RM7000 Strong on Cost/Performance

First QED-Branded Chip Targets Low-End Workstations, High-End Embedded

by Brian Case

With its RM7000 MIPS-based microprocessor, QED continues its tradition of staying one step behind the industry in processor-core complexity to achieve industry-leading metrics in other areas. In his presentation at last week's Microprocessor Forum, John Kinsel described the new memory-heavy chip with a total of 288K of cache: split 16K L1 caches and a 256K L2 cache.

The new chip combines low cost with strong performance. QED claims the chip, when samples are available sometime in the second half of 1997, will boast a modest 80-mm^2 die size in a 0.25-micron CMOS technology, an operating frequency of 300 MHz, and performance of greater than 10 on both SPECint95 and SPECfp95 (base).

The RM7000 moniker stands for RISCMark7000. QED chose the RM designation to supplant the anonymous R prefix used in previous MIPS-architecture part numbers because it gives the company a simple way to distinguish its MIPS implementations from those of other vendors.

One reason QED sought a distinguishing trademark is because the company will, for the first time, become a chip vendor. Previously, QED was simply a small, nimble design house; with the introduction of the RM7000 and its derivatives, the company will purchase chips from foundries and sell them directly through its own sales force. This decision marks a dramatic change in QED's business focus, possibly signalling a desire to have more control over its destiny and to participate more materially in the fruit of its labor.

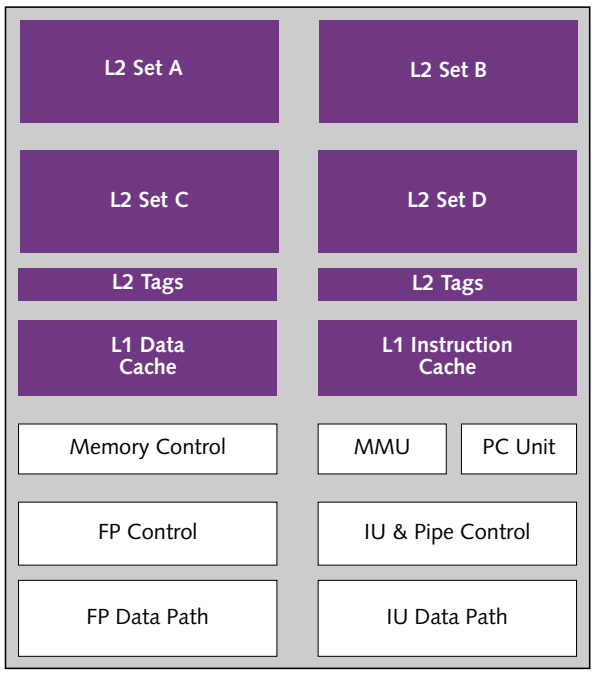

Figure 1. This RM7000 die plot shows that, as with QED's R4600 and R5000, the die area of the new design is dominated by cache. Note that 256K of unified cache is squeezed into less than half the 80-mm^2 die area.

RM7000 Continues QED's Design Philosophy

The evolution of QED's MIPS microprocessors tells an interesting tale. Beginning with the R4600 "Orion" chip *(see* **061507.PDF***)*, QED established a design philosophy resulting in chips that the company characterized as "an SRAM with an attached processor, not a processor with an attached SRAM." The goal of the philosophy is to achieve an efficiency "sweet spot" by leveraging semiconductor technology. In addition to emphasizing on-chip cache, QED leverages process technology by minimizing the complexity of the processor core.

QED's application of this philosophy has generated remarkably consistent results. Table 1 compares the three CPU core designs from QED over the past four years. By targeting a conservative die size in state-of-the-art technology, QED's chips achieve high clock rates, moderate performance and price/performance, and an attractive profit margin. With each generation of process technology, QED's designers deliberately avoid "great ideas" that imply long on-chip signal paths and deep pipelines and make it difficult to reuse previous design efforts. In each design, cache consumes roughly 50% of die area (see Figure 1 for the RM7000).

Even in the third-generation RM7000, QED remains staunchly loyal to the classic five-stage pipeline, as shown in Figure 2. While most other processor cores in its performance class implement aggressive superscalar techniques,

	R4600	R5000	RM7000
SPECint95 (base)	—	5.5 int	>10 int
SPECfp95 (base)	—	5.5 fp	>10 fp
On-Chip L1 Cache	16K I, 16K D 2-way	32K I, 32K D 2-way	16K I, 16K D 4-way
On-Chip L2 Cache	none	none	256K unified 4-way
Off-Chip Cache	—	0–2M, 1-way	0–8M, 1-way
Instruction Issue	Single	Int + FP	Any two
Core Frequency	133 MHz	200 MHz	300 MHz
Die Size	77 mm^2	87 mm^2	80 mm^2
IC Process	0.5µ, 3M	0.35µ, 3M	0.25µ, 4M
Price @ Intro	$250	$300	$300*
Est Mfg Cost*	$20	$25	$35

Table 1. QED has advanced the performance of its MIPS chips with die size and price staying fairly constant. (Source: QED except *MDR estimates)

deep pipelines, branch prediction, and speculative execution, the RM7000 core implements simple (by today's standards) symmetric dual-instruction issue with no branch prediction.

Thanks to an effective combination of compact die size, high performance, reasonable power consumption, and relatively simple system interface, the chips designed by QED have successfully straddled the fence between low-end workstations and high-end embedded control. It has been relatively easy for the company to create targeted derivatives, such as the R4650 for embedded control and the R4700 for low-end workstations.

Wisely, QED is aiming the RM7000 squarely at the next generation of high-performance embedded applications. Silicon Graphics is likely to use the RM7000 in its low-end MIPS-based workstations, but by the time the RM7000 is shipping in volume, at least a year from now, the low-end workstation market will be under attack from commodity x86 systems. QED can keep busy designing RM7000 derivatives for specific embedded market segments.

Caches Matched to Core Enhancements

The most noteworthy characteristic of the RM7000 is its massive, 256K on-chip level-two cache. This write-back unified cache is four-way set-associative, has a 32-byte line size, and is nonblocking.

The current record for on-chip L2 cache is 96K, held by DEC's 21164 Alpha chip, and the current record for total on-chip cache is 160K, held by IBM's P2SC. With 288K of on-chip cache, the RM7000 will likely set a new record when it is available, and the chip will pack its caches into just over 40 mm^2 of die area.

Even with a 0.25-micron process with four metal layers and one poly layer, this miraculous feat of density could not be achieved with traditional SRAM design techniques. For comparison, the 0.35-micron R5000 crams 64K of cache into about 42 mm^2 of die area using a three-metal/two-poly process. QED was able to tune its SRAM cell layout to achieve a 15% reduction in size, but even generously assuming the denser process and cell design result in a factor of two improvement in density, the RM7000's L2 cache should hold only 128K.

QED will state publicly only that it has a patent pending on the memory cell used in the RM7000's L2 cache. To achieve such high density, the proprietary cell design is possibly some sort of DRAM-like structure.

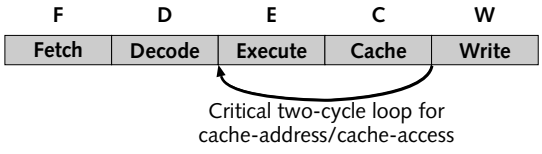

Figure 2. The RM7000's designers focused on the critical two-pipeline stage loop comprising the computation of a data-cache address and the subsequent cache access. With modern process technology, the other phases of a traditional pipeline can be made as fast as needed, but cache access needs special attention.

If desired, an RM7000 system can include an R5000-compatible off-chip cache of up to 8M in size. In an R5000 system, the off-chip cache is at the second level, but it acts as an L3 cache for an RM7000 system. To implement the cache, only external SRAMs are required: the RM7000 provides all cache control. If implemented, the write-through L3 cache is direct mapped.

To help make room for the large L2 cache and to improve yield at 300 MHz, the sizes of the separate L1 instruction and data caches were scaled back to 16K, half that of the R5000's caches. Normally, cutting the L1 cache in half would hurt performance, but, according to QED, any performance impact is offset by the increased associativity and the large on-chip L2 cache. Where the R5000 has two-way set-associative L1 caches, the RM7000's caches use a four-way organization. To service a miss in the L1 caches, the R5000 must access an off-chip cache (if the system implements it), but the RM7000's on-chip L2 cache can satisfy most L1 misses without the penalty cycles needed for an off-chip access. An L1 cache access takes one cycle; an L1 miss that is satisfied by the L2 caches takes only three cycles.

Like the L2 cache, the RM7000 L1 caches have a 32-byte line size and are write-back. The RM7000 is the first QED processor with nonblocking L1 caches. Nonblocking capability improves the efficiency of the RM7000's symmetric superscalar issue because the processor can have up to two outstanding loads: the CPU can issue a load even if a previous load caused an L1 or L2 cache miss.

RM7000 Improves Superscalar Symmetry

Like the R5000 *(see* **100102.PDF***)*, the RM7000 implements the MIPS IV instruction set architecture (ISA), which adds floating-point multiply-add, a register+register addressing mode for floating-point loads, and conditional move instructions to the full 64-bit architecture of the MIPS III ISA.

Both the R5000 and RM7000 are dual-issue superscalar machines, but unlike its predecessor, the RM7000 has symmetric superscalar capability. The R5000 could issue two instructions simultaneously only if one used the integer pipeline and the other used the floating-point pipeline. As

Issue Combination	R5000	RM7000
ALU-ALU	—	Yes
ALU-Load/Store	—	Yes
ALU-Branch	—	Yes
Load/Store-Branch	—	Yes
Branch-Branch	—	—
Load/Store-Load/Store	—	—
FP-Load/Store	Yes	Yes
FP-ALU	Yes	Yes
FP-Branch	Yes	—
FP-FP	—	—

Table 2. The RM7000 implements much more symmetric superscalar issue than its predecessor the R5000.

Price & Availability

The RM7000 is expected to tape out in early 1997, with samples available in the second half of the year. Contact QED (Santa Clara, Calif.) at 408.565.0300 or visit the company's Web site at *www.qedinc.com.*

Table 2 shows, the RM7000 implements far more general dual-issue capability.

One exception is the ability to dual-issue a floating-point operation with a branch. The RM7000 excludes this case because branches and floating-point operations use the same pipeline. In the R5000, integer and floating-point instructions were directed at physically different units. The RM7000 makes up the lost ground, however, in most situations by being able to dual-issue both a branch and a loop-update instruction. Note that for the combinations implemented, the RM7000 is completely symmetric; e.g., for ALU-load/store, the instructions in the pair can be in either order and still be issued simultaneously.

Pipeline Is Boring But Effective

In keeping with QED's philosophy of reusing as much as possible from previous design efforts, the RM7000 implements the same basic five-stage pipeline as the R4600 and R5000. Conceptually, QED considers the first two pipeline stages as the "front end" and the last three as the "back end." The RM7000 pipeline (see Figure 2), then, has a traditional back end and slightly enhanced front end.

When tuning the design to achieve the highest possible clock rate, QED endeavors to ensure that the two-cycle cache-address computation/cache-access loop is the critical path in the design. QED tunes the two-cycle cache-access critical path for highest speed and then makes sure the rest of the design can run at this speed. To tune this critical path, the company concentrated on reducing the R-C delay (delay associated with resistance and capacitance) in the buses between the caches and execution units.

To get the most out of the symmetric dual-issue capability without resorting to out-of-order execution, branch prediction, and speculative execution, the RM7000 always fetches two instructions per cycle, which usually allows the processor to fetch beyond an upcoming conditional branch. Then, when the branch is being decoded, the processor fetches from the branch-taken path.

Figure 3 illustrates how the RM7000 fetches from both paths of a conditional branch and then can execute the branch with minimal throughput impact. Each pair of horizontal rows shows the two pipelines for a given clock cycle. In the first cycle, the two instructions before the conditional branch are being fetched.

In the next cycle, the fetcher is fetching the branch and branch-delay instruction. The first pipeline is decoding one of the instructions before the branch, but the second pipeline decoder is idle because of, say, a resource constraint. This allows the instruction fetcher to get ahead; a buffer saves prematurely fetched instructions.

In the next cycle, the first two instructions on the fall-through path are being fetched while the instruction just before the branch is being decoded. In this cycle, the second decoder is again presumed idle, allowing the fetcher to get two instructions ahead.

In the next cycle, the branch is being decoded. The RM7000 pipeline knew the branch would enter the decode stage at the very end of the previous cycle, so in this cycle, the fetcher brings in the two instructions at the target of the conditional branch. At the end of this cycle, the branch direction is known, and the RM7000 pipeline uses the branch direction to select which two instructions—from the buffer for the fall-through path or the just-fetched instructions from the taken path—should go to the decoder in the next cycle (shown as $D_{correct}$ in Figure 3). The decision also selects the fetch address for the next cycle.

This example is an optimistic one, because the alignment of the branch target address is important and the fetcher must be able to get ahead of the decoders. Under most circumstances, however, the fetcher does get ahead of the decoders, and at least some of the time branch targets are correctly aligned. For performance tuning, future compilers may allow programmers to request favorable alignment for important code loops.

Without branch prediction or this enhancement to the front-end of the RM7000 pipeline, either taken branches or non-taken branches would cause a pipeline bubble and thus reduce performance.

The floating-point pipeline of the new chip is the same as in the R5000, which means the RM7000 will have somewhat lower peak performance per MHz than some other processors in its class for double-precision operands, but the large on-chip cache and the high operating frequency should boost absolute performance in many applications. At least by

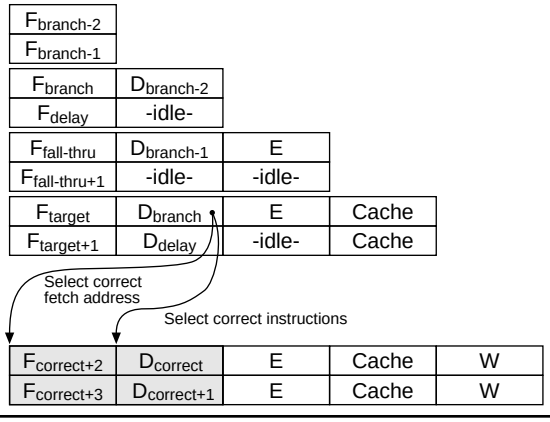

Figure 3. Under favorable circumstances, when the instruction fetch logic gets ahead of the execution phases of the pipeline, dual issue can be sustained even when a branch is executed.

today's standards, the projected 10 SPECfp95 is impressive—faster than MDR's projections for a 300-MHz P6—and should satisfy most applications that demand a general-purpose microprocessor.

System Interface Changes Little

To fully support the nonblocking caches, the RM7000 implements a new PRQST pin. On the R4600 and R5000, when a read is issued to the external interface, the system owns the interface until the read is completed. With the PRQST pin, the RM7000 can reclaim ownership of the interface even if an outstanding read is in progress. After the RM7000 issues a second outstanding read, the system owns the interface again until one of the reads completes; outstanding reads can complete out of order.

System designers will appreciate another improvement in the system interface: the ability to use bus-to-internal-clock ratios from 2.0 to 5.5 in increments of 0.5. Previous QED chips supported only integral ratios. With the RM7000, it will be possible to upgrade systems using a 100-MHz bus with a 250-MHz processor instead of forcing the choice of either a 200- or 300-MHz chip. The 4.5 ratio will allow a system to use a relatively relaxed 66-MHz external bus with a high-performance 300-MHz core. For many embedded applications, the large on-chip L2 cache will reduce the performance impact of the mismatch between the internal and external speeds.

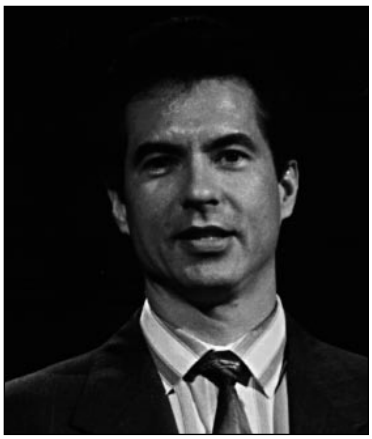

MICHAEL MUSTACCHI

At the Forum, John Kinsel of QED discusses the price/performance tradeoffs of the RM7000.

Aside from the new PRQST pin and the support for more bus-clock ratios, the RM7000 system interface is compatible with that of the R4600 and R5000. The package pinout may end up being the same as that of the R5000, depending on the final placement of logic blocks on the RM7000 die. To allow the RM7000 to be dropped into existing designs with at most a change in board layout, the new chip has a mode bit that disables the new features.

QED Grapples With Power Issues

Typical embedded systems need to minimize power dissipation, and embedded microprocessors have traditionally sported low to moderate power requirements. With the clock rates of embedded microprocessors hitting 300 MHz, however, it is getting increasingly difficult to keep power dissipation reasonable.

According to QED, the maximum power dissipation of the RM7000 is 12.5 W at 300 MHz. This power level is achieved only when the chip is able to sustain dual issue of integer and floating-point instructions, but it illustrates that supplying power and removing heat will be significant system design issues for the forthcoming generations of high-end embedded microprocessors.

Under typical operating conditions, the RM7000 will probably dissipate less than 10 W, and integer-only code might reduce power requirements even more. In addition, it seems strange that QED is planning to operate the RM7000 from a 3.3-V power supply when most other vendors of 0.25-micron microprocessors are planning to move to 2.5-V—or even 1.8-V—supplies *(see* **101203.PDF***)*. Perhaps the circuit design of the memory cells in the L2 cache requires the higher voltage. If the RM7000 were able to operate at a lower voltage, power dissipation would plummet to perhaps no more than 5 W.

In any case, by satisfying most of the processor's requests for data and instructions, the large on-chip caches of the RM7000 serve double duty: they improve performance and reduce total power dissipation. Power is reduced because the power-hungry external signal drivers are used less often to satisfy cache misses.

High-End Embedded Market Ready

The RM7000 is a potent combination of large caches, high-frequency operation, a high-performance core, and small die size. Its clock rate of 300 MHz is almost shocking in an embedded processor, since only high-end Alpha processors currently use a faster clock. MDR's cost estimate of $35 gives the RM7000 a significant cost advantage over other parts with midrange performance, such as the PowerPC 60x chips, PA-7300LC, and UltraSparc-2i.

On the other hand, even acknowledging QED's good track record of getting into production on time, the chip will not be in production until late next year at the earliest. While its estimated performance of greater than 10 on both SPECint95 and SPECfp95 would allow SGI to build very attractive low-end workstations today, this performance level will probably not be sufficient to compete against high-end x86 chips a year from now. Still, SGI may need low-end systems in its product line, and the RM7000 will likely be the best chip for such systems.

The high-end embedded market, however, offers more than enough challenge and opportunity for QED. For example, Cisco currently uses high-end embedded MIPS chips in its network-control products, and with the exploding popularity of the Internet and corporate intranets, Cisco and other infrastructure suppliers represent a profitable and expanding market. Similarly, high-speed color laser printers need lots of processing power and represent another expanding market. QED will be able to further reduce costs for some of these applications by eliminating the floating-point unit. With the addition of the RM7000 and its derivatives to its impressive portfolio of MIPS microprocessors, QED is well positioned to grab a significant share of important embedded markets. M