cost, while providing up to 38 percent higher performance than existing high-density code, which can be

used to prolong battery life or to enrich the product feature set (Figure 2).

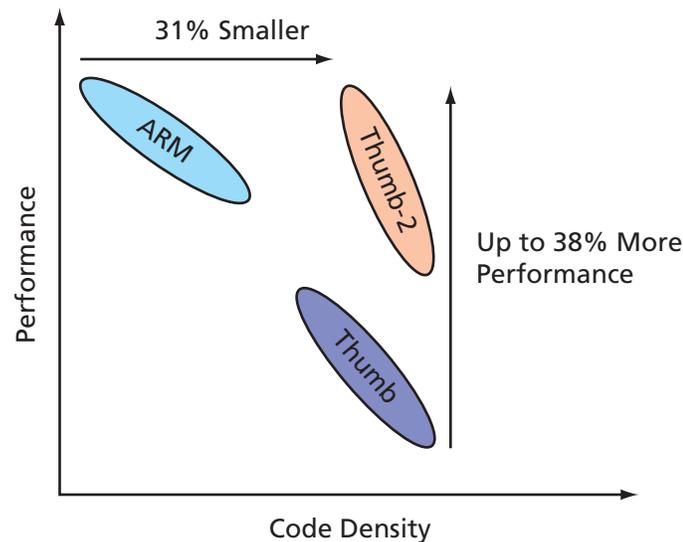


Figure 2: Thumb-2 Performance and Code Density

Cortex-M1 executes the ARMv6-M instruction set, which is a full subset of the Thumb-2 (ARMv7) instruction set that is used across the rest of the Cortex family. ARM recognizes the benefit of having a large volume of legacy code available for customers to use, so they made the Cortex family upward-compatible with Thumb code written for their legacy cores (ARM7, ARM9™ and ARM11™). Existing Thumb code can be run without change on Cortex-family processors, including Cortex-M1 devices. This is a big advantage for designers. Most of the important processing on legacy ARM processors was done in subroutines written in Thumb code, so designers can take advantage of their existing code for ARM processors.

One of the benefits of Thumb-2 over previous ARM instruction set architectures is that 16- and 32-bit instructions are executed in the same mode. In older ARM architectures, Thumb instructions were primarily used in subroutines with the 32-bit ARM instructions used to service interrupts. This often resulted in long latency between the time an interrupt was received and the time it was serviced. In Thumb-2, ARM merged the 16- and 32-bit operating modes so that interrupts could be serviced without the need to switch from 16-bit mode. It is a big advantage to be able to freely mix 16- and 32-bit instructions. This greatly simplifies the programming task and eliminates the need to profile the code to minimize code size and maximize throughput. It also results in increased performance, because the instructions can be optimally mixed without having to cluster them in 16- and 32-bit groupings.

Fast Memory Access Improves Performance

Cortex-M1 has a separate memory interface from the external AMBA (Advanced Microcontroller Bus Architecture) peripheral bus interface. This is similar to the high-performance ARM9 architecture but different from ARM7, which has a combined memory and peripheral bus. The separate memory interface on Cortex-M1 is actually implemented as two interfaces, giving separate access to instruction and data tightly coupled memory spaces (ITCM and DTCM). This increases the performance of the processor, because it can fetch an instruction from the ITCM on every clock cycle (Figure 3 on page 6), and it is never stalled due to data memory accesses, or reads and writes to the peripherals on the AMBA bus.

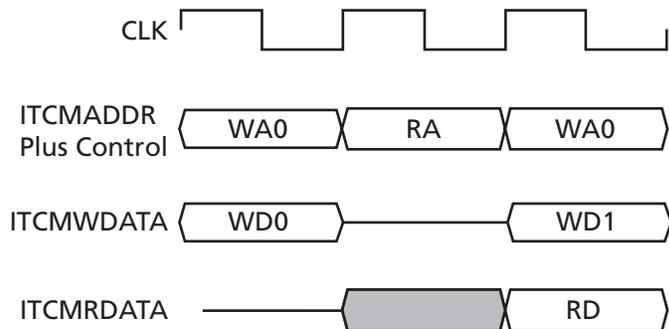


Figure 3: Tightly Coupled Memory Signal Timings

Efficient Interrupts Reduce Latency

A configurable Nested Vectored Interrupt Controller (NVIC) is available on Cortex-M1 that facilitates low-latency interrupt and exception handling and simplifies programming. The NVIC supports reprogrammable interrupts and is closely coupled to the processor core to support low-latency interrupt processing and efficient processing of late-arriving interrupts. In many applications, especially those that operate in real-time, low-interrupt latency is critical. The NVIC in Cortex-M1 also gives users the option of individually determining the priority of interrupts, supporting four interrupt levels, allowing designers to give a system event priority over other events. All of the capabilities built into the NVIC give users maximum control over how the processor responds to events as they occur in an application.

Simplified Programmer's Model Eases Coding

Cortex-M1 implements a subset of the Thumb-2 architecture with two operating modes—Thread mode and Handler mode. For normal processing, Thread mode is entered through reset or exception return. The Handler mode is entered as the result of an exception. The processor is designed for embedded applications where additional operating modes are not necessary. By limiting the programmer's model to a few operating modes, the size of the processor has been significantly reduced.

To ease programming and make transition between Thread mode and Handler mode as seamless as possible, the processor has two stacks ([Figure 4 on page 7](#)). Out of reset, all code uses the main stack. An exception handler can change the stack used by Thread mode from the main stack to the process stack by changing the value it uses on exit. The stack pointer, R13, is a banked register that switches between the main stack and the process stack. The processor has been architected to maximize the designer's control over the flow of data and processing through the core and to simplify programming, especially when different engineers do the software and hardware development.

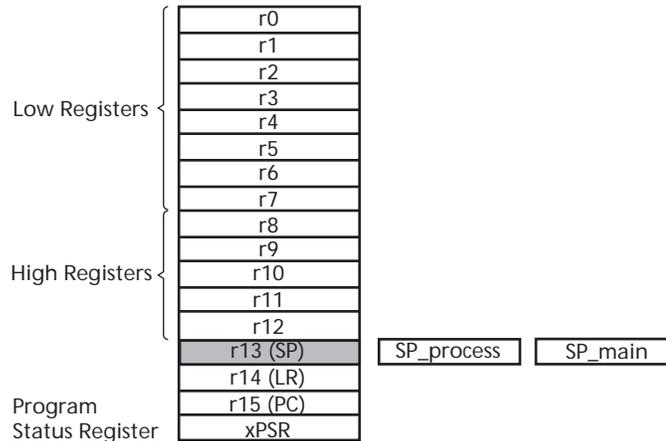


Figure 4: Cortex-M1 Register File

Implementation in Microsemi M1-Enabled FPGAs

FPGA usage is growing in part because of the flexibility that these devices offer. Engineers can tailor the function of a device exactly to their application by adding or removing soft IP. This is similar to what can be done with an ASIC. With FPGAs, however, a design can be developed and running in the application within a few hours, whereas ASICs require many months and large up-front, non-recurring engineering (NRE) charges.

When developing Cortex-M1, ARM and Microsemi made the processor highly configurable. The tightly coupled memory size, the size and speed of the multiplier, the number of external interrupts, the endianness, and whether the debug circuitry and OS extensions are included are all selectable by the user. This gives designers the control to select the minimum processor configuration that best meets their application requirements. Even better, because Cortex-M1 is being implemented in an FPGA, designers can quickly configure the core and program it into an M1-enabled, flash-based device and test it in their end application. If a change is required, it can be modified and reprogrammed into the device within minutes. In this way, engineers can modify and test their design many times within a few hours to find the optimal implementation for their product.

The M1-enabled Microsemi FPGAs allow seamless use of the Cortex-M1 processor core. The M1-enabled ProASIC3 and Fusion devices offer all the benefits of Microsemi nonvolatile, flash-based families—single-chip, reprogrammable, live-at-power-up, secure, firm-error immune, and low-power. The first M1-enabled devices offered by Microsemi are the M1AFS600 Fusion PSC and the M1A3P1000 ProASIC3. Additional devices will be supported in coming quarters.

The 600,000-system-gate M1AFS600 is a member of the world's first mixed-signal Fusion PSC FPGA family, which integrates a 12-bit analog-to-digital converter (ADC), as many as 40 analog I/Os, up to 8 Mbits of flash memory, and FPGA fabric—all in a single device. The Cortex-M1 processor can be implemented in as few as 4,300 tiles, less than 30 percent of the FPGA logic in an M1AFS600 device. When used in conjunction with a soft processor, such as Cortex-M1, the Microsemi Fusion™ device offers a powerful soft mixed-signal MCU platform.

The one-million-system-gate M1A3P1000 device offers 144 kbits of SRAM and 300 digital I/Os. ProASIC3 FPGAs are reprogrammable and offer fast time-to-market benefits at an ASIC-level unit cost. These features enable engineers to create Cortex-M1–based high-performance, high-density system applications using existing FPGA design flows and tools. The Cortex-M1 processor can be implemented in less than 20 percent of the FPGA logic in an M1A3P1000 FPGA.

Cortex-M1 Tools

When using a processor in an embedded application, the available development tools are the connection linking capability to implementation reality and a finished product. If the number of available tools is limited, as is the case with proprietary processors, the development effort will take more time and involve greater risk. Knowing this, ARM and Microsemi have made broad tool support a major part of the development of the Cortex-M1 processor. Both companies have spent years refining their tools and those of their third-party partners to offer a seamless and easy-to-use development environment that maximizes the designer's efforts while minimizing time-to-market.

To support an ARM processor in a Microsemi FPGA, tools are required for developing and debugging the programs that run on the processor as well as tools for developing and debugging the circuit that is programmed into the FPGA. A major benefit for users of an ARM microprocessor is the huge ecosystem of tools and design support as well as the large volume of embedded software programs that exist for it. To this ecosystem, Microsemi brings the CoreConsole IP Deployment Platform (IDP), the world-class Microsemi Libero[®] Integrated Design Environment (IDE) development tools, the SoftConsole program development environment, and a complete board-level development and debug environment (Figure 5). ARM is offering support for Cortex-M1 in both the RealView[®] Development Suite and the RealView Microcontroller Development Kit (MDK).

In addition to the tools available from ARM and Microsemi, the Cortex-M1 processor is supported by tools from Aldec, CriticalBlue, CodeSourcery, IAR, Impulse C,[™] Mentor Graphics,[®] Synplicity,[®] and others. Additional tools will be announced as they become available, significantly broadening the range of tools available for development of embedded applications with the Cortex-M1 processor.

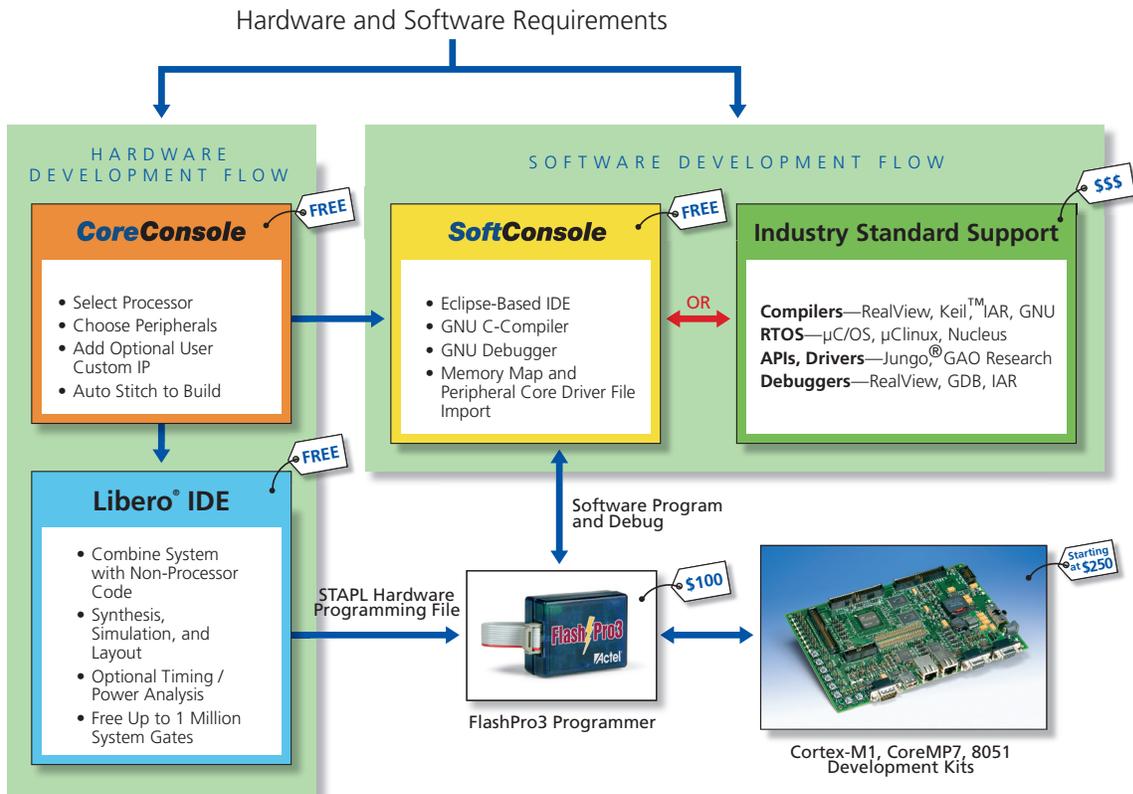


Figure 5: Processor Tools

Conclusion

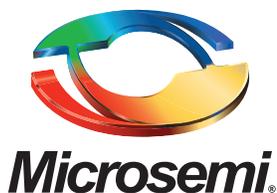
The movement from ASICs to FPGAs is driving the usage of 32-bit processors in FPGAs as embedded applications transition to programmable logic. By working together, ARM and Microsemi have developed an efficient FPGA-optimized 32-bit industry-standard processor, giving designers a powerful new solution for their embedded applications. The features of the Cortex-M1 processor—balanced three-stage pipeline, sophisticated interrupt controller, and tightly coupled memories—are targeted to give users maximum embedded performance while minimizing size and cost. Based on an industry-standard architecture, users are able to exploit the huge volume of existing code, extensive industry knowledge and support, and a vast ecosystem of tools. The free delivery of the Cortex-M1 processor for use in Microsemi flash-based, M1-enabled Fusion and ProASIC3 FPGAs provides all designers with a flexible and cost-effective platform for the development of low-cost, high-performance embedded applications.

List of Changes

The following table shows important changes made in this document for each revision.

Date	Changes	Page
Revision 1 (August 2015)	Non-technical Updates.	NA
Revision 0 (May 2007)	Initial Release.	NA

*Note: *The revision number is located in the part number after the hyphen. The part number is displayed at the bottom of the last page of the document. The digits following the slash indicate the month and year of publication.*



Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996

E-mail: sales.support@microsemi.com

© 2015 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi Corporation (MSCC) offers a comprehensive portfolio of semiconductor and system solutions for communications, defense & security, aerospace and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, Calif., and has approximately 3,600 employees globally. Learn more at www.microsemi.com.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.