

IBM PowerPC 440 Hits 1,000 MIPS

High-Performance Embedded Core Implements Book E Architecture



by Tom R. Halfhill

IBM Microelectronics is slinging the fastest gun in the West—at least, according to the easy-to-please Dhrystone 2.1 benchmark. IBM's new PowerPC 440 core is the first officially announced embedded-processor core that's projected to hit 1,000 Dhrystone MIPS. The 440 may not be top gun for long, though, because future cores from Mips Technologies and others are aiming for the same target.

The 440 achieves some other firsts as well. It's the first core to implement Book E, the new embedded PowerPC architecture defined by IBM and Motorola (see MPR 5/10/99, p. 9). And it's the first core to use a 128-bit version of IBM's on-chip CoreConnect bus (see MPR 7/12/99, p. 8).

Principal architect Tom Sartorius described the 440 at this month's Microprocessor Forum. Although he said the design team didn't ignore such concerns as power consumption and die area, their primary goal was to exceed the performance of IBM's PowerPC 405 core (see MPR 10/26/98, p. 26) by about 3×, just as the 405 core exceeds the performance of the 401 core by the same factor. Actually, the 440 outperforms the recently announced 405GP chip (see MPR 7/12/99, p. 8) by nearly 4× on the Dhrystone benchmark. The 440 is expected to hit 1,000 MIPS at 555 MHz, while the 405GP delivers 252 MIPS at 200 MHz.

At its nominal clock frequency of 555 MHz, the 440 will typically consume about 1.4 W at 1.8 V, or 2.5 mW/MHz. The core, not including caches, occupies 4 mm² in IBM's CMOS-7SF, a 0.18-micron copper IC process. It uses only

four of the six metal layers available in that process, leaving two layers for interconnects with an auxiliary coprocessor (such as a floating-point unit) or additional devices in a system on a chip (SOC).

The primary caches can range in size from 0 to 64K, with 32- to 128-way set-associativity. Assuming dual 32K caches for instructions and data, the 440 has about 5.5 million transistors. The hard core is available now for custom designs, and IBM plans to sample the first product—probably an SOC similar to the 405GP—in 2Q00.

A Simple Approach to Superscalar

IBM designed the 440 for high-end applications: set-top boxes, information appliances, network computers, digital cameras, network printers, RAID controllers, routers, ATM switches, and cellular base stations.

Sartorius described the 440 as the “simplest possible superscalar design” consistent with the performance goals, but it's a design that not long ago would have been considered state of the art in desktop and server processors. The 440 is an out-of-order machine with three execution pipelines, dynamic branch prediction, 24 digital-signal-processing (DSP) instructions, and an ALU that can execute a 32×32 -bit integer multiply or a $16 \times 16 \rightarrow 32$ -bit multiply-accumulate (MAC) with single-cycle throughput.

As Figure 1 shows, there are two ALUs and a load/store unit. Only one ALU has a multiplier, so it handles complex instructions, such as MACs. The 440 can dispatch instructions to any two of these units in parallel.

To reach higher clock frequencies, the 440's pipelines are seven stages long, compared with five stages in the 405 and three stages in the 401. IBM added a predecode stage and a second execute stage for data-cache access. The predecode stage allows an auxiliary coprocessor, such as an FPU, to tie directly into the 440's pipelines.

When the 440 fetches and partially decodes an instruction that's intended for an auxiliary coprocessor, it immediately detours the instruction to the auxiliary unit. After executing the instruction, the coprocessor can feed results back into the 440's main pipeline during the write-back stage.

Auxiliary coprocessors enjoy equal status with the 440's function units. They can access the 440's general-purpose registers (GPRs) to read and write an instruction's source and destination operands, or they can have their own register files. They also participate in superscalar dispatching. The 440 can dispatch two instructions in parallel to a dual-pipelined coprocessor, or it can dispatch one instruction to a coprocessor and another to its own pipelines. This tightly coupled interface allows coprocessors to work as seamlessly as any

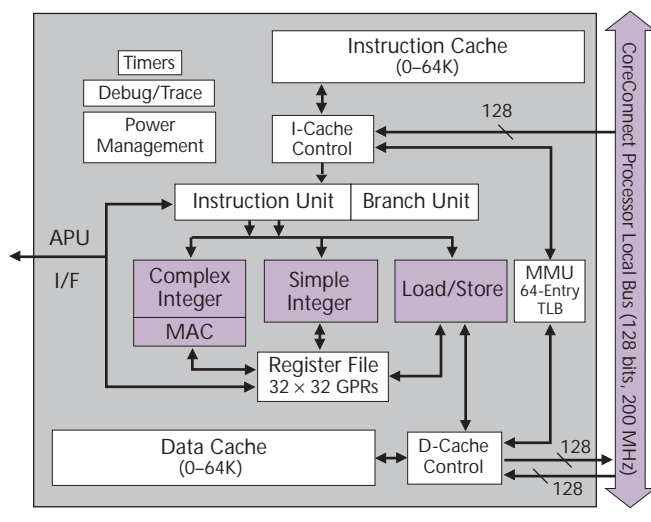


Figure 1. The PowerPC 440 is a two-way superscalar core with three execution pipelines (highlighted in purple).

other function units in the core, so ASIC designers can integrate FPUs or other auxiliary units for special purposes.

A Transient Hotel for Data

IBM took an unusual approach to the primary caches that should boost the 440's performance with certain kinds of multimedia and network applications. The 440's caches are blocks of content-addressable-memory RAMs (CAMRAMs). The blocks can be 1K, 2K, or 4K in size. Each block has 32-byte lines and is fully associative. ASIC designers can join as many as 16 of these blocks to create an instruction or data cache ranging in size from 8K to 64K. Therefore, an 8K cache would be 32-way set-associative; a 16K or 32K cache would be 64-way set-associative; and a 64K cache would be 128-way set-associative.

The CAMRAMs should improve performance in network routers, because address-table lookups will be much quicker. But IBM went further by allowing programmers to partition the caches in two different ways.

One scheme, which is common in embedded processors, allows programmers to lock any part of a cache to keep critical instructions or data resident. The 440 can lock a cache in increments of one way—so the smallest increment for a 32K cache would be 16 lines (512 bytes).

The second partitioning scheme is less common: programmers can define any part of a cache (again, in per-way increments) as a transient region. The CPU protects this region from normal victim replacements (which follow a round-robin policy), but doesn't lock it entirely. Transient regions are ideal for processing network packets or decoding MPEG video streams, because the CPU can perform repetitive operations on chunks of data that move sequentially through the region.

Running With the Big Dogs

To make sure the core isn't held back by slow interfaces to on-chip peripherals, IBM endowed the 440 with the first 128-bit version of its CoreConnect bus. Functionally, it's identical to the CoreConnect bus in the 405GP, except it's twice as wide, allows twice as many pending transactions (four), and supports twice as many bus masters (eight). It maintains compatibility with existing 32- and 64-bit devices. The 440 has three 128-bit interfaces to this bus, as Figure 1 shows.



IBM's Tom Sartorius, principal architect, described the PowerPC 440 core at Microprocessor Forum.


MICHAEL MUSTACCHI

Price & Availability

IBM's PowerPC 440 core is available now to ASIC developers and is scheduled to appear in the first standard product from IBM in 2Q00. The price depends on the customer's agreement with IBM. For more information, go to www.chips.ibm.com/products/powerpc/.

Table 1 compares the 440 with some other high-end cores recently announced by Hitachi, IDT, Intel, and Mips. The 440 clearly offers the best Dhrystone performance. IDT's new RISCORE 64600 (see MPR 9/13/99, p. 11) comes the closest and could narrow the gap if IDT squeezes out a bit more clock speed. The new 5Kc core from Mips (see MPR 10/25/99, p. 22) fares the worst in this comparison, but it's a soft core that offers ASIC developers more flexibility in return for the compromises imposed by synthesis tools, and the Dhrystone benchmark doesn't measure the advantages of its 64-bit architecture. The race will be closer next year when Mips rolls out its 20K hard core, code-named Ruby, which is expected to deliver about 1,000 MIPS.

Power-consumption comparisons are premature at this point because not all of the vendors have released power estimates for fully loaded cores with caches. But we expect Intel's second-generation StrongArm (see MPR 5/10/99, p. 1) to have the best power/performance ratio, though probably not the best raw performance.

The 440 offers IBM's customers a big step up from the 405 and 401 cores. The 440's cache organization is well suited to its intended applications, and the 128-bit CoreConnect bus allows developers to build fast SOCs. Together with the Book E improvements, these features make the 440 a superlative addition to IBM's line of embedded PowerPC cores. 

	PPC 440	5Kc	RC64600	SA-2	SH-8000/ST50
Vendor	IBM	Mips	IDT	Intel	Hitachi / ST
Architecture	PPC Book E	MIPS64	MIPS	StrongArm	SH-5
Arch Width	32 bits	64 bits	64 bits	32 bits	64 bits
Synthesizable?	No	Yes	No	No	No
Superscalar?	2-way	No	2-way	No	No
Issue Order	Out of order	In order	Out of order	In order	In order
Branch Prediction	Dynamic	Static	Dynamic	Dynamic	None
L1 Cache (I/D)	0-64K/0-64K	0-64K/0-64K	64K/64K	32K/32K	32K/32K
Core Frequency	555 MHz	300 MHz	400-500 MHz	600 MHz	400 MHz
Dhrystone 2.1	1,000 MIPS	360 MIPS	> 800 MIPS	> 700 MIPS	604 MIPS
IC Process	0.18μ	0.18μ (soft)	0.18μ	0.18μ	0.15μ
CPU Availability	Samples 2Q00	Samples 1Q00	Samples 1Q00	2H00*	2001*

Table 1. Among recently announced high-performance embedded cores, the PowerPC 440 has a clear advantage on the Dhrystone 2.1 benchmark. (Source: vendors, except *MDR estimates.)