Electronic Design

ALAN R. EARLS. Contributing Editor

Arm's Cortex-R Core Tightens Its Grip on Automotive-Grade Silicon

This article takes a deep dive into the architecture of Arm's real-time Cortex-R5 MCU and its relevance for the automotive industry.

Not every computer system can cut it under the hood.

These days, dozens of electronic control units (ECUs) can be spread out around a modern vehicle. Each unit generally only needs enough computing power to do a single task in areas ranging from body control to the powertrain. In many cases, these computer modules must be able to run safety-critical operations without interruptions. That means taking advantage of compact, real-time, and automotive-grade microcontrollers (MCUs).

Built on the same power-saving architecture used in everything from IoT devices to high-end smartphones, Arm's Cortex-R family of real-time CPU cores is becoming one of the main building blocks of modern cars.

Many of the largest purveyors of Arm Cortex-M MCUs also supply a wide range of chips based on real-time Cortex-R cores, including Infineon, Microchip Technology, NXP Semiconductors, Renesas Electronics, Silicon Labs, STMicroelectronics, and Texas Instruments. However, in recent years, these companies have been upping their game in real-time performance with safety-critical MCUs based on Arm's Cortex-R5 and its counterpart, the Cortex-R5F.

Some companies also offer multicore variants of the chips, including quad-core Arm Cortex-R5F MCUs. The real-time CPU cores inside them deliver sufficiently high performance for computing tasks under the hood and, of course, the real-time behavior that's critical in automotive uses. Its real-time computing power is also a strength for functional safety.

The Cortex-R5 builds on the basic features of the R4. They include enhanced error management, extended functional safety, and SoC integration features intended as a good basis for highly embedded real-time and safety-critical automotive systems.

Safety is delivered, according to Arm, by having a highly flexible second core that can be used for redundancy or improved performance. Arm said the Cortex-R5 also stands out from its better bus protection and error correction code (ECC).

Cortex-R5: Inside Arm's Real-Time CPU Architecture

The Cortex-R5 is based on what's called the "Harvard" memory architecture, which means that it features separate storage and signal paths for both instructions and data *(see figure)*. The real-time CPU core adds optional integrated instruction and data-cache controllers, while the amount of cache memory inside it is flexible. These caches can be configured for between 4 and 64 kB, while both write-back and write-through are options for cache lines.

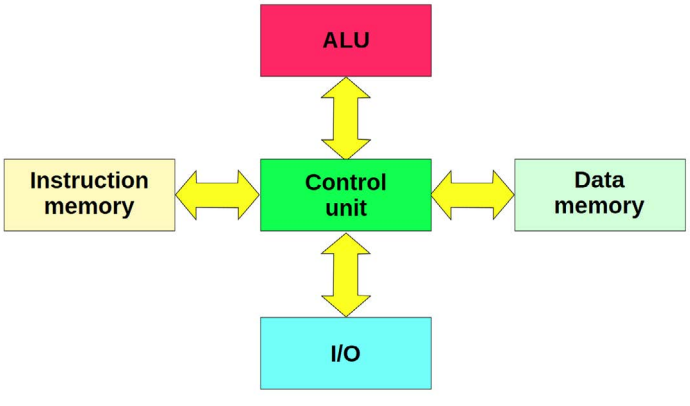

The Harvard architecture is a core feature of many MCUs, including Arm's Cortex-R5F MCU. Image credit: Neesa Ios, Creative Commons

The instruction set at the heart of the CPU core is the Thumb-2.

It folds 32-bit instructions into the 16-bit Thumb architecture and supports digital signal processing. The Cortex-R5 is based on an eight-stage pipeline microarchitecture with instruction pre-fetch, branch prediction, and selected dual-issue execution. Parallel execution paths are also part of the package to run MAC, shift-ALU, load-store, divide, and floating-point (FP) operations.

The Cortex-R5 contains up to a pair of tightly coupled memories (TCMs). The TCMs are dedicated, low-latency, on-chip memory areas directly connected to the processor core, offering faster access compared to standard external RAM and often used for critical code sections such as interrupt handlers and/or real-time workloads where deterministic timing is vital. The A and B TCMs can be used for any combination of code and data and can be configured up to 8 MB. TCM B is also provisioned with two physical ports, B0 and B1. This provides interleaving of incoming direct-memory-access (DMA) data streams.

Highly deterministic or low-latency workloads can take advantage of the optional TCM interface in situations such as instruction code for interrupt service routines and data requiring intense processing that otherwise might not respond well to caching.

The Other, Optional Building Blocks of the Cortex-R5 CPU

The Cortex-R5 includes an optional MPU that can be set up to use either 12 or 16 regions. In turn, these regions (which can overlap) have resolution down to 32 bytes. The highest numbered region is granted the highest priority.

Another optional aspect of the R5 is a floating-point unit (FPU). This implements the Arm Vector Floating Point architecture VFPv3, which features 16 double-precision registers. This complies with IEEE 754. Optimized performance for the FPU is provided for single-precision calculations and there's optional support for double precision.

According to Arm, the FPU can handle operations including add, subtract, multiply, divide, multiply and accumulate, square root, conversions between fixed- and floating-point, and floating-point constant instructions.

Other important facets of the real-time CPU relate to the Accelerator Coherency Port (ACP), the dual-core design, and debug capabilities. The 64-bit AXI subordinate port inside the CPU helps establish coherency between the processor(s) and intelligent peripherals, ranging from DMA controllers to Ethernet and even FlexRay-type interfaces. The dual-core processor configuration can work in coordination (for instance, for fault tolerance) or running separate programs. Also, a Debug Access Port is incorporated.

Interrupt handling is one of the keys to real-time processing. The Cortex-R5 is specifically designed for a standard interrupt (IRQ) as well as non-maskable fast interrupt (FIQ). The GIC interrupt controller is placed inside to handle more complex, priority-based interrupt handling. Low-latency interrupt technology is baked directly into the processor. This allows for long multi-cycle instructions to be interrupted and restarted. Lengthy memory accesses can also be deferred when necessary.

Where Real-Time CPU Cores Fit into the Automotive Industry

The Cortex-R5 represents a step up for real-time computing and overall performance compared to Arm's Cortex-M family of MCUs. The Cortex-R5 also manages to combine improved energy efficiency and real-time responsiveness. But the most telling factor may be that it was designed with input from multiple players in the automotive market, helping it meet the parameters the industry seeks in terms of performance, reliability, and real-time response times.