

Motorola and IBM Unveil PowerPC 603

Slightly Slower than 601 at One-Third the Power Consumption

by Michael Slater



At last week's Microprocessor Forum, Motorola and IBM unveiled the PowerPC 603, the second in a line of PowerPC microprocessors being developed by the two companies. The new chip has a remarkably compact die size of 85 mm²—about 70% of the size of the 601, and only slightly larger than Intel's 486DX2.

Motorola and IBM made a full technical disclosure but did not announce pricing. The MPR Cost Model (*see 071004.PDF*) estimates the 603's production cost to be about \$66—roughly 35% below that of the 601. With the 601 currently priced at \$450 in quantities of 1,000, the 603 seems likely to be priced closer to \$300—less than half of Pentium's price and comparable to competitive MIPS and Alpha chips. Volume pricing will have to be far below that level if Apple is to meet its goal of putting the chip in systems selling for less than \$2000, and eventually closer to \$1000.

Motorola and IBM said that they have functional silicon now and expect to begin general sampling in 2Q94, with production in 3Q94. The 603 is expected to be available at 66 and 80 MHz; slower versions for lower cost and power are a possibility.

Sources claim that the first silicon is performing well, and that the delay from today's early samples to general sampling more than six months from now is due to a desire to use the final production version for these samples. The two key customers, IBM and Apple, already have early samples and will consume all the initial production volume, so there is no need for a widespread early sampling program.

Simulated performance with a 1M second-level cache, as shown in Table 1, is 60 SPECint92 and 70 SPECfp92 at 66 MHz, or 75 SPECint92 and 85 SPECfp92 at 80 MHz. This would make the 66-MHz chip approximately 7% slower than Pentium on integer code and 23% faster on floating point. The 603 requires a second-level cache to match the 601's performance without one.

	601 without L2 cache	601 with 1M L2 cache	603 with 1M L2 cache
66 MHz	62/72*	75/91	60/70
80 MHz	77/93	85/105	75/85

Table 1. SPEC ratings for the PowerPC 601 and 603, shown as SPECint92/SPECfp92. All are estimates, based on simulations, from IBM and Motorola, except: * measured result on an IBM RS/6000 Model 250 workstation. (Source: IBM)

The 603 is intended primarily for portable systems, although it also will be used in some low-cost desktop systems. It operates from a 3.3-V supply but can interface directly to 5-V logic. Maximum power dissipation at 80 MHz is expected to be under 3 W. At 66 MHz, it should be around 2.5 W—less than one-third of the 9 W consumed by the 601 at this clock rate, and a mere 20% of Pentium's power consumption. (Pentium will get closer when a 3.3-V version is available, but it still is likely to use twice as much power as the 603 at the same clock rate and power supply voltage.) The design also includes a number of power-management features that allow system designers to reduce average power consumption even further.

First True PowerPC Chip

The 603 is the first "true" PowerPC chip in that it is the first to be designed from scratch to implement that architecture. The 601, on the other hand, was a modification of IBM's earlier single-chip "RSC" design, and it implements the PowerPC instruction set and most POWER instructions. The 603 traps POWER instructions that are not part of the PowerPC architecture, allowing them to be emulated by system software for compatibility with older applications. Eliminating the logic for these instructions, many of which have significant implementation costs, simplified the CPU design to some degree, but the 603's more-aggressive microarchitecture increased the CPU complexity nonetheless.

The smaller die size of the 603 is due partly to a denser process technology, and partly to cutting the total cache size in half. While the 601 has a 32K unified cache, the 603 has separate 8K instruction and data caches.

The PowerPC 603 is also the first chip to be designed using the jointly developed tool set created for the Somerset design center. Both Motorola and IBM will manufacture and market the 603, providing true alternate sources and ensuring a competitive supply and pricing environment. The 601, on the other hand, is marketed by both companies but manufactured only by IBM. The 603 is fabricated in a 0.5-micron, four-layer-metal CMOS process that was jointly developed by the two companies. This process is slightly denser than the 0.65-micron process used for the 601, but it has about the same performance. Figure 2 shows the die photo.

The 603 will be offered by Motorola and IBM in a 240-pin CQFP package. The chip is close to the power level that would enable it to be packaged in a PQFP, which could help bring down the cost. Unlike the 601,

which distributes its pads across the surface of the chip, the 603 has both a traditional pad ring and a set of C4 solder bumps. IBM's chips use the C4 bumps to connect the pad ring to the CQFP package, but Motorola uses conventional wire bonding.

Revamped Microarchitecture

Figure 1 shows a block diagram of the PowerPC 603. It implements an entirely new microarchitecture that differs from the 601 in many respects. Table 2 summarizes the key differences between the two designs. The smaller cache size will make the 603 a less expensive chip than the 601, but it actually has a more aggressive microarchitecture. (See *061401.PDF* and *070602.PDF* for details on the 601.) In some respects, such as the use of register renaming and a separate load/store unit, it more closely resembles the multichip POWER processor designs. While the 601 was constrained by time-to-market demands to follow IBM's earlier design as much as possible, the 603 designers had the benefit of the being able to start with a clean slate.

The pipeline is similar to that in the 601: four stages

(fetch, decode, execute, write back) for integer instructions, five stages (fetch, decode, address generation, cache, write back) for load/store instructions, and six stages (fetch, decode, execute 1, execute 2, execute 3, write-back) for floating-point instructions. Branch instructions use a short two-stage pipeline (fetch, decode/execute).

The maximum number of instructions that can be issued in each cycle is the same for both the 601 and the 603, but the percentage of cycles for which this maximum can be achieved should be higher for the 603 because of the separate load/store unit and more flexible dispatch capabilities. The gains from the faster CPU core are more than offset by the loss from the smaller caches, so the 603 is slower than the 601. The die size is small enough, however, that a version with larger caches would be practical; a 603 with its caches doubled in size should out-perform a 601 at the same clock rate.

The 603's separate instruction and data caches enable high instruction-fetch bandwidth without very wide data paths or deep queues. While the 601 has an eight-word-wide path from its unified cache to an eight-word

prefetch buffer, the 603 has a two-word-wide path to a six-word prefetch buffer. Because there is no contention from data accesses, the narrower path still provides adequate instruction-fetch performance. Both the instruction and data caches are two-way set-associative and have a 32-word line size. The caches are physically indexed and physically tagged.

A three-state (modified, exclusive, invalid) cache consistency model is maintained with bus snooping. The cache tags are single-ported, so snooping accesses cause processor accesses to stall. The full four-state MESI model is not implemented because the chip is not designed for multiprocessor systems; the consistency features are designed to support DMA controllers or non-cached bus masters, such as I/O processors.

Following the shift in cache strategy, the 603 replaces the 601's large, unified translation look-aside buffers (TLBs) with two separate 64-entry, two-way set-associative TLBs. Unlike the 601, which provides hardware page-table searching on a TLB miss, the 603 relegates this task to software but provides a hardware assist. The hardware assist includes address generation logic and shadow registers for four of the general-purpose integer registers, which allows them to be used by the TLB handler without saving them first. These hardware functions enable the miss-handler routine to fit in two cache lines. A 52-bit virtual address is supported, but physical addresses are limited to 32 bits.

The 603 supports some speculative and out-




Figure 1. Block diagram of the 603, showing dedicated branch and load/store units.

	PowerPC 601	PowerPC 603
Transistors	2.8 million	1.6 million
Die size	120 mm ²	85 mm ²
Process	CMOS 4S shrink 0.65-micron, 4+ level metal	CMOS 5L 0.5-micron, 4 level metal
Fabricated by	IBM only	IBM and Motorola
Package	304-pin CQFP	240-pin CQFP
General sampling	Now	2Q94
Production	Now	3Q94
Clock speed	50, 66, 80 MHz	66, 80 MHz
Max. power @ 66 MHz	9 W	2.5 W
Supply voltage	3.6 V	3.3 V
Data bus width	64 bits	32 or 64 bits
Clock input	2 ×	1 ×
Caches	32K unified	8K data, 8K instr.
Cache line size	64 bytes with 32-byte sub-block	32 bytes
Cache associativity	8-way	2-way
Cache consistency	Four-state, for MP support	Three-state, for DMA
Load/store unit	No	Yes
Branch unit	Yes	Yes
Register renaming	No	Yes
Peak issue rate	3, including branches	3 effective (2 + branches processed earlier)

Table 2. Key differences between the PowerPC 601 and 603 processors.

of-order execution. The completion unit tracks the execution of all instructions. While instructions can be started out-of-order, they are always forced to complete in program order. A FIFO buffer holds completion information for up to five outstanding instructions. If this buffer is full, then instruction dispatch is stalled. In each cycle, two instructions can be dispatched and two can be completed.

More Execution Units than 601

The 603 adds two new execution units to the three present in the 601: a load/store unit (LSU) and a system register unit (SRU). The other three execution units are the branch processing unit (BPU), integer unit (IU), and floating-point unit (FPU).

The execution units are fed from the instruction unit. The dispatch unit (part of the instruction unit) takes two instructions from the instruction queue each cycle, checks them for dependencies, and dispatches both of them unless restricted by a dependency or a resource conflict. Each of the execution units has a reservation station that can hold one instruction waiting for execution. This allows the dispatch unit to continue dispatching instructions, even though previously dispatched instructions still may be awaiting execution. Thus, some out-of-order execution is allowed.

The integer, floating-point, and branch units each include rename registers: five for the IU, four for the FPU, and five for the condition register. These registers can be overlaid on any of the standard registers, reducing the number of instances in which dependencies cause

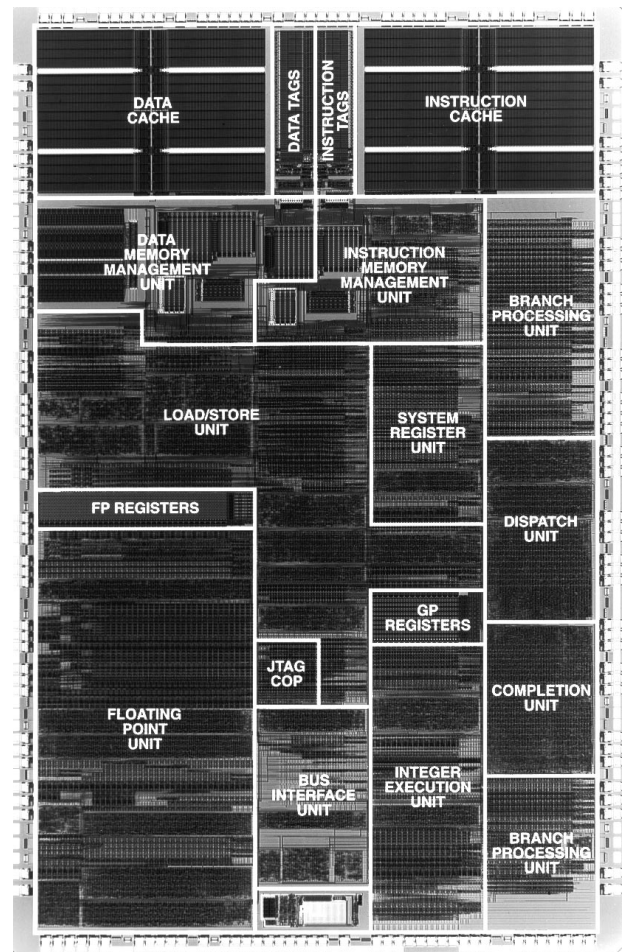



Figure 2. Die photo of the PowerPC 603, which incorporates 1.6 million transistors on an 85-mm² die.

instruction stalls. Rename registers prevent the CPU from stalling when one instruction wants to store data into a register from which a previous instruction is still waiting to read data. They are also a key part of the mechanism for supporting speculative execution.

As in the 601, branches are sent to a branch processing unit (BPU). The configuration is slightly different, however; instructions are fed either to the BPU or to the instruction queue in the 603, whereas the 601 design feeds all execution units, including the BPU, from the output of the queue. The 603 design approach allows the BPU to see branches earlier, so it can redirect the fetch sequence. Both processors have the same effective dispatch rate; the 601 dispatches up to three instructions from its instruction queue, while the 603 can dispatch only two, but branch instructions have already been removed from the instruction stream. Thus, as in the 601, the effective maximum dispatch rate is three instructions per cycle (IPC), provided that one of the three is a branch. Because of the fetch and completion limitations, the maximum sustained rate is two IPC.

The BPU uses static branch prediction, based on a single "predict" bit set by the compiler and encoded in all



Art Arizpe on Motorola's role in designing the PowerPC 603.

conditional branches. When a branch is encountered, the BPU causes the prefetch stream to follow the predicted direction of the branch. Instructions beyond the branch can move through the pipeline to the register write-back stage; writes are inhibited until the branch is resolved.

Integer multiply takes two to five cycles (data-dependent), and integer divide takes 37 cycles; this is slightly faster than the 601 for multiply, and slightly slower for divide. As in the 601, floating-point add, multiply, and multiply-add-fused (MAF) are fully pipelined for single precision; for double precision, add is fully pipelined, but multiply and MAF have a two-clock issue rate. In both designs, the floating-point pipeline has a basic latency of three cycles for single-precision operations, four cycles for double-precision. The 603 FP core is slightly faster because unlike the 601, it does not require a result to be written to the register file before it can be used by another operation; the result is forwarded directly from the write-back stage to feed the next instruction. FP divide takes 18 cycles for single-precision, or 33 cycles for double-precision—slightly slower than the 601.

The load/store unit (LSU) handles all load and store instructions. It calculates effective addresses, aligns data, and sequences load- and store-multiple and move-string instructions. It has a 64-bit path to the data cache, allowing a double-word read or write to complete in a single cycle.

The LSU speculatively executes loads, as long as there are no pending writes to the same address. The loaded data is written to a rename register; only when the completion logic indicates that the load is no longer speculative is the rename register assigned to the appropriate logical register number. Stores cannot, of course, be issued speculatively, but are held in a single-entry store buffer until the completion logic signals that the store can proceed.

The system register unit (SRU) is a specialized execution unit that handles system-level instructions, including logical operations on the condition register and moves to or from special-purpose registers. These functions were given their own unit for ease of implementation, not for performance reasons; they are infrequently executed.

Power Management Added

Several new features were added to the 603 to re-

duce power consumption. Automatic power management operates without software intervention and puts idle functional units in a low-power state; it extracts no performance penalty. For additional power savings, three software-controlled low-power modes are provided: doze, nap, and sleep.

In doze mode, most of the processor is shut down. The only parts that continue operating are the time base counter (which can produce a periodic interrupt), the bus snooping logic, and the phase-locked loop. Bus snooping continues, keeping the on-chip cache consistent and allowing external DMA activity to continue. Events that will wake up the processor include an external interrupt, a timer interrupt, reset, or a machine check (typically used to signal a memory error, which could occur during a DMA transfer). The transition to the full-on state takes only a few cycles.

Nap mode is identical to doze mode except that snooping is disabled. This is a lower-power mode for use when external bus activity does not threaten to create cache consistency problems.


Sleep mode shuts down the chip entirely. Once this mode is entered, external logic optionally can shut down the clock input and the on-chip phase-locked loop. If the clock has been stopped, the processor cannot be restarted until the clock and PLL are enabled and the PLL has been given time (about 50 μ s) to relock. After the PLL is locked, an interrupt, exception, machine check, or reset will return the processor to the operating mode.

With automatic power management enabled, maximum operating power is about 3 W at 80 MHz; typical is near 2 W. At the same frequency, doze mode cuts power to under half a watt, and nap mode reduces it further to about 200 mW. In sleep mode with the clock shut down, only the leakage current of a few μ A is drawn.

Compatible System Interface

As in the 601, the system interface is based on Motorola's 88110. It supports split transactions and allows one level of address pipelining. Unlike the 601, which requires a 64-bit data bus, the 603 provides a pin-selectable 32-bit bus option for low-cost system designs. The 603 also lacks some of the multiprocessor cache consistency signals present in the 601.

The 603 has an on-chip PLL that allows it to use a 1 \times clock input. As in the 601, the bus clock can run at 1/2,



Jim Kahle gives IBM's side of the PowerPC 603 story.

Price & Availability

Motorola and IBM have not announced pricing for the PowerPC 601. An announcement will be made when general sampling begins in 2Q93. Production is planned for 3Q93.

For more information from Motorola, call 800.845.MOTO or contact your local Motorola sales office. For information from IBM Microelectronics, call 800.IBM.0181

1/3, or 1/4 of the on-chip rate. A typical 66-MHz system would use a 33-MHz bus clock. The PLL is designed to run as slow as 16 MHz for low-power applications.

Market Opportunities

The 603 begins to broaden the PowerPC family. The 601, impressive as it is, required some compromises to meet its time-to-market demands. The 603 illustrates the high performance level that can be delivered on a modest power budget when the implementation is optimized with that goal in mind.

Motorola says that the 604, which is intended to deliver twice the performance of the 601 and will include multiprocessor support, is about one calendar quarter behind the 603. General sampling of the 604 is promised for mid-'94. The high end of the line will be filled out by the 620, the first 64-bit implementation of the architecture. General samples of this chip are expected late next year, with production in early '95. Embedded derivatives are also in the works at both IBM and Motorola.

The 603's primary role is expected to be in portable systems. For Apple, the chip will enable a PowerPC PowerBook with Pentium-level performance. The 603 at 66 MHz offers about ten times the performance of the 33-MHz 68030 in today's high-end PowerBooks, at about twice the power consumption. Thus, it won't do anything for battery life, but it will give a tremendous speed boost for native PowerPC applications.

The chip's significance in the broader PC market will depend on what software support emerges; the much-rumored port of Windows NT still has not been formally acknowledged, although numerous sources indicate that it is under way.

Given the required OS port—and, more significantly, a range of application support—the 603 could be the heart of an outstanding Windows NT portable. (This assumes, of course, that one wants a Windows NT portable. Given the much higher memory requirements of NT, the vast majority of portable systems will run Windows 3.1 or its successor, Chicago.) The 603's key competitors in this arena are the MIPS R4200 and R4600 processors, which deliver similar performance at roughly the same power consumption. A Pentium pro-

80-MHz 601 Announced

Just before their announcement of the PowerPC 603, IBM and Motorola both announced an 80-MHz version of the 601 processor. Samples from both companies are available now, and volume production is planned for January '94. The 50- and 66-MHz versions are already in production.

Only IBM manufactures 601 silicon, so it is no coincidence that both companies have announced the higher clock speed at the same time. Motorola prices the chip at \$500 in quantities of 20,000. The 50- and 66-MHz versions cost \$280 and \$374, respectively, in the same quantity. IBM quotes the 80-MHz chip at \$490 in quantities of 25,000, and \$275 and \$350 for the 50- and 66-MHz speeds.

At 80 MHz, the 601 operating without an external cache is rated at 77 SPECint92, and 93 SPECfp92. With a 1M level-two (L2) cache, the rating is 85 SPECint92 and 105 SPECfp92. This puts it in the same range as the 150-MHz R4400, and far above the 66-MHz Pentium's level of 64 integer, 57 FP. Compared with a 150-MHz Alpha chip, the 80-MHz 601 is a bit faster on integer code and slower on floating-point.

IBM also released new estimates for 66-MHz performance with a 1M level-two cache of 75 SPECint92 (a 25% jump from the earlier estimate) and 91 SPECfp92 (a 21% increase). The improvement over the year-old estimates comes partly from the conservatism of those simulation-based estimates, and partly from compiler enhancements.

cessor, even at 3.3-V, is unlikely to be competitive in terms of power consumption (or price), and no Alpha implementations yet come close to the required power level.

While penetrating the Windows PC market with RISC processors is going to require great patience, Apple's PowerPC Macs should give the PowerPC enough volume to fund continued development efforts. How significant IBM's consumption is depends on how far IBM pushes its Power Personal Systems into the PC market; IBM's workstation business won't generate high volumes from a semiconductor perspective.

Ironically, much of Motorola's and IBM's efforts to broaden PowerPC's role in the PC market will be directed toward competing with the Macintosh, since Apple still refuses to license its system software.

With the 603, the Somerset design center has delivered its first ground-up design. The coming year should bring a dramatic broadening of the family, along with significant new operating system and peripheral chip support. While breaking the x86's lock on the volume PC market remains a tough challenge that will require great patience, PowerPC's backers appear to have the design skills, financial resources, manufacturing capability, and determination that give PowerPC the best chance of any architecture to achieve this goal. ♦