

# MICROPROCESSOR *report*

Insightful Analysis of Processor Technology

## SiFIVE U8 TAKES RISC-V OUT OF ORDER

*U84 Is First in New Series of High-Performance RISC-V CPUs*

By Bob Wheeler (October 28, 2019)

Continuing its annual cadence, SiFive disclosed its highest-performance RISC-V core to date at the Linley Fall Processor Conference. The U8-series builds on the 64-bit superscalar U7-series by adding out-of-order (OOO) execution, a first for commercial RISC-V intellectual property (IP). The first core in the new series is the low-power U84, which includes a double-precision floating-point unit (FPU) and offers the same multicore configurations and Linux compatibility as the U74. In mid-2020, SiFive plans to release the U87, its first CPU to sport a vector unit.

The company designed the U84 to target the same performance as Arm’s Cortex-A72, but the latter CPU is about five years old. It overdesigned portions of the U8 architecture, however, to support higher-performance variants that will follow the U84. For example, the U8 front end can decode up to four instructions per cycle (IPC), but the U84 instantiates only three decoders. The U84’s 12-stage pipeline achieves an estimated 2.3GHz in a fully synthesized 16nm design and typical process corner. Like the U74, the U84 omits private L2 cache, relying instead on an L2 shared by all cores.

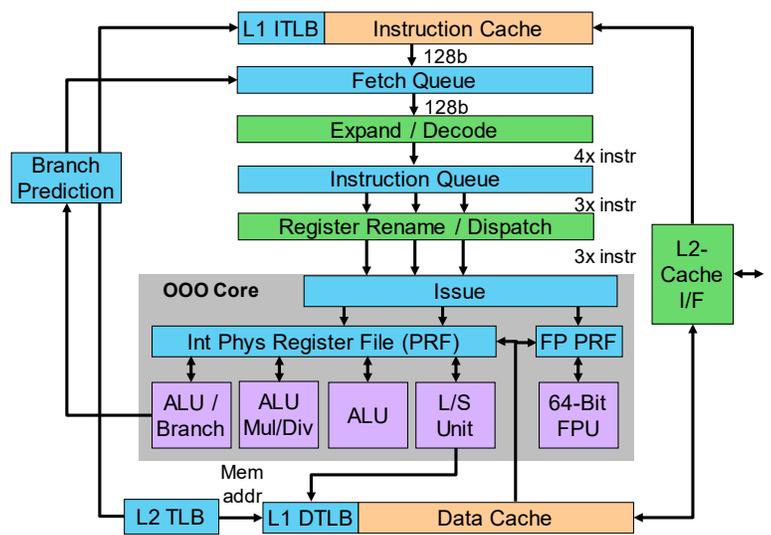
The U84 slots in between Arm’s newest little and big Cortex CPUs in per-clock performance. Customers who want the best single-thread performance in the newest process node can select Cortex-A77. If they need better area efficiency, however, the U84 delivers an attractive compromise.

### Marching Three by Three

Although SiFive plays the David to Goliath Arm, it’s much bigger than newer companies that the RISC-V ecosystem has been hatching. SiFive has grown rapidly, separating it from the RISC-V field. In June, the company raised \$65 million and welcomed Qualcomm as a

new strategic investor. It now claims 130 design wins supported by 550 employees across 16 offices, including about 200 employees from its 2018 acquisition of Open-Silicon. In addition to CPUs, the company now offers peripheral IP including an HBM2E controller, 400Gbps Ethernet MAC, USB3.1 controller, and DMA controller. In April, it taped out a 7nm chip to validate some of these cores along with its U7 and S7 CPUs.

The U84 fits the same multicore SoC “chassis” as the U74, SiFive’s first superscalar CPU (see [MPR 11/12/18](#), “SiFive Raises RISC-V Performance”). It’s a 64-bit design that implements the RV64GC instruction set, meaning it has a double-precision FPU and handles compressed (C extension) instructions. The company designed the U8-series to



**Figure 1. U84 microarchitecture.** The CPU can fetch and decode four RISC-V instructions per cycle and has three integer execution units.

### Price and Availability

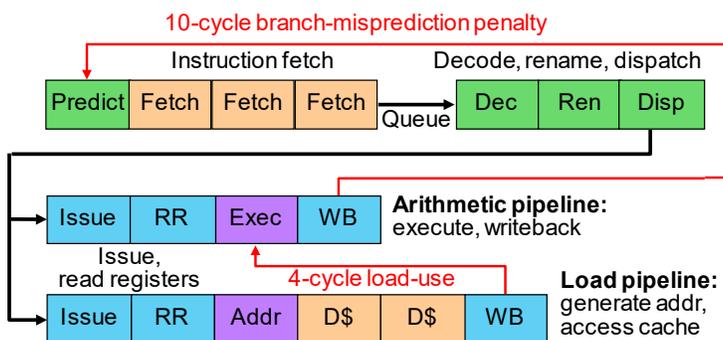
RTL for the U84 is available now to lead customers; general availability is expected in 2Q20. SiFive withheld U8-series license fees. For more information, point your browser to [tinyurl.com/y4qb77eg](https://tinyurl.com/y4qb77eg).

be highly configurable, including variable fetch, decode, and dispatch widths. As Figure 1 shows, the U84's front end fetches and decodes four instructions per cycle. The decoders expand compressed (16-bit) instructions into 32-bit equivalents. Like the U74, the U84 features two-level branch prediction using a branch target buffer (BTB) and branch history table (BHT), both with configurable sizes.

The instruction queue feeds three instructions per cycle to the register-rename function, which renames logical integer registers to physical registers. A reorder buffer (ROB) tracks the state of instructions in the out-of-order core, such as by maintaining a program counter for each instruction. The flexible-issue queue holds dispatched operations and can issue integer instructions to the FPU when its pipeline is empty.

The U84 has three integer execution units, a load/store unit, and one floating-point unit (a second FPU is optional). One integer unit includes an ALU and handles branches, a second handles multiply and divide, and the third is a simple ALU. The load/store unit handles data-memory access through the level-one translation lookaside buffer (DTLB). The FPU performs 32- and 64-bit math using a pipelined multiply-accumulate (MAC) unit, iterative divider, and various other units. Lacking SIMD, it completes one MAC operation per cycle regardless of precision. Customers can optionally configure the U84 with a second FPU.

Figure 2 shows the U84 integer pipeline, which has 10 stages for arithmetic operations and 12 stages for load/store. SiFive added a fetch cycle relative to the U74 to make the cache nonblocking; the U84 includes miss-status holding registers to handle multiple outstanding cache misses ("miss under miss"). The new core also implements prefetching, a



**Figure 2. U84 pipeline.** The CPU requires 10 stages for basic integer instructions and 12 stages for a load instruction that hits the data cache.

first for SiFive. Following the fetch queue, instructions spend one cycle each in the decode, rename, and dispatch stages. Once an operation is queued for a particular execution unit, the scheduler for that unit waits until all operands are available before issuing the operation.

Whereas the in-order U74 includes secondary ("late") ALUs to eliminate load-to-use latency, the U84 always executes these operations as early as possible. Thus, an instruction in the arithmetic pipeline will suffer a minimum four-cycle load-to-use penalty if it depends on a load in the memory pipeline. The branch-misprediction penalty is 10 cycles.

SiFive designs its cores using a high-level tool called Chisel, which was developed at UC Berkeley. Because Chisel automatically generates RTL, the company can more rapidly iterate cores. Thus, if a customer requires higher clock speeds, it could create a future U8-series core with a deeper pipeline than the U84. The U84 represents the low end of likely U8-series designs, but SiFive withheld its roadmap beyond the vector-enabled U87.

### Building Out an SoC

Like the U74, the U84 supports 512GB of virtual-address space using a 39-bit (SV39) memory-management unit (MMU), which comprises an L2 TLB and table walker. It also has a physical memory-protection unit (PMP) with eight regions. The CPU connects with the L2 cache and the rest of the SoC through a cache-coherent TileLink inner bus. (SiFive says private L2 cache is on its roadmap for future U8-series cores.) It supports multicore configurations with up to nine CPUs, potentially including a mix of application and microcontroller cores. For example, SiFive can instantiate four U84s for Linux applications plus one small 64-bit S21 for system management. It developed the S21 after learning that some customers preferred 64-bit over 32-bit microcontroller cores owing to their better software consistency.

The application cores share the directory-based L2 cache, which can be up to 8MB and is ECC protected. The U8-series introduces cache stashing, thereby enabling an external agent to write data directly into the L2 for faster processing. The multicore design offers local and global interrupts; a core-local controller (CLIC) generates the former, meaning they're associated with that core's thread (which RISC-V calls a hart). A platform-level interrupt controller (PLIC) distributes up to 132 global interrupts with seven priorities.

The multicore complex attaches to the SoC through memory, system, peripheral, and front ports that implement the TileLink coherent interconnect. The L2 cache's outer port handles cached TileLink, whereas the other ports aren't cache coherent.

### They're Out of Order

SiFive aims the U84 at a broad range of embedded applications in wireless infrastructure, autonomous

vehicles, and home entertainment. The area-efficient U84 requires only 0.28mm<sup>2</sup> in 7nm technology, whereas Arm’s Cortex-A76 is about 0.9mm<sup>2</sup> (excluding L2 cache). The A76 includes SIMD (Neon), however, which the U84 lacks. SiFive expects the U84 will clock at up to 1.8GHz in 28nm and 2.6GHz in 7nm. As Table 1 shows, the U84 trails the A76 by about 27% in SPECint per gigahertz, and Arm’s deeper pipeline enables about 15% higher clock speeds (see [MPR 6/4/18](#), “Cortex-A76 Revamps Core Design”). SiFive claims the slower U84 is 50% more efficient than Cortex-A72 but provided no power estimates.

The RISC-V design most similar to the U84 is Esperanto’s Maxion, which is also a 64-bit superscalar OOO CPU (see [MPR 12/10/18](#), “Esperanto Maxes Out RISC-V”). That company started with the open-source Berkeley Out-of-Order Machine (Boom), instantiating a front end that decodes four instructions per cycle. Still, its delivered SPECint per gigahertz is about 10% lower than the U84’s. Esperanto plans to sell chips rather than license IP, however, so it may be more conservative in its performance estimates.

Smaller RISC-V IP vendors are emerging, and two have discussed OOO CPUs, likely also starting with Boom as a baseline. Syntacore (Cyprus) offers a range of 32-bit CPUs and is adding 64-bit CPUs this year. Its initial SCR7 configuration decodes only two instructions per cycle, however, so

its IPC falls below that of the U84. CloudBear (Russia) sells the BI-671, which also implements a two-wide front end and targets lesser IPC than the U84. Both CloudBear and Syntacore claim better-than-Boom performance, adding value over open source.

### Know Your Audience

Because tapeout costs are rising in each new process generation ( $n$ ), a growing number of designers will employ  $n-1$  or even  $n-2$  nodes for their SoCs, particularly for lower-volume ones. SiFive can serve these customers by offering area-efficient cores that also carry lower licensing fees than similar Arm cores. The U84 extends its reach into high-performance SoCs while maintaining a consistent 64-bit instruction set even in its low-performance microcontroller cores. The company has additionally expanded its IP portfolio to include memory and network controllers unavailable as open source.

The RISC-V ISA presents some pitfalls for SiFive and its customers. It remains immature, with several extensions (including vector/SIMD) missing or in draft form. Absent specifications, vendors have developed their own extensions, leading to ISA fragmentation. Andes, for example, has implemented DSP extensions that it has also proposed for standardization (see [MPR 4/15/19](#), “Andes Strengthens Its RISC-V Arsenal”).

As the community grows and established designs become entrenched, consensus becomes more difficult, slowing progress. SiFive must now balance customer requirements with RISC-V stewardship. Sophisticated customers may choose to design their own CPU, starting with open source. Disclosed examples include Esperanto and Western Digital (see [MPR 2/4/19](#), “WD Rolls Its Own RISC-V Core”).

Car companies have long understood the value of a “halo” product to promote their brand—Chevy’s Corvette is a classic example. These products have low volumes but show off cutting-edge technologies. In the near term, the U8-series will likely serve the same function for SiFive. It shows off RISC-V’s performance potential in a midrange application-class CPU. It’s also a foundation for future high-performance cores. Finally, it keeps SiFive well ahead of others offering RISC-V CPUs. ♦

	SiFive U84	Esperanto Maxion	Arm Cortex-A76
Instruction Set	RISC-V	RISC-V	Arm v8
Architecture Size	64 bits	64 bits	64 bits
Extensions	Compression	Compression	Comp., SIMD
Max Decode	4 instr	4 instr	4 instr
Load/Store Units	1 L/S	1 L/S	2 L/S
Integer ALUs	3 ALUs	3 ALUs	3 ALUs
FP/SIMD Units	2x 64 bits	1x 64 bits	2x 128 bits
Reorder Window	Undisclosed	128 instr	~160 $\mu$ ops
Pipeline Length	12 stages	12 stages	17 stages
Max Speed*	2.6GHz	2.0GHz+	3.0GHz‡
SPECint/GHz†	6.9	6.2 (est)	9.4‡
SPECint (est)†	17.9	12.5	28.1‡
Die Area*	0.28mm <sup>2</sup>	<0.50mm <sup>2</sup> ‡	0.90mm <sup>2</sup> ‡
RTL Availability	4Q19	Not applicable	2017

**Table 1. Comparison of high-performance CPUs.** The SiFive U84 can’t match Cortex-A76’s IPC, and its shorter pipeline results in a lower top speed. \*In TSMC 7nm, excluding L2 cache; †SPECint\_2006. (Source: vendors, except ‡The Linley Group estimate)

To subscribe to *Microprocessor Report*, access [www.linleygroup.com/mpr](http://www.linleygroup.com/mpr) or phone us at 408-270-3772.