

Arm's new Cortex-R82 is its first 64-bit real-time processor

Enlarge / The code in the background of this image is a bit too high-level for a typical RTOS... but it might run under Linux, on the MMU side of a Cortex-R82 real-time CPU.

reader comments

16 with 15 posters participating

Share this story

Last Friday, Arm announced the newest generation of its real-time processor series, the Cortex-R82. Most people are more familiar with the Cortex A-series of CPUs, which are used as the primary processors in devices such as smartphones and tablets, running full-fledged operating systems. The Cortex-R series, by contrast, is typically used for high-performance applications demanding “real-time” performance—meaning simple, predictable, and extremely low-latency response loops—in much simpler software stacks.

Earlier Cortex R-series processors weren't capable of running full-fledged modern operating systems, including Linux, because they utilized a simple Memory Protection Unit (MPU) rather than the more complex Memory Management Unit (MMU) needed to support functions such as virtual memory. This generally wasn't seen as a problem, since Real-Time Operating System (RTOS) workloads generally need to have much simpler, more predictable control loops.

Real-time versus multitasking

Enlarge / The Cortex-R82's cores can be dynamically assigned to either real-time or applications processing—the same device might have a different profile during the day than it does overnight.

The Cortex R-82 still offers a simple MPU, but it can be optionally configured with an MMU as well—and the CPU's cores can be individually, and dynamically, assigned to either. Arm's Neil Werdmuller hypothesizes storage controllers that might operate with different profiles during peak and off-peak hours, reassigning cores from real-time “pure SSD” duties to “computational storage”—likely meaning onboard AI analysis—as needed.

Flexible design

When used for cache-friendly applications such as storage controllers—in which Arm says the R-series has 85 percent of the market—the 32-bit architecture of previous designs placed sharp performance limits. Modern SSDs routinely feature 2GiB onboard DRAM, which is uncomfortably close to the 4GiB theoretical maximum address limit of a 32-bit CPU.

The R-82's 64-bit architecture brings it inline with modern mainstream processors, enabling it to address a physical address space up to 1TiB. The new architecture also brings support for optional Neon units for Single Instruction, Multiple Data (SIMD) processing, allowing for greater flexibility in storage designs.

Finally, Cortex-R82 designs with the optional MMU will be capable of running full-stack operating systems such as Linux, as well as—and even in parallel with—simpler RTOS. This will allow reuse of existing higher-level code developed for full Linux environments in what traditionally have been much more restrictive environments.

Higher performance

Enlarge / The new processor shows modest improvements in generic benchmarks—and even better results for applications

using the new Neon SIMD instructions.

The new CPU design brings traditional performance uplift as well as architectural changes. In the chart above, Arm shows significant improvement in generic benchmarks, with massive improvement in tests utilizing the new Neon SIMD instructions. The “Consumer” column is the EEMBC Consumer benchmark, which takes advantage of the SIMD instructions when compiled with the Arm Compiler 6.14; Customers 1 through 4 are unspecified Arm customers’ real-world applications.

<https://arstechnica.com/?p=1704541>