

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

VOLUME 8 NUMBER 5

APRIL 18, 1994

PPC 604 Powers Past Pentium

PowerPC Chip Will Open Performance Gap, Possibly Permanently

by Linley Gwennap

IBM and Motorola have taken the wraps off their next processor, the 604, revealing the true performance potential of the PowerPC architecture. The chip, which the companies expect to ship in volume by the end of the year, is not only far faster than current Pentium processors but is likely to exceed the performance of all future Pentium chips as well. According to each company's roadmap, Intel will not surpass the 604's estimated 160 SPECint92 performance until the P6 generation—by which time PowerPC will have moved to even faster processors. Thus, the 604 will open a performance gap that Intel will find difficult to close.

While the 604 bears some resemblance to its little sibling, the 603, most aspects of the design have been pumped up to improve performance. For example, the 604 can issue four instructions per cycle, twice as many as the 603, and includes an extra integer ALU. The 604 uses register renaming and out-of-order execution more extensively than the 603. Its on-chip caches are twice as large, weighing in at 16K for instructions and 16K for data. Pushing the clock rate to 100 MHz and beyond also lifts the 604's performance.

As usual, this added performance comes at a price. The die size of the 604 is more than twice that of the 603 and 20% larger than the P54C Pentium. Neither vendor has released pricing information for the new chip, but it appears that the 604 will not replace the 601, as originally planned. Rather, the 604 will be positioned above the new 100-MHz 601, offering significantly more performance than Intel's Pentium for a similar price.

The 604 will appear in PCs and workstations from IBM as well as in Apple's second-generation Power Macs. The workstations will probably ship in 4Q94, with Macs and PCs shipping in 1Q95. Low-cost 604-based systems, running Apple's Mac OS and IBM's Workplace OS, will give the PowerPC boxes a performance advantage over Pentium PCs, but the 604 systems will be too expensive for most users, at least at first.

Emphasis on Integer Performance

Like the PowerPC 603, the 604 was designed from scratch at the jointly owned Somerset design center, and both Motorola and IBM will manufacture and market the processor. First silicon was received in January, and these chips are currently being tested by Apple and IBM, the two lead customers. The companies say that general sampling will not begin until 3Q94, but Canon's PowerHouse subsidiary (*see 0804MSB.PDF*) and other interested parties probably will receive early samples. Volume production is planned for 4Q94.

The 160-SPECint92 estimate assumes a 100-MHz 604 with a 1M external cache and a 66-MHz system bus. Ultimately, the chip may do even better, as 100 MHz is the "center frequency" of the design; some parts may run at higher clock rates. The 601, for example, was designed to a center frequency of 66 MHz and is now shipping at speeds up to 80 MHz in the original process. A hypothetical 120-MHz 604 could reach 190 SPECint92.

For floating-point applications, the 100-MHz 604 is estimated to achieve 165 SPECfp92. This represents a major improvement in SPECfp92 per MHz compared with earlier designs, although it does not match the larger relative increase in integer performance. The 604 designers spent more effort and die area on improving integer performance but did not neglect floating-point.

By the time the 604 is shipping, Intel expects 100-MHz Pentium chips to be available in volume desktop systems. At the same clock rate, the 604 will deliver 60% better integer performance and nearly twice the floating-point performance of Pentium, according to Somerset's estimates. Typical desktop systems will not include the expensive caches used to generate the quoted performance figures, but the ratio between the 604 and Pentium should be similar in less expensive designs.

Speculative and Out-of-Order Execution

Because the 603 and 604 were designed in parallel, the 604 is not derived from the 603, but the microarchi-

tectures are similar in many ways. The 603 (*see 071402.PDF*) can fetch and issue two instructions per cycle to four function units: an integer unit, floating-point unit, load/store unit, and system-register unit. As Figure 1 shows, the 604 adds a second integer unit and combines the integer multiplier and divider with the system registers to create a complex integer unit. The designers claim that this configuration includes three integer units, but in fact only two units handle typical (single-cycle) integer operations; the third unit handles less common multicycle integer operations and accesses to the special registers.

The two chips handle branch instructions somewhat differently. The 603 detects and handles branches early in the pipeline, removing them from the instruction stream; thus, the 603 can essentially execute three instructions per cycle if one is a branch. The 604 dispatches up to four instructions per cycle but, like the 601, has a separate unit to handle branches.

As in the 603, when the 604 encounters a branch, it predicts the outcome and begins to speculatively issue and execute instructions based on the prediction. The results of speculative operations are kept in rename registers or other temporary storage until the branch prediction is verified. The 604 includes 12 integer rename registers and 8 FP rename registers, about twice as many as the 603. Load instructions can be speculatively executed; speculative stores are kept in a six-entry queue and are not written to the cache until the branch prediction is verified.

Up to four instructions per cycle are read from the instruction cache into the four-entry decode buffer. These instructions are decoded in a single cycle, and the

results flow into the four-entry dispatch buffer. The dispatcher can dispatch up to four instructions per cycle. Rename registers are assigned and source values are read during the dispatch cycle. This dispatch cycle adds an extra stage to the pipeline used in previous PowerPC chips; more accurate branch prediction reduces the impact of this added stage.

Instructions flow between the two buffers in pairs, simplifying the dispatch logic. For example, if only one instruction is dispatched in a particular cycle, three instructions (not four) will be available for dispatch in the next cycle. There are a few other dispatch restrictions. Instructions are always issued in order, and only one instruction can be sent to each function unit per cycle. Instructions after a branch are not issued until the following cycle.

The dispatcher does not check for data dependencies; the function units handle this task. Each function unit contains two reservation stations (twice as many as in the 603) for instructions that cannot be executed immediately due to data dependencies. For the branch, load/store, and floating-point units, these reservation stations are simple FIFO queues, which prevent long-latency instructions from stalling the dispatcher.

The integer units have out-of-order queues that allow instructions ready for execution to bypass those still waiting for data. As a result, integer instructions can execute out of order relative to each other and relative to memory and floating-point instructions. As operations complete, their status is sent to the completion unit, which uses a 16-entry buffer to track pending instructions. Instructions are always retired in order, and up to four can be retired per cycle. If an instruction generates an exception, the exception is logged in the buffer and issued when the excepting instruction is retired, ensuring that exceptions are handled precisely and in correct program order.

Dynamic Branch Prediction

The 604 is the first PowerPC chip to use dynamic branch prediction, a technique used in Pentium, Alpha, and other microprocessors. Designers had resisted this technique on earlier chips because the PowerPC architecture includes a "predict" bit used by the compiler to signal the probable branch direction to the hardware. Dynamic prediction, however, increases the success rate and helps compensate for the added latency of the dispatch pipeline stage.

The 604 contains two structures to assist in branch prediction. The fully associative, 64-entry branch target address cache (BTAC) holds the target addresses of branches that have been recently executed and are predicted taken. The branch history table (BHT) contains 512 two-bit entries that are used to predict the direction of conditional branches. The BHT contains no tags and




Figure 1. The 604 can fetch and dispatch four instructions per cycle to six function units. Each function unit has two "reservation stations" for instructions that cannot be executed immediately.

thus always hits. If two or more branches are mapped to the same entry, their branch histories will be combined; the designers claim that this rarely happens.

During the fetch stage, the fetch address is issued to the BTAC in parallel with the instruction cache access. If the address hits in the BTAC, the resulting target address is used for the next instruction fetch, since the BTAC contains only branches that are predicted taken. Thus, for branches that hit in the BTAC, there is no penalty for correctly predicted taken branches.

Once the instructions have been decoded, further prediction occurs during the dispatch stage. If a branch tests the Counter register (typically used for loops), its outcome is predicted based on the current Counter value. For other branches, the BHT is accessed and the branch is predicted based on the two history bits using the Smith and Lee algorithm (see *070402.PDF*). In either case, if the prediction indicates a different address than the current fetch stream, the new address is sent to the instruction cache immediately. For taken branches that miss the BTAC, the penalty is 1–2 cycles.

The branch instruction eventually ends up in the branch unit, which waits for the operands to become available and determines the actual target address and outcome of the branch. If these results do not match the prediction made in the dispatch stage, the correct address is sent to the instruction cache. The branch unit also notifies the completion unit to invalidate the results of any instructions that were executed speculatively but incorrectly. The penalty for a mispredicted branch is at least three cycles and can be far greater if the branch operands are not immediately available when the instruction reaches the branch unit.

Double-Precision FPU

The dual integer units execute all simple register-to-register operations in a single cycle. The complex integer unit handles integer multiply and divide along with all instructions that access the special-purpose registers. The multiplier is pipelined and can sustain one multiply per cycle if one of the operands is 17 bits or less (i.e., the upper bits are all zeroes or all ones) or one instruction every two cycles if both operands are larger. The multiplier uses an “early termination” algorithm that detects the smaller operands and completes more quickly. Table 1 shows the latency and throughput of all complex integer operations.

The 604 is the first PowerPC chip to implement a full double-precision floating-point unit. Other designs use a single-precision FPU that takes less die area but requires more time for double-precision operations. As Table 1 shows, the 604 will significantly boost throughput for applications that make heavy use of double-precision arithmetic. Like other PowerPC chips, the 604 supports the fused multiply-accumulate instruction.

Price and Availability

Both IBM and Motorola expect to sample the 604 in 3Q94 with volume production in 4Q94. Neither company has announced pricing. For more information, contact IBM Microelectronics (Essex Junction, Vt.) at 800.POWERPC, or Motorola at 800.845.MOTO, or contact your local Motorola or IBM sales office.

The load/store unit calculates effective addresses and controls memory accesses. It performs address translation in parallel with the on-chip cache access, keeping the load-use penalty to only one cycle for integer data, as in most RISC processors. But because floating-point data is always stored on-chip in double-precision form, FP loads take an extra cycle to check the data type and possibly convert from SP to DP. The register file provides enough read ports for all units to operate simultaneously: eight for integer data and three for FP data.

Memory Design Avoids Stalls

As Table 2 shows, the caches and TLBs in the 604 are similar in structure but twice as large as those in the 603. One exception is that the caches in the 604 are four-way set-associative, increasing their hit rate. Compared with other RISC processors, the 604 is unusual in supporting parity for the on-chip caches and a true “least recently used” (LRU) replacement scheme rather than the “not recently used” algorithm implemented in most caches that have more than two sets.

The 604 includes much more extensive buffering of memory accesses than the 603, reducing the opportunities for loads or stores to stall the processor. As Figure 2 shows, loads that miss the on-chip data cache are kept in a four-deep queue while waiting to be processed. When the miss is handled, the requested data is loaded, critical word first, into one of two line-fill buffers. As soon as data arrives in the buffer, it can be sent directly to the CPU, allowing processing to resume more quickly. Once full, a buffer takes four cycles to transfer its data into the cache; the chip attempts to perform these refill cycles when the cache is not otherwise being accessed, since the

	PPC 604		PPC 603	
	Latency	Thruput	Latency	Thruput
Int Multiply (32 x 16)	3 cycles	1 cycle	2 cycles	2 cycles
Int Multiply (32 x 32)	4 cycles	2 cycles	5 cycles	5 cycles
Int Divide	20 cycles	19 cycles	37 cycles	36 cycles
FP Add / Mult (SP)	3 cycles	1 cycle	3 cycles	1 cycle
FP Add / Mult (DP)	3 cycles	1 cycle	4 cycles	2 cycles
FP Divide (SP)	18 cycles	18 cycles	18 cycles	18 cycles
FP Divide (DP)	31 cycles	31 cycles	33 cycles	33 cycles

Table 1. On a per-cycle basis, the 604 delivers twice the throughput of the 603 on integer operations and on most double-precision floating-point calculations.

	PPC 604	PPC 603	PPC 601
Cache Size	16K Instr 16K Data	8K Instr 8K Data	32K unified
Cache Associativity	4-way	2-way	8-way
Cache Line Size	32 bytes	32 bytes	64 bytes
Cache Coherency	MP support	DMA only	MP support
TLB Entries	128 Instr 128 Data	64 Instr 64 Data	256 unified
TLB Associativity	2-way	2-way	2-way
H/W Table Walk?	Yes	No	Yes

Table 2. The cache and TLB design of the 604 is similar to that of the 603, but most structures are twice the size.

cache is blocked during these cycles.

All stores are sent to the six-entry store queue; if the processor is executing speculatively, as it most often is, stores are not processed until the associated branch instruction is completed. Once the store at the front of the queue is ready, it attempts to write to the data cache; if the address misses the cache, the miss is handled as soon as the bus is available. The cache can accept subsequent stores while a store miss is being processed.

Unlike the 603, the 604 is designed for multiprocessor systems. The new chip uses a MESI protocol to keep the data cache coherent with other caches in the system. The 604 has dual-ported cache tags, so snoops from other processors do not stall CPU accesses.

Compatible System Interface

The 604 system bus is compatible with the 64-bit bus used by the 601 and 603. The new design supports a wider variety of clock multipliers than its predecessors: 1×, 1.5×, 2×, and 3×. The 604 is not pin-compatible with the 603, as it requires additional power and ground pins to feed its 10-W (typical) power consumption, nor is it pin-compatible with the 601, which uses the same 304-pin CQFP package but different supply voltages, as Table 3 shows.

Figure 3 shows a die photo of the 196-mm² 604, which has more than twice the die area of the 603 despite using the same 0.65-micron CMOS process. (IBM




Figure 2. Load and store queues prevent memory accesses from stalling the CPU and assist in speculative execution.

	PPC 604	PPC 603	PPC 601+	PPC 601
Clock Rate	100 MHz	66–80 MHz	100 MHz	50–80 MHz
Total Cache	32K	16K	32K	32K
Peak Issue Rate	4 instrs	3 effective (2+branch)	3 instrs (2+branch)	3 instrs (2+branch)
Load/Store Unit?	Yes	Yes	No	No
Dual Integer ALUs?	Yes	No	No	No
Reg. Renaming?	Yes	Yes	No	No
Branch Prediction	Dynamic	Static	Static	Static
Supply Voltage	3.3 V	3.3 V	2.5 V	3.6 V
Maximum Power	13 W	3 W	4 W	10 W
Clock Input	1×...1/3×	1×...1/4×	2×	2×
Package (CQFP)	304-pin	240-pin	304-pin	304-pin
IC Process (drawn)	0.65 μm 5 metal	0.65 μm 5 metal	0.50 μm 5M + local	0.72 μm 4M + local
Transistors	3.6 million	1.6 million	2.8 million	2.8 million
Die Size	196 mm ²	85 mm ²	74 mm ²	121 mm ²
Est. Mfg. Cost*	\$180	\$50	\$75	\$85
Availability	4Q94	3Q94	4Q94	4Q93
Estimated SPEC92	160 int	75 int	105 int	85 int
@ Max Clock Rate	165 fp	85 fp	125 fp	105 fp

Table 3. The 601, 603, and 604 provide a similar system interface but cover a wide price/performance range. (Source: Somerset except *MPR Cost Model)

calls this a 0.5-micron process, but the gate length is comparable to most vendors' 0.65-micron processes; see [080504.PDF](#) for details.) The increase in size comes from a variety of factors. The caches and TLBs have twice the capacity, and the FPU uses twice the area to handle DP arithmetic. The instruction and completion buses are wide enough for four instructions rather than two, and the memory unit is expanded by the load and store queues. The register files are larger due to the extra read ports; as the die photo shows, the general registers alone occupy 6% of the die, while they are barely visible on the 603 die photo. Interestingly, the second integer ALU adds less than 3% to the 604 die size.

The complexity of the control logic for the four-way instruction issue and out-of-order execution should not be underestimated. The dispatch and completion unit on the 604 consumes about 22 mm², more than four times the 5 mm² used by this unit in the 603 with the same IC process. Although the combined effect of these changes significantly increases the cost of the 604, it also boosts the chip's performance.

The MPR Cost Model (see [071004.PDF](#)) estimates that the 604 will cost about \$180 to build, compared with just \$60 for the 603. The 604's cost is comparable to that of other high-performance processors.

Neither Motorola nor IBM has announced pricing for the 604. The companies' stated goal is to deliver twice the performance of an x86 processor at the same price. The 100-MHz 604 should have nearly twice the integer performance of a 90-MHz Pentium, which sells for \$849 today. By the time the 604 ships, however, the Pentium price will have dropped to \$700–\$750.




Figure 3. The PowerPC 604 incorporates 3.6 million transistors on a 12.4×15.8 mm die using 0.65-micron, five-layer-metal CMOS.

Performance Advantage over Intel

Based on the 160-SPECint92 estimate, the 604 will outrun any microprocessor shipping today, although its estimated floating-point performance trails HP's PA-7150 and Digital's 21064, which peak at 200 SPECfp92. Digital estimates that its 275-MHz 21064A, due this summer, will hit 170 SPECint92, which may prevent the 604 from gaining the performance lead. Early next year, however, a variety of RISC chips are likely to leave the 604 in the dust with SPECint92 ratings of more than 200; these include the MIPS T5, Digital's 21164, Ultra-Sparc, and the PowerPC 620. The 604 should hold a cost advantage over these higher-speed processors.

The primary target of the 604, like other PowerPC chips, is Intel's product line. Figure 4 shows how the two families stack up. The 603 and new 100-MHz 601 (see "601" sidebar) will cover the performance range of existing Pentium chips (shown in dark gray) but have a much lower manufacturing cost. This translates into a price advantage as well: a 90-MHz Pentium, at \$849, costs 60% more than an 80-MHz 601 in comparable volumes.

The 604 establishes a performance range well beyond current Pentium chips. (The chart assumes that the 604 will ship at 80, 100, and 120 MHz.) In fact, it

PowerPC 601 Hits 100 MHz

Separately, IBM and Motorola have announced that they will ship a 100-MHz version of the PowerPC 601 around the end of the year. The new chip is logically identical to the original 601 (see [070602.PDF](#)) but is built using IBM's 0.5-micron CMOS-5X process. The new process delivers a significant speed increase beyond the original 0.7-micron CMOS-4S process, allowing the chip to reach 100 MHz—and possibly beyond. At the new clock speed, the 601 is estimated to achieve 105 SPECint92 and 125 SPECfp92 using a 1M external cache and a 50-MHz system bus.

The 5X process also reduces the die size considerably, from 121 mm^2 to 74 mm^2 . The new die is smaller than a 486DX2 or DX4 and smaller than the PowerPC 603, yet it contains 1.2 million more transistors. The new process has smaller gates than the CMOS-5L process used by the 603 and 604 and includes a local interconnect that greatly reduces the size of the cache arrays (see [080504.PDF](#)). The MPR Cost Model estimates that the new 601 die, which we call the 601+, will cost about \$75 to build, 15% less than the original 601.

The 601+ is pin-compatible with the 601 except that the 5X process requires a core voltage of 2.5 V, compared with 3.6 V for the 601. The 601+ permits the I/O pad ring to run from a separate voltage of up to 5 V for compatibility with standard memory and system-logic chips. The lower voltage and smaller transistors cut the maximum power dissipation to 4 W at 100 MHz, half the power used by a 66-MHz 601.

Like all 601 chips, the 100-MHz 601 will be manufactured exclusively by IBM but marketed by both Motorola and IBM. Although Motorola has the ability to build CMOS-5L parts such as the 603 and 604, it does not yet have a fab capable of running the 5X process. The companies plan to sample 100-MHz parts in 2Q94, with volume production in 4Q94. Neither vendor has yet released pricing; we expect the 100-MHz version, by the time it begins shipping, to sell for around \$500 in 1,000-unit quantities.

should easily exceed the performance of the 150-MHz Pentium that Intel plans for late 1995. The estimated manufacturing cost of the 604 is only slightly higher than Pentium's, and given Intel's high margins, Motorola and IBM should have little trouble maintaining a significant price/performance advantage—unless Intel decides to reduce its margins.

Intel's first release of its next-generation P6 processor, also expected in late 1995, could match or even exceed the performance of a hypothetical 120-MHz 604. IBM and Motorola say that the 620 will ship well before the P6 and deliver significantly better performance. Preliminary information indicates that the 620 may be less costly than the P6, which could be quite expensive.

Both the 603 and 604 could be improved beyond what is shown here by using the same CMOS-5X process as the 601+. IBM plans to have this process on line late this year and will be motivated to shrink the parts in order to lower cost and increase clock speed. Motorola currently does not have plans to implement the 5X process but has licensed rights to it. Should IBM aggressively move to shrink its PowerPC parts, Motorola will have to do the same to remain competitive.

Carving a Slice for PowerPC

The 604 should put PowerPC among the leaders in performance and create a substantial advantage over Pentium. Its first impact will be to hasten the transition from IBM's own POWER architecture. The 604 will outperform the Power1 processor, which is still used in the vast majority of IBM's RS/6000 systems, on nearly any application and will beat the Power2 processor on all but the most memory- or FP-intensive programs. As a single-chip CPU, the 604 carries a far lower manufacturing cost than either of the multichip POWER processors.

The 604 will be used in a variety of RS/6000 workstations and servers. These systems will increase the price competitiveness of the RS/6000 compared with other RISC workstations. The AS/400 line, currently CISC-based, is also moving to PowerPC and will receive a tremendous price/performance boost, although IBM will struggle with the CISC-to-RISC transition.

The biggest volumes, of course, will come from the PC market. Apple is currently shipping systems based on the 601 (see *080401.PDF*) and plans to ship 604-based systems in early 1995. These future products will, for the first time, give Apple an integer performance advantage over x86-based PCs, while Apple's low-cost manufacturing should keep its costs competitive with those of PC vendors. Systems using the 604 will double the performance of the initial Power Macs within a year of their release, greatly improving their competitiveness.

These 604 systems should attract power-hungry PC buyers who are not wedded to Microsoft Windows. The majority of users, however, may not want to give up their




Figure 4. With chips shipping around the end of 1994 (dark gray) and in 1995 (light gray), PowerPC should maintain a significant performance advantage over x86 with comparable factory costs. (Sources: 1994 performance—vendors; 1995 performance—MPR estimates; manufacturing costs—MPR Cost Model)

Windows software; although the Mac can emulate x86 applications, it requires 16M of memory and can still barely manage 486DX performance. Thus, Apple's impact on the overall PC market will continue to be limited.

The 604 also will be featured in PCs from IBM, Canon, and others. These systems will run Windows NT or IBM's Workplace OS with an OS/2 shell, operating systems that are closer to the look and feel of Windows. The speed of the 604 should allow it to emulate x86 software at the speed of a fast 486, respectable performance for most existing Windows applications.

Many questions remain to be answered. Can IBM, Canon, and others deliver full software compatibility with the existing base of Windows applications using a reasonable amount of DRAM? Can they, like Apple, deliver RISC systems at PC price points? Can they convince key application vendors to port their software to yet another new platform? The 604 does provide the answer to one important question: it shows that PowerPC can outperform Intel's x86 chips. ♦