

# PowerPC G4 Gains Velocity

## Motorola Adds Pipe Stages and On-Chip L2 to AltiVec Processor



by Keith Diefendorff

Feeling cautiously optimistic after Apple's resurrection, Motorola is making a major upgrade to the microarchitecture of its MPC7400 with AltiVec (née G4), which is at the heart of the Macintosh G4 systems that Apple began shipping in September. Speaking earlier this month at Microprocessor Forum, Naras Iyengar, manager of the project at Motorola's Somerset Design Center in Austin, described the new microarchitecture.

The two most significant enhancements over the 7400 include deepening the pipeline to achieve frequencies of more than 700 MHz and adding an on-chip 256K L2 cache to boost memory-system performance. The team is also tacking on new features for the embedded market and is preparing the design for fabrication in Motorola's forthcoming 0.18-micron copper HyperMOS-6 (HIP6) process. While Iyengar described the technical details of the enhancements, he stopped short of announcing any products or schedules. Sources indicate, however, that the new part taped out about two months ago, indicating a production date of mid-2000.

### For Lack of a Better Name

Confounded by a former partner with a different marketing agenda and an arrogant customer with little sympathy for Motorola's branding needs, Motorola's feeble efforts to name its processors have resulted in a confusing snarl of part names

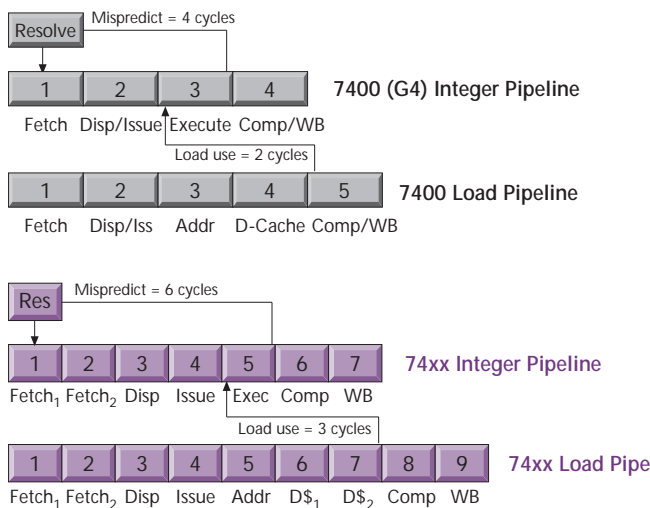


Figure 1. The 74xx superpipelines the instruction and data cache access, breaks dispatch and issue into separate cycles, and also breaks completion and writeback into two cycles. Cycle time is shown to scale for 500-MHz operation of the 7400 and for 700-MHz operation of the 74xx.

and numbers (see sidebar, next page). Motorola's latest attempt to establish AltiVec as the brand name for its SIMD extensions, for example, was dealt a blow when Apple undermined Motorola with its own name for the same feature: the Velocity Engine. Although that name is probably one Apple's customers can relate to more easily than AltiVec, the disparity will inevitably cause some confusion in the market.

Apparently gun-shy, Motorola was unable to settle on a name for the processor in time for its presentation at the Forum, referring to it simply as the next-generation G4. Realizing the folly of keeping it nameless, however, Motorola has since decided to forge ahead with its plans to call the new part the 74xx, differentiating it only slightly from the current 7400, or G4 (see MPR 11/16/98, p. 17). The "xx" designator apparently implies a family of processors based on the new microarchitecture.

The 74xx moniker is curious. Since the microarchitectural changes made between the 7400 and 74xx are significant—more significant than those Intel made between Pentium II and Pentium III—it isn't clear why Motorola has elected to label it as a fourth-generation PowerPC processor. A less conservative approach would have been to market the new design as the G5, but that name would have given Apple less than a year of air time with the G4. According to Motorola's long-range marketing roadmap, the G5, which it will label 75xx, is not due until sometime in 2001.

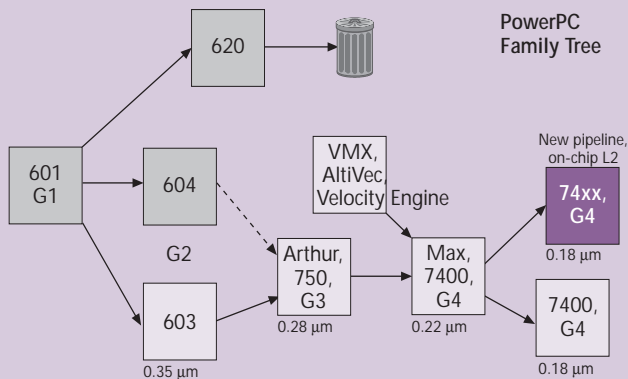
### Longer Pipeline Lifts G4 Speed Limit

Even though PowerPC chips don't compete directly with x86 chips, the frequency mania that has swept the PC market like a plague has also infected the Macintosh market. To some extent, Apple has successfully deflected the issue with emotional appeals to the Macintosh faithful, translucent colored plastics, and bogus benchmarks (e.g., Bytemarks). With the Velocity Engine, Apple finally has a legitimate claim to performance leadership, at least in the multimedia niche, but this advantage will be difficult to sell in the face of large frequency deficits. Although it may not be necessary to have higher frequency than x86 processors, Apple's processors cannot lag conspicuously behind if the company is to have any realistic hopes of recapturing market share from PCs.

So far Motorola has fought the frequency battle with advanced IC processes. Using the IC-process skills it was forced to develop during its rivalry with IBM over Apple's microprocessor sockets, Motorola has pushed the 7400's pipeline to the brink. Applying more pressure to its fabs could inflate wafer costs or lead to serious yield problems, like those AMD had with its K6-2. Eventually, however, copper wires, strong-phase-shift masks (see MPR 2/15/99, p. 4),

## An Update on the PowerPC G4

As the figure below shows, the lineage of the 74xx began with the 603 (see MPR 10/25/93, p. 11), a second-generation PowerPC processor and the first to be designed entirely at the joint IBM/Motorola Somerset Design Center. The 603 begat the 750 (aka Arthur, aka G3) (see MPR 2/17/97, p. 10), which used basically the same micro-architecture as the 603 but added an external backside L2 cache. The 750 begat the 7400 (aka G4) (see MPR 11/16/98, p. 17), which has exactly the same pipeline as the 750 but adds AltiVec (aka VMX, aka Velocity Engine) (see MPR 5/11/98, p. 1). Thus, the current 7400 uses a pipeline that is well over six years old—surpassing even the age of Intel's P6 pipeline.



Considering the advanced age of the pipeline and the fact that the 7400 was underpipelined to begin with, Motorola will do well to coax it to 500 MHz. While this frequency is still far below that of Intel's fastest 733-MHz Coppermine Pentium III (see MPR 10/25/99, p. 1), the 7400 has to overcome the large handicaps of a four-stage pipeline and a 0.22-micron process compared with Pentium III's 12-stage pipeline and 0.18-micron process.

and other process tricks will play out; Motorola has had to take more drastic steps.

It has done what every other PC-microprocessor manufacturer has already done: lengthen the pipeline beyond five stages. From an analysis of the 7400's critical timing paths, Iyengar's team concluded that L1 cache accesses, instruction issue, and instruction completion all stood in the way of reducing cycle time. As a result, they superpipelined the instruction fetch and data-cache access stages and separated instruction dispatch from issue and instruction completion from operand writeback, as Figure 1 shows.

Although the execution latency of all the simple scalar and vector (SIMD) instructions remains one cycle, as in the 7400, the latency of several of the more complex execution units had to be increased. The scalar double-precision

The 7400-500 will score about 24 on SPECint95 (base), roughly on a par with the Katmai-based Pentium III-600, but about 30% below Coppermine-733, which, with its on-chip L2 cache, will score about 35. On floating point, the 7400 with a large external L2 cache will score 21 SPECfp95 (base), about 20% behind Coppermine's score of 27.

On multimedia applications, however, the 7400 reigns supreme. According to Motorola's tests, AltiVec raises performance on multimedia and DSP algorithms by up to 20×, and routinely more than 8×, compared with a 7400 not using AltiVec.

Apple says its AltiVec-enhanced signal-processing library, the one that will be included in the next release of MacOS 9, running on a 7400 outperforms Katmai by an average of 3.5× at equivalent clock speeds on six common algorithms (bsqr1, bMpy2, DotProd, FFT, bFir, and Convolution) reported by Intel for its SSE-enhanced signal-processing library (see [www.apple.com/powermac/processor.html](http://www.apple.com/powermac/processor.html)). This advantage remains at about 2.3×, even after adjusting for Coppermine's much higher frequency and the multimedia-performance boost Intel reports for the on-chip L2.

The 7400 also boasts a small die size and low power. At 83 mm<sup>2</sup> in 0.22-micron HIP5, the 7400 is about 35% smaller than a 0.25-micron Katmai. In 1H00, Motorola will respin the 7400 in its 0.18-micron HIP6 process. This part, which the company will likely call the 7410, should reduce die size to about 50 mm<sup>2</sup> and increase frequency to 600 MHz. The 1.8-V HIP5 7400-500 will dissipate about 14 W (max), but the 1.5-V HIP6 version should dissipate only about 10 W, making it suitable for notebooks. The new 500-MHz mobile Coppermine dissipates about 14 W at 1.3 V. Sources indicate that in the 7410, Motorola may elect to take advantage of the 128-bit backside-cache option, which is pinned out to only 64 bits in the 7400.

floating-point latency was upped by 66%, from three cycles to five. Also, the vector permute unit was increased from one cycle to two, and the vector complex-integer unit (e.g. multiply-sum) was bumped from three to four. The vector floating-point units remain at four cycles. The throughput of all vector instructions, and of all scalar computational instructions except divides and integer multiplies, remains one cycle.

A higher number of cycles per instruction due to greater execution-unit latencies could theoretically hurt performance on some applications. Motorola, however, believes the impact will be small for existing code and virtually nonexistent on recompiled code. It says that code potentially influenced by greater latencies is generally rich in parallelism, and the PowerPC architecture provides enough registers (32 scalar integer, 32 scalar floating point, and 32 vector)

for the compiler to cover the additional latency. In any event, the increase in frequency will be a net performance boost, even if the greater latencies cannot be fully absorbed.

### New Stages Threaten IPC

In microarchitecture as in life, there is no free lunch. Simply lengthening the pipeline may have had the desired marketing effect on frequency, but it would not have increased performance. Increasing pipeline length from four to seven cycles, would, ideally, increase frequency by 75%. In practice, however, the effect is much smaller. Pipeline latch overheads and clock skews conspire to reduce usable cycle time. In addition, the granularity of pipeline stages and the uneven distribution of work among stages limit frequency to well below the theoretical speedup.

Even if the theoretical frequency gains could be realized, simply lengthening the pipeline probably would not have increased performance much, if any. The number of cycles required to execute each instruction would have increased by the same ratio that frequency increased, and the longer instruction pipeline would have been less efficient, owing primarily to larger branch-misprediction penalties, longer execution-unit latencies, and longer load latency.

To realize an actual performance gain along with its frequency increase, Motorola made a series of enhancements to offset what otherwise would have been a significant loss of IPC (throughput in instructions per cycle). First, it increased

the instruction-issue bandwidth and added new execution units to eliminate some structural hazards. Next, it increased the instruction-reorder depth to exploit more instruction-level parallelism (ILP), and the branch predictor was beefed up to reduce stalls on control-flow hazards.

Motorola's simulations show that together these improvements kept the 74xx's IPC on a par with that of the 7400, allowing all the frequency gained from the longer pipeline to fall through to real performance gains. Beyond these performance gains, however, Motorola went on to reduce average memory-access time by including a fast 256K L2 cache on chip, as Figure 2 shows, and by converting the 7400's backside L2 cache into an L3.

### Transistors to the Rescue

These new features all require transistors and silicon area to implement. Fortunately, if the 7400 has anything to spare, it's silicon area. In equivalent IC processes, a 7400 would be about 25% smaller than a Katmai-based Pentium III, according to our analysis. Even with an on-chip 256K L2 cache, a 7400 would be about 25% smaller than Coppermine (see MPR 10/25/99, p. 1). Thus, assuming Motorola's process is similar to Intel's, the 74xx should have a good deal of silicon available for new features before it would approach the size of Coppermine, which, at 106 mm<sup>2</sup>, is a reasonably small die itself.

Indeed, our analyses of Motorola's HIP6 process (see MPR 9/14/98, p. 1) and Intel's P858 process (see MPR 1/25/99, p. 22) indicate they are similar. Even though Motorola—probably for misguided marketing reasons—insists on referring to its process as “0.13 micron,” HIP6 features are very similar to those of Intel's P858 and IBM's CMOS-8S, which those companies more accurately call “0.18 micron.”

Although its feature sizes, and probably its transistor speeds as well, are similar to those of P858, HIP6 does have a couple of advantages. One is copper interconnect layers, which improve interconnect speed and, with greater current-carrying capability and fewer vias, also increase density. In addition, the Motorola process uses a tungsten local-interconnect (LI) layer. LI increases overall circuit density slightly but has a very large effect on SRAM cell size; indeed, a HIP6 SRAM cell is 4.5 μm<sup>2</sup>, 20% smaller than a P858 cell. This is a significant advantage for chips like the 74xx and Coppermine that have large on-chip caches. It's especially important for the 74xx, which has about 25% more SRAM bits than a Coppermine (due to larger L1s and its L3 tags).

The 74xx requires a total of about 33 million transistors, over 75% of which are in the L1 and L2 caches and the L3 tags. Of these, the L2 requires about 17 million transistors. Motorola attributes another 2 million to the 74xx's modular design (described later), primarily due to a larger than otherwise necessary memory-subsystem controller. This leaves about 6.5 million transistors in the 74xx core, 80% more than in the 7400 core. According to our calculations, this would make the 74xx about 95 mm<sup>2</sup> in HIP6.

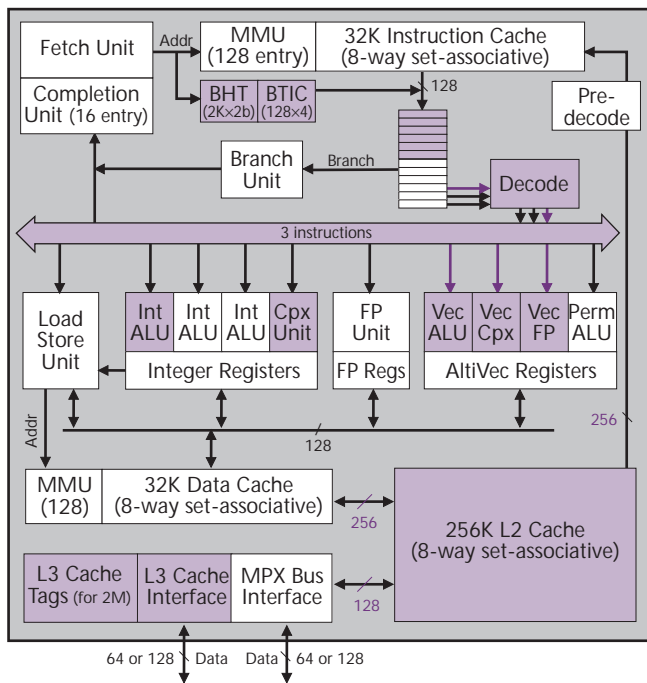


Figure 2. The 74xx microarchitecture has many enhancements (purple) over the 7400, including a longer seven-stage pipeline, wider instruction dispatch (three plus one branch), two new integer execution units, a less restrictive AltiVec issue matrix, and a 256K on-chip L2 cache.

### Issue Width, Throughput Increased

Much of Motorola's effort to avoid IPC degradation went into increasing instruction throughput. The instruction-fetch buffer was deepened from 6 entries in the 7400 to 12 in the 74xx. The deeper buffer allows the instruction fetcher to stay ahead of instruction dispatch, which was widened from two instructions plus one branch to three instructions plus one branch. To sustain the higher peak instruction-dispatch rate, two scalar integer execution units were added: one simple ALU, bringing the total to three, and one complex unit, which prevents iterative instructions such as integer divides from blocking dispatch of three simple instructions. With these additions, the instruction window was increased by 33%, from six instructions to eight.

The number of Altivec units remains at four, as in the 7400, but the units have been decoupled from each other to allow greater dispatch flexibility. In the 7400, the vector-permute unit was a separate dispatchable unit, while the vector-ALU, vector-complex, and vector-floating-point units were grouped into a second dispatchable unit. This grouping precluded, for example, the dispatch of a vector compare and a vector multiply in the same cycle. In the 74xx, even though only two vector instructions can still be dispatched per cycle—due to register port limitations—the four vector units are now orthogonal, eliminating many annoying and seemingly arbitrary structural hazards. The change should significantly improve performance on tight DSP inner loops.

To cover the longer execution unit latencies, the number of rename registers in the 74xx was increased from 6 for each register file to 16. In addition, the number of instruction-completion buffers was doubled from 8 to 16, allowing twice as many instructions to be in flight at a time. This enhancement allows the 74xx to search further ahead in the instruction stream to find independent instructions to issue, thus ameliorating the longer execution-unit latencies and increasing execution-unit utilization. This feature should significantly improve the 74xx's ability to exploit ILP.

Branch latency is another serious problem imposed by the 74xx's longer pipelines. Because pipelined processors must often predict control-flow direction before the instructions that determine the direction are complete, a misprediction can create a long stall while incorrectly fetched and executed instructions are purged and the pipeline refilled. On the 7400, the mispredict penalty was only four cycles and, due to the very short pipeline, most branch conditions were evaluated ahead of the branch that tested them, avoiding the need to predict these branches altogether.

The mispredict penalty of a 74xx, however, is 50% longer than that of the 7400, and the pipeline is 75% longer, requiring the branch predictor to be souped up to maintain a reasonable branch penalty ( $\text{mispredict\_rate} \times \text{mispredict\_penalty}$ ). To achieve this goal, Iyengar's team added several features. First, to cover the additional branches that have to be predicted because of the longer pipeline, the size of the branch-history table (BHT) was quadrupled from 512 to

2,048 two-bit entries, and the size of the branch-target instruction cache (BTIC) was quadrupled by doubling the number of entries from 64 to 128 and doubling the capacity of each entry from two instructions to four.

Next, Iyengar's team added an eight-entry return stack to predict subroutine return addresses. Finally, unlike the 7400, which allows only one speculative instruction stream to be in progress at a time, the 74xx allows up to three, resolving all three pending branches in a single cycle. Motorola says the combined effect of these improvements was to boost prediction accuracy by about 5% on SPECint95.

### On-Chip L2, Off-Chip L3 Cuts Latency

Because the latency of L1-cache accesses had to be increased by one cycle for frequency's sake, something had to be done to keep average memory latency from suffering. Motorola's solution was to add a level to the cache hierarchy. Unlike Intel's Coppermine, which simply moved the external L2 cache onto the processor, Motorola not only added an on-chip L2 but also kept the 7400's external cache as an L3, as Figure 3 shows. The off-chip L3 is one of the most obvious distinctions between Coppermine and the 74xx.

The 74xx's on-chip 256K L2 cache is eight-way set-associative and uses a copyback write policy. The cache is non-blocking, allowing L2 hits to be serviced while prior L2 misses are being brought from the L3 or main memory. The cache-access time is six cycles beyond an L1 miss, but the cache is fully pipelined and able to transfer a complete 32-byte line to the CPU and to the L1s every cycle, giving it an impressive bandwidth of more than 22 GBytes/s at 700 MHz. Data in the cache is protected by byte parity.

The L3 cache is similar to the L2 on the 7400. And, like the 7400, which was offered both with and without an L2 (750 and 740, respectively), the 74xx will likely also be offered both with and without an L3 (presumably the 7450 and 7440 if the naming convention holds). A 7440 would surely be more popular than the 740, however, since the on-chip L2 will

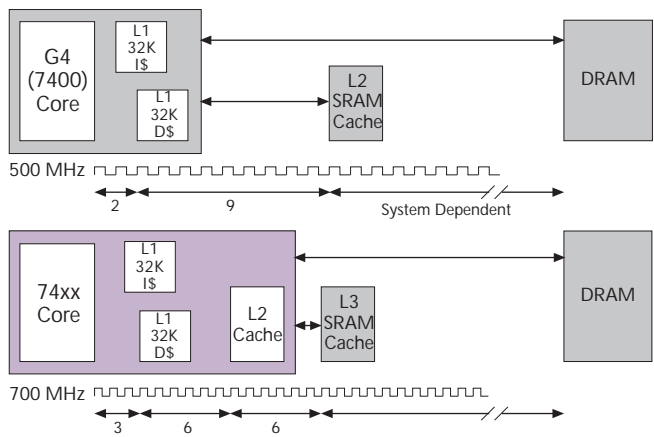


Figure 3. Even though the L1 load-use penalty had to be extended one cycle, the average load-use penalty was kept about the same (in cycles) by adding a fast 256K on-chip L2 cache as well as an external L3 cache.



be more than adequate for most desktop systems and nearly all mobile systems.

As with the 7400's L2, the 74xx's L3 tags are completely on chip, allowing hit/miss detection and cache-coherency operations to be performed at on-chip speeds. The L3 tags are two-way set-associative with 8,192 tags per way. The cache lines are each 32 bytes and unsectored for a 512K cache; 64 bytes and two-way sectored for a 1M cache; and 128 bytes and four-way sectored for a 2M cache. (A sectored cache maintains one tag per line but separate status bits for each sector in a line.)

As on the 7400, the external-cache interface is optionally either 64 or 128 bits wide. The interface operates at a variety of half-clock bus ratios and is fully pipelined. Data is fetched critical doubleword first and is immediately forwarded directly to the CPU to minimize L3 latency. The load-use penalty for an L3 hit is 15 cycles with the cache operating at half speed. The L3 interface has been improved over the 7400's L2 interface to support high-performance DDR and late-write SRAMs. Low-cost pipelined-burst (PB2) SRAMs and PC-DDR SRAMs are also supported. The L3 data and addresses are protected with byte parity.

The 74xx's system bus is virtually identical to that used on the 7400. This bus, however, just went through a major upgrade between the G3 and the G4, and Motorola now refers to it as the MPX bus rather than the 60x bus, as it was previously called. The MPX bus, while remaining backward compatible with the 60x, adds full support for out-of-order transactions, increases pipelining depth, eliminates dead cycles between many transactions, improves data streaming, and offers a 128-bit option for the data bus.

Apple has reported that even in its 64-bit form, the MPX bus delivers from two to three times the sustained throughput of the G3's 64-bit 60x bus at the same frequency. This improvement gives a significant performance boost to applications like Photoshop that process huge image files. The 7400 offers bus-clock ratios from 3× to 9× in half-clock increments. Motorola says the bus can operate well above the current 100 MHz, but, in typical fashion, Apple has failed to deliver chip sets that can take advantage of higher clock rates.

The only memory-system feature Motorola gave up in going from the 7400 to the 74xx is the reserved state in the 7400's five-state MERSI (modified, exclusive, reserved, shared, invalid) cache-coherence protocol. The "R" state offered a shared-intervention protocol with the capability for direct cache-to-cache transfers between processors. The state, however, added significant complexity to the multiprocessor logic and was deemed not to have enough performance benefit to justify carrying it forward. Thus the 74xx implements only the more conventional four-state MESI protocol—not a huge loss.

## Embedded Features to Boot

Motorola long ago recognized that Apple didn't represent a large enough market to justify all the investments Motorola would have to make to keep PowerPC chips competitive. As a result, it has made a major push into the high-end embedded-processor market with PowerPC. To support this strategy, Motorola added a few simple, but important, embedded features to the 74xx: parity protection and cache way-locking were added to the L1s; the L3-cache tags can be disabled, allowing the L3 port to be used for a high-speed 2M memory; and a software tablewalk was added to support a wider variety of memory-management systems, such as those found in embedded real-time operating systems.


A somewhat surprising feature added for the embedded space was 36-bit physical-memory addressing. This feature was needed to allow the 74xx to be used in large RAID systems as well as in large network switches (see MPR 5/10/99, p. 9). The feature was simple to implement; four unused bits in the page-table descriptors were simply implemented in the TLBs and pinned out on new address-bus wires.

Not a feature of the chip per se, but nonetheless a valuable capability for serving the embedded space, is the 74xx's modular design. In this approach, each major function block is designed as an independent module. With this technique, Motorola can rapidly produce 74xx derivatives with, for example, an SDRAM or RDRAM interface instead of an L3, a larger or smaller L2, or a different bus. Motorola did not say how far the concept extends into the CPU core. It is not clear, for example, whether the floating-point or AltiVec units are modular.

## Good, But Good Enough?

Although the 74xx is a significant improvement over the 7400, it is not likely to deliver the Pentium III—toasting performance that Steve Jobs would like to claim. The new seven-stage pipeline will provide some frequency relief, but even allowing for the 74xx's predecode and the stages it saves by not having to decode x86 instructions, the 74xx's effective pipeline length is still a couple stages shy of Pentium III's and Athlon's. Since Motorola's HIP6 process is not likely to be sufficiently faster than Intel's P858 to make up the difference, and since AMD will also use the same HIP6 process for Athlon, we expect the 74xx to lag both those chips in frequency by 10–20%. This handicap will probably prevent the 74xx from gaining any clear performance advantage over x86 PCs on general-purpose integer or floating-point code.

Even with this frequency deficit, however, on multimedia applications the 74xx should enjoy a performance margin over Pentium III and Athlon similar in size to the one the 7400 has over Katmai today. Unfortunately for Motorola



MICHAEL MUSTACCHI

Naras Iyengar, speaking at the Forum, described the pipeline enhancements his team is making to the G4.

## Price & Availability

No products have been announced using the 74xx architecture, but we expect a 74xx chip to begin shipping around mid-2000. The current 7400-450 lists for \$355 (quantity 1,000), while the 400- and 350-MHz versions go for \$275 and \$210 respectively. The 7400-500 will be available in 1Q00. For more information on the 7400, go to Motorola's Web site at <http://motorola.com/SPS/PowerPC/products/semiconductor/cpu/7400.html>.

and Apple, however, this advantage will be hard to prove, as there are no legitimate cross-platform multimedia benchmarks to rely on. Furthermore, to realize the full effect of the 74xx's multimedia hardware, Apple and its ISVs (independent software vendors) must invest the same time and effort to recode for AltiVec that Windows ISVs are investing in SSE and 3DNow. That, it would seem, is unlikely.

To match the performance of Pentium III and Athlon, Motorola must get more aggressive than it has with the 74xx. According to an internal roadmap that accidentally found its way to the Web late last year, Motorola had originally slotted a chip multiprocessor (CMP), code-named V'Ger, for 1H00. We suspect that the 74xx core is simply the remnant of that design; Motorola probably exercised its modular design capability to eliminate the other core(s) under cost pressure from Apple. The pressure might backfire; although Apple may get the low cost it sought, it will certainly not get the Pentium III-crushing performance it needs. When Intel ships Willamette late next year, it's the 74xx that will be crushed.

The upside, of course, is that the modular design of the 74xx core and caches was probably accomplished with the CMP design in mind. If so, Motorola might be able to quickly resurrect a CMP based on the 74xx core. Indeed, maybe Motorola is saving the G5 moniker for just such a part. A dual-processor die would probably be practical, thanks to the 74xx's small size. A two-core version of the 74xx with double the L2 cache and L3 tags would still probably be less than 170 mm<sup>2</sup>, about 8% smaller than today's Athlon.

Software wouldn't be an obstacle either; MacOS 9 has already been retrofitted with preemptive multitasking and full multiprocessor support, and the MacOS X kernel was MP ready from the get go. Since the 74xx is not likely to keep up with either Intel's or AMD's parts in frequency, and since multimedia is still too small a niche for the performance advantage of AltiVec to offset the difference, Apple and Motorola may want to revisit the CMP strategy.

A dual-processor chip, especially for Apple's traditional publishing markets, could offer a real performance advantage, and a marketing advantage as well. Applications in these markets typically have plenty of thread-level parallelism

Feature	Motorola			Intel
	G3 (750)	G4 (7400)	G4 (74xx)	Pentium III
Instr Decode	2 + 1 br	2 + 1 br	3 + 1 br	3 x86
Instr Issue	6 instr	6 instr	8 instr	5 ROPs
Instr Thruput	2 + 1 br	2 + 1 br	3 + 1 br	3 ROPs
Reorder Depth	6 instr	6 instr	16 instr	40 ROPs
Pipeline (int/ld)	4/5 cycles	4/5 cycles	7/9 cycles	12/14 cy
BHT	512 × 2b	512 × 2b	2,048 × 2b	512, 2-level
BTIC	64 × 2	64 × 2	128 × 4	512 (BTAC)
16-bit Int SIMD	None	12 GOPS	17 GOPS	5.5 GOPS
FP SIMD	None	4 GFLOPS	5.6 GFLOPS	3 GFLOPS
L1: I/D (ways)	32K/32K (8)	32K/32K (8)	32K/32K (8)	16K/16K (4)
Load Use	2 cycles	2 cycles	3 cycles	3 cycles
Int L2 (ways)	None	None	256K (8)	256 (8)
Load Use	-	-	9 cycles	7* cycles
Bandwidth	-	-	22.4 GB/s	11.7 GB/s
Ext Cache	1M (2-way)	2M (2-way)	2M (2-way)	None
Load Use	11 cycles	11 cycles	15 cycles	-
Bandwidth	1.6 GB/s	4 GB/s	5.6 GB/s	-
Sys Bus (max)	800 MB/s	1.6 GB/s	1.6 GB/s	1.1 GB/s
Transistors	6.5 million	10.5 million	33 million	28 million
IC Process	0.27μ 5-AI	0.22μ 6-Cu	0.18μ 6-Cu	0.18μ 6-AI
Die Size	67 mm <sup>2</sup>	83 mm <sup>2</sup>	95 mm <sup>2</sup>	106 mm <sup>2</sup>
Power (max)	5W @2.5V	14W* @1.8V	14W* @1.5V	19W* @1.6V
Frequency	400 MHz	500 MHz	700 MHz	733 MHz
SPEC95b Int/FP	18.8/12.2	23.8/21	37*/35*	35*/27*
Production	Now	Now	Mid-2000*	Now
Est Mfg Cost*	\$30	\$35	\$40	\$40

Table 1. The 74xx will have a longer pipeline than the previous 750 and 7400 processors as well as a fast on-chip L2 cache like Intel's new Coppermine Pentium III. (Source: vendors except \*MDR estimates)

(TLP) that a tightly-coupled CMP could exploit for a performance benefit. Sun, with its new MAJC 5200 chip (see MPR 10/25/99, p. 18), is taking the CMP approach on just this basis. Although Motorola will definitely be on the defensive with a 700- or 800-MHz 74xx against a 1-GHz Willamette, it may just be able to take the marketing high-ground with two 800-MHz processors on a chip. (In Steve Jobs's marketing terms, that's 1.6 GHz.)

While Apple may try to delude itself into believing Macs don't compete with x86 PCs, in reality, they do. Unfortunately for Apple, Macintosh application software is mostly a subset of Windows software, and the customers Apple covets are the same as those targeted by Compaq, Dell, and other PC vendors. For some customers, primarily at the low end, Apple has demonstrated it can sidestep the issue of processor performance by applying sexy industrial design. For more sophisticated customers, however, the issue of processor performance is not so easily swept under the rug.

If a 74xx processor were to materialize early next year at 700 MHz, it would probably be within Apple-marketing distance of Pentium III-733, assuming Apple plays its Velocity Engine card to the full extent. By mid-2000, however, both Pentium III and Athlon could be over 800 MHz, making a 74xx-700 far less compelling. But if Motorola can deliver early and then stay within 5%, or at most 10%, of Intel and AMD on frequency, the 74xx should keep the Apple platform competitive on general-purpose applications, and well ahead on multimedia applications, for most of 2000. □