



ARM Cortex-R4(F)

Embedded Processors for Real-Time Applications

The ARM Cortex™-R4 processor is the first deeply embedded processor to be based on the ARMv7 architecture and is targeted at very high-volume, deeply embedded applications such as hard-disk drives, inkjet printers, and automotive safety systems.

The ARM Cortex-R4 processor provides key savings in cost and power consumption for system developers, offering substantially higher performance than any other processor with similar die size.

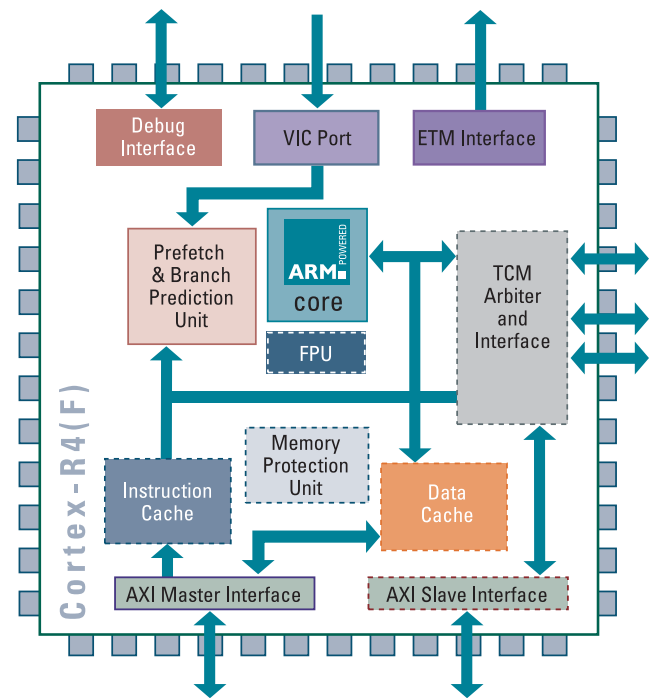
Along with the ARM1156T2-S and ARM Cortex-M3 processors, the ARM Cortex-R4 processor completes comprehensive coverage for the diverse needs of the embedded microprocessor market. Furthermore, the ARM Cortex-R4 processor supports substantial synthesis time configurability that enables designers to match the processor precisely to the application requirements.

In addition to the ARM Cortex-R4, ARM has introduced the ARM Cortex-R4F, which contains a Floating Point Unit (FPU). The ARM Cortex-R4F processor's FPU performs floating-point calculations that allow a greater dynamic range and accuracy than fixed-point calculations. The FPU is backward compatible with earlier ARM FPUs (VFP9/10/11), and is optimized for the single-precision processing most commonly used in automotive and control applications. The FPU is particularly useful in sophisticated control applications, where algorithms are often modeled in an environment such as Simulink or ASCET-SD, and code is auto-generated using tools such as Real Time Workshop Embedded Coder, ASCET-SE, or dSPACE Targetlink.

The ARM Cortex-R4 processor is capable of running at clock speeds of up to 400 MHz on typical 90 nm processes, and the focus throughout the design is on efficiency and configurability.

Technical Innovations ▶

- Thumb®-2 technology; an innovation that has enabled partners to combine the minimal memory footprint of 16-bit Thumb code with the high performance of 32-bit ARM code
- AMBA 3 AXI protocol; a set of major enhancements to AMBA for high-performance on-chip interconnect, the ARM Cortex-R4 processor integrates a 64-bit master port as well as a 64-bit DMA port for direct access to the Tightly Coupled Memories (TCM)
- A selective superscalar eight-stage pipeline that provides more than 1.6 DMIPS/MHz in an efficient low gate count implementation
- Non-Maskable Interrupts (NMI); many real-time applications demand this and the ARM Cortex-R4 supports a configurable NMI pin ARM Cortex-R4(F)
- CoreSight™ technology; a framework for complete system debug and trace; this includes the ETM-R4 embedded trace macrocell and many other CoreSight components



- A significantly improved local memory architecture for TCM and DMA; TCM can now be unified into a single logical address space and can run as fast as cache memory
- Enhancements over the ARMv6 architecture include improvements in interrupt handling and the memory protection scheme; new instructions for managing interrupts reduce the critical early-interrupt handler code, and the worst-case interrupt latency is vastly improved to only 20 clock cycles
- Performance monitoring support; very useful for refining and tuning a system through advanced profiling of the system performance
- Architected support for parity in the caches and parity or ECC in the TCMs; soft errors are an increasing concern in embedded systems and either parity or ECC is now essential in many systems

- A very efficient branch prediction and prefetch unit provide a branch accuracy of more than 90 percent for typical C code
- The overall aim of the ARM Cortex-R4 processor is to provide around 40 percent more efficiency than the ARM9 family whilst increasing the maximum clock speed, supporting the use of low-power, dense RAMs for cache and TCMs, and delivering an efficient Thumb-2 engine

Architectural Features ▶

The ARM Cortex-R4 processor's sophisticated pipeline architecture is based on low-cost dual-issue pipeline, eight stages with advanced dynamic branch prediction achieving 1.6 DMIPS/MHz; the ARM Cortex-R4 processor is fully ARMv7 architecture-compliant and includes:

- Thumb-2 technology for greater performance, energy efficiency, and code density
- Hardware divide instructions for control applications
- Optimized level-one caches and TCM
- Synthesis optional cache controllers (with optional cache parity) and TCM ports for flexibility
- Full wait and error support on TCM interfaces
- Flexible configuration at synthesis time of major level-one features
- A Memory Protection Unit (MPU) can be removed or an eight- or 12-region one selected
- Either one, two, or three TCM ports can be included
- A number of breakpoints and watchpoints can be selected

- Dynamic Branch Prediction
 - Enabled by branch target, global-history buffers, and a function called return stack
 - Achieves 90 percent accuracy across industry benchmarks
- Single-cycle load-use penalty for access to the L1 cache and TCM
- A single 64-bit AXI master port for easy integration into the SoC interconnect
- An AXI slave port to allow direct access to TCMs by DMA controllers and other processors in the system
- Vectored Interrupt Controller (VIC) port for fast connection to interrupt management peripherals

Core area, frequency range, and power consumption are dependent on process, libraries, and optimizations. The numbers quoted above are illustrative of synthesized cores using general-purpose TSMC process technologies and ARM Artisan standard cell libraries and RAMs.

The speed-optimized implementations refer to the library choices and synthesis flow decisions and tradeoffs made in order to achieve the target frequency performance. The area-optimized implementations refer to the library choices and synthesis flow decisions and tradeoffs made in order to achieve a target area density.

The cache sizes are specified as InstructionCache/DataCache. The area without cache numbers quoted exclude RAM area, but include all logic including memory management, cache control, and debug. The area with cache numbers quoted includes the core, the specified instruction and data caches and all necessary RAMs.

Performance Characteristics ▶	0.13 μm	90 nm	
	Area Optimized	Speed Optimized	Area Optimized
Standard cells	SAGE-HS	Advantage-HS	Metro
Memories	HS	Advantage	Metro
Frequency* (MHz)	300	500	210
Area with cache (mm ²)	3.35	2.50	1.50
Area without cache (mm ²)	1.99	1.66	0.80
Cache size	16K/16K	16K/16K	16K/16K
Power** with cache (mW/MHz)	-	0.41	0.22
Power** without cache (mW/MHz)	-	0.33	0.16

*Worst-case conditions—0.18 μm process—1.62V, 125°C, slow silicon; 0.13 μm process—1.08V, 125°C, slow silicon; 90 nm process—0.9V, 125°C, slow silicon

**Typical-case conditions—0.18 μm process—1.8V, 25°C, typical silicon; 0.13 μm process—1.2V, 25°C, typical silicon; 90 nm process—1V, 25°C, typical silicon

For more information on ARM certified technical training courses, visit www.arrow-nac.com/atc.