Jalapeno Powers Cyrix's M3

Hot New Core Drives Fastest Integrated PC Processor Yet

by Keith Diefendorff

In pursuit of its vision of a PC on a chip, Cyrix (a division of National Semiconductor) previewed details of the M3 and its hot new Jalapeno core at Microprocessor Forum last month. The M3, due in 1H00, integrates the CPU core, 256K of L2 cache, a Direct RDRAM interface, and a 3D-graphics rendering pipeline on a single 0.18-micron die.

With Jalapeno, Cyrix seeks to reduce the frequency gap between its current processors and those offered by Intel and AMD. As Figure 1 shows, Cyrix's MediaGX/MMX, based on the 5x86 core, currently tops out at 266 MHz, while the MXi, with its Cayenne core (see MPR 10/27/97, p. 22), is expected to reach 350 MHz in 1Q99. By that time, AMD and Intel will ship 400- and 500-MHz processors, respectively, and by 1Q00 we expect both to field processors of 700 MHz or more. Jalapeno, Cyrix hopes, can match the clock rates of these high-end processors. This is a change in strategy for Cyrix, whose previous cores sought high IPC (instructions per clock), trading off frequency.

Unlike Intel's P6, AMD's K7 (see MPR 10/26/98, p. 1), and Rise's new mP6 (see MPR 11/16/98, p. 1), all of which decode and issue up to three x86 instructions per clock, Jalapeno does only two. Greg Grohoski, project manager for Jalapeno, says that the two-issue design requires significantly less silicon area yet achieves nearly the same performance in practice. His claim may be credible in light of the miserably low instruction-level parallelism (ILP) in PC software and the latency of PC memory systems.

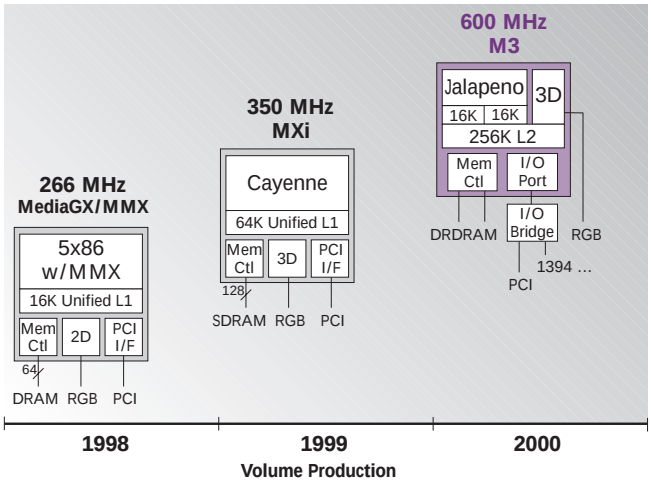

Figure 1. The M3, the third in a series of integrated processors from Cyrix, integrates the new Jalapeno core with 256K of L2 cache, a 3D graphics engine, and two Direct RDRAM channels.

Jalapeno's 11-stage pipeline is designed to achieve a clock rate of more than 600 MHz in National's 0.18-micron CMOS-9 process, even though that process is not as fast as Intel's 0.25-micron P856.5 (see MPR 9/14/98, p. 1). The M3 die will weigh in at about 120 mm^2, which, in its 400-contact CBGA package, will cost about $60 to manufacture, according to the MDR Cost Model. Since CMOS-9 will have been in production for a year by the time the M3 samples in 4Q99, National should have no trouble achieving yields consistent with a high-volume, low-cost product. Cyrix expects the core to reach 1 GHz in National's 0.15-micron copper CMOS-10 process, which should be in production before the M3 samples, allowing a rapid migration to that process.

Although sources indicate that Cyrix has been contemplating a Slot 1–compatible processor, the company gave no indication that Jalapeno would ever appear in anything other than the integrated M3 chip disclosed at the Forum. With its on-die L2 cache, an M3-like derivative would probably perform well in a Socket 7 environment. But in 2000, rigor mortis will be gripping the Socket 7 market, making it unlikely that Cyrix will invest time and effort in such a product.

More Out of Order Than Before

Jalapeno improves on previous Cyrix cores by offering more instruction-reordering capability. Whereas the dual-pipeline 6x86 and Cayenne cores allowed execution past an instruction stalled in one pipe, Jalapeno provides six pipelines that can each continue execution while others are stalled. Registers are renamed to avoid false pipeline hazards.

Because the execution-unit pipelines process instructions in order, Jalapeno's reordering capability is less extensive than either the P6's or the K7's. But Jalapeno's design will cover many common pipeline stalls, giving it more tolerance to long-latency operations than previous Cyrix processors.

Jalapeno provides for limited memory-access reordering with a 12-entry store reservation station. This facility allows stores to wait on data without blocking load accesses to the cache. Store data in the reservation station can be forwarded directly to subsequent loads, eliminating their cache access. Store forwarding is an important feature for x86 processors, which must often store and immediately reload data due to the limited number of architectural registers.

The L1 data cache supports hits-under-misses, with loads accessing the cache out of order. Load data, however, is queued and returned to the execution units in strict program order, increasing the average latency of loads somewhat. As with instruction reordering, Jalapeno's memory reordering is less aggressive than the P6's or the K7's, but it is probably adequate to address many of the important cases.

Memory Hierarchy Minimizes Latency

Since Jalapeno's limited instruction and memory-access reordering capabilities give it little tolerance to memory latency, the M3 invests a large number of transistors in a low-latency memory system. As Figure 2 shows, the processor uses a three-level memory hierarchy consisting of 16K L1 instruction and data caches, a 256K on-chip level-two cache (L2), and a direct connection to main memory.

The instruction cache is four-way set-associative with 32-byte lines and sports a 256-bit access path. The cache controller handles up to two pending misses. The data cache is also four-way set-associative with 32-byte lines, but it is dual ported, allowing one read and one write each cycle. The cache is truly dual ported, so there are no bank conflicts, as can occur with the P6's and the K7's multibanked caches.

The data cache implements a write-through store policy. Write-through caches are normally less efficient than write-back caches, because every store must be written to the next higher (and slower) level of the memory hierarchy. The M3's on-chip write-back L2 cache, however, makes the performance loss negligible, and the write-through L1-cache design is simpler. The data-cache controller can queue up to four outstanding misses before it stalls the load/store pipeline.

Effective addresses are translated to physical memory addresses via separate instruction and data TLBs. Both TLBs have only 32 entries but are fully associative, giving them a hit rate similar to larger set-associative TLBs. The primary TLBs are backed by a unified 512-entry four-way set-associative second-level TLB.

The 256K L2 cache is eight-way set-associative and eight-way interleaved. The interleaving allows the cache to service two accesses to different banks on each cycle—one 256-bit read and one 64-bit write. Up to eight transactions can be queued waiting to access the L2. A simple prefetch engine watches L2 transactions and prefetches cache lines into the L2 when it detects certain access patterns, such as a linear sequence of addresses.

The two L2-cache ports are shared between the CPU core and the 3D-graphics unit. Although 3D-unit accesses and core accesses can interfere, the separate read/write ports, the wide data paths, a priority scheme, and the core's access queue avert a bottleneck. Cache ways can be locked, giving software a means for optimizing cache usage.

Any of the cache's eight ways can be dedicated to the 3D unit for texture cache or a compositing buffer. Allocating a single way gives the 3D unit twice the buffer space provided on ATI's new Rage 128 (see MPR 9/14/98, p. 16), which is the fastest single-chip 3D engine on the market today.

As Figure 3 shows, the load-use penalty for an L1 cache hit is four cycles—an unusually long penalty. While this long access time may ensure that data cache accesses do not limit frequency, it will cause pipeline stalls that hurt IPC—especially in x86 code, which depends heavily on memory in lieu of registers. Cyrix says this penalty reduces performance by only a few percentage points compared with the three-cycle penalty in, for example, the K7. But it also whispers that it may try to remedy the situation in a future revision of the chip.

In addition to the L1 load hit time, seven more cycles are required for L1 misses that hit in the L2. L2 misses incur an 11-cycle penalty plus however many cycles are required to access the DRAMs. Cyrix says the on-chip DRAM controller cuts 30 ns off the best memory-access time that can be achieved with a DRAM controller on a conventional north-bridge chip. The M3's controller can keep 32 DRAM pages open at a time, the same as Intel's high-end 440BX.

At the Forum, Cyrix described the M3 with a two-channel Direct RDRAM memory interface. But the company

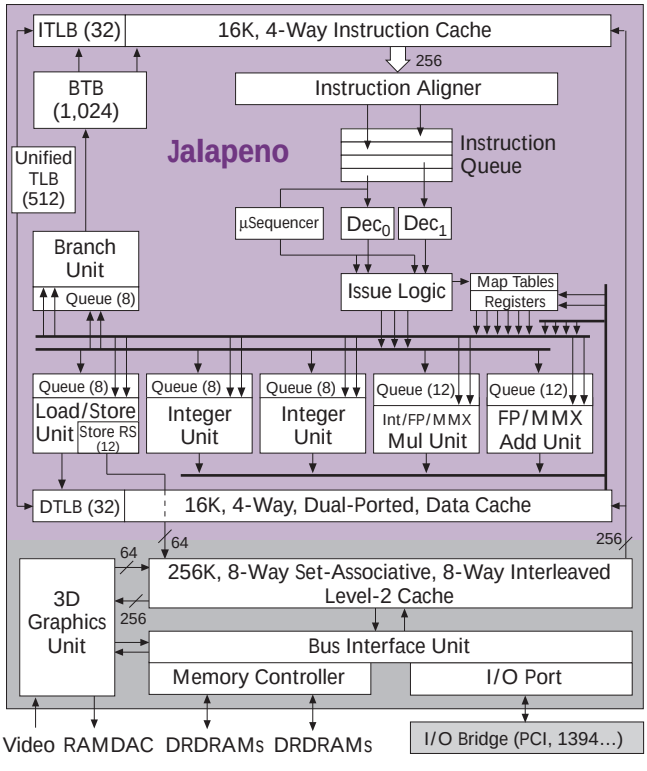

Figure 2. The M3 integrates the Jalapeno core with a 256K L2 cache, a two-channel Direct RDRAM interface, and a 3D-graphics rendering pipeline on a single die.

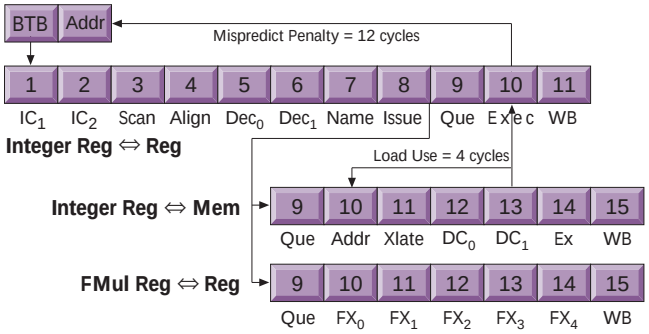

Figure 3. Jalapeno's long 11-stage pipeline should enable it to achieve its target frequency of more than 600 MHz in 0.18 micron.

says that the memory controller module could easily be replaced by a 128-bit SDRAM or an SLDRAM interface, depending on customer demand.

Pipelined for Speed

Many architects seem to be converging on similar microarchitectures for the front end of an x86 pipeline. The P6, the K7, and now Jalapeno all use the first few stages of their 10- to 14-stage pipelines to fetch, align, and decode x86 instructions into RISC-like operations (ROPs). This conversion to ROPs was avoided in previous Cyrix processors and is an approach for which previous Cyrix architects have publicly expressed disdain (see MPR 9/11/95, p. 17). Possibly for this reason, Cyrix is careful not to label Jalapeno's ROPs as such, but instead calls them "nodes." For all intents and purposes, however, nodes appear similar to other vendors' ROPs.

In stages 1 and 2 of Jalapeno's pipeline, 32 bytes are fetched from the instruction cache and deposited into a linear 96-byte buffer. In stage 3, the buffer is inspected to locate the beginnings and ends of the next two x86 instructions. Unlike the P6 and the K7, Jalapeno does not predecode instructions on the way into the instruction cache to locate their boundaries. Cyrix says predecoding is totally unnecessary with its two-issue design—even at these high frequencies—and prefers to save the extra 35% instruction-cache area that would be required to hold the predecode alignment information.

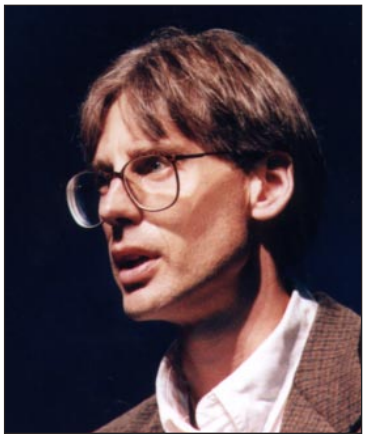

MICHAEL MUSTACCHI

Chief architect Greg Grohoski describes Cyrix's M3 and its new Jalapeno core at the Forum.

In stage 4 of the pipeline, the two x86 instructions are aligned and placed in a four-deep instruction queue, each entry of which can hold an x86 instruction of any length. In this queue, the various instruction fields are aligned by type (e.g., the *modR/M* fields are all aligned) for easy decoding. In stages 5 and 6, the instructions are decoded into ROPs.

Jalapeno has two parallel decoders that are symmetric, with the exception that only one can handle a memory operation. While this restriction would appear to present a decode hazard, Cyrix claims that it doesn't in practice. The reason is that instructions tend to occur in sequences of memory ops separated by one or more register-to-register ops, so once the sequence gets aligned, decode flows smoothly.

According to Grohoski, the asymmetry of the decoders, surprising as it may sound, does not show up as a factor in simulations of chip performance. This could be true because the decoder restriction is similar to the P6's, and compilers try to avoid offending sequences. Such limitations, however, have a nasty way of surfacing in the critical inner loop of important applications, seriously degrading performance. It is interesting that AMD went to great lengths to avoid this restriction on the K7, indicating some differences in opinion between architects on this issue.

Jalapeno's decoders can each produce three types of ROPs: integer, FP/MMX, and branch. One of the decoders can also generate a memory ROP. ROPs each comprise an opcode, a source, and a destination specifier.

Each x86 instruction falls into one of three classes. The first includes instructions that produce a single result, such as reg-op-reg or load-op-store. These are handled in both decoders, so Jalapeno can decode two per cycle (as long as only one is a memory op). The second class includes instructions that generate two results, such as pushes and pops that update the stack pointer and also perform a memory operation. These instructions require both decoders, so they are processed at a rate of only one per cycle. The third class includes such instructions as string ops and far calls that require microcode sequencing. These instructions iterate, theoretically generating up to seven ROPs per cycle, although three to four are more typical.

Renaming Supports Speculation

In stage 7 of the pipeline, registers are renamed to avoid anti- and output dependencies (register name conflicts). Jalapeno provides 64 physical integer registers, 32 of which are initially assigned to the eight x86 architectural registers and to a number of hidden and temporary registers. The remaining 32 comprise the free pool, out of which new destination registers are allocated as instructions are issued. Registers are returned to the free pool as instructions complete and retire from the machine. A parallel structure is used for the floating-point registers.

Renaming is performed using a two-level register-mapping scheme. Two register-mapping tables track how the architectural registers are mapped onto the physical registers. The first table, called the speculative map table, defines the current mapping. Instructions locate operands by using the architectural register specifier as an index into the speculative map table to get the name of the physical register containing the desired operand.

The second table , or committed map table, tracks the mapping of completed instructions. Thus, on any exception, interrupt, or branch misprediction, the machine can instantly be rolled back to its precise architectural state by a broadside transfer of the committed map table into the speculative map table. This transfer repairs all modifications made to the machine state by instructions that executed speculatively or completed out of program order.

Dual Dispatch to Six Units

Jalapeno has six execution units: two integer units, one multiply unit, one floating-point/MMX ALU, one load/store unit, and one branch unit. The multiplier and floating-point

units each have 12-entry queues; the others have 8-entry queues. As many as three ROPs can be dispatched into the instruction-unit queues on each cycle. Each queue entry holds the operation type and the operand tags indicating where operands will come from.

ROPs are issued from each queue to the associated execution unit in the same order they were placed in the queue (FIFO). One ROP can issue from every queue each cycle, putting six ROPs per cycle into execution. Operand availability is tracked by presence bits in the dispatch logic. ROPs waiting to issue will read their operands from the register file if the corresponding presence bits are set; if not set, the ROP waits until the operands become available. At that time, the needed operands are forwarded from the result buses indicated by the operand tags, and the ROP issues. Forwarding paths are provided from all four result buses to all execution units, so ROPs issue as soon as possible.

The two integer units are identical; both are fully pipelined with single-cycle latency for all instructions except carry-form rotates and bit scans (BSF and BSR), which have a two-cycle latency, and integer divides, which are not pipelined and iterate at one bit per cycle with early out.

The multiply unit handles all integer, x87 floating-point, MMX, and 3DNow multiplies, as well as floating-point divides and transcendentals. The unit is fully pipelined for multiplies of all precisions. Table 1 gives the latency of various operations in all of the units.

Two-Level Predictor Simpler Than Most

Jalapeno implements a 1,024-entry four-way set-associative branch target buffer (BTB) that is accessed in the first stage of the pipeline in parallel with the instruction-cache fetch. A 16-entry return stack predicts subroutine return address.

The BTB index is a byte address that is hashed to achieve uniform usage of BTB entries. Each BTB entry contains a 7-bit history and the target address of the next sequential branch to be predicted. Each entry also identifies the last byte of the branch, which is used to re-index the BTB for the next prediction. Branches that have not previously been mispredicted, and are therefore not in the BTB, are, by default, predicted to be not taken. This scheme allows Jalapeno to prefetch instructions ahead of execution, although not as far ahead as the more aggressive scheme Rise uses in the mP6 (see MPR 11/16/98, p. 1).

Unlike two-level predictors that use a global branch history to access pattern tables of dynamic 2-bit saturating up-down counters, Jalapeno's pattern table is fully static. Jalapeno simply uses its 7-bit history field to index a 128-entry static-prediction ROM. This scheme is simpler than dynamic ones, and Cyrix says the saving in complexity is easily worth the small 1–2% loss in prediction accuracy over a 2-bit dynamic predictor.

Jalapeno's branch unit processes branches in strict program order. The unit maintains an eight-entry queue to reorder condition flags, which can return from the execution units out of order. The condition flags are renamed using a scheme similar to that used for registers.

Branches that turn out to have been correctly predicted update the BTB history bits without introducing any bubbles into the instruction pipeline. Mispredicted branches incur a penalty of at least 12 cycles to compute the correct fetch address, update the BTB history bits and the target address (if the address was also mispredicted), start fetching the correct instruction stream, and refill the instruction pipeline. Such a harsh misprediction penalty requires an accurate branch predictor.

3DNow, High Bandwidth Supply 3D Engine

For 3D geometry and lighting calculations, the M3 relies on Jalapeno's 3DNow capability (see MPR 6/1/98, p. 18), which can deliver up to 2.4 GFLOPS of floating-point horsepower at 600 MHz. Even at 50% efficiency, this FP performance should be ample to feed the M3's 3-Mtriangle/s rendering pipeline—at least for simple lighting situations. A benefit of the M3's architecture is the tight coupling between the Jalapeno core and the 3D rendering engine, both of which share a common memory system, beginning at the L2 cache. The L2 cache provides a total of 24 Gbytes/s of bandwidth that is shared between the core and the 3D unit.

With two Direct RDRAM channels, the M3 provides a total of 3.2 Gbytes/s of memory bandwidth—twice that provided by today's high-end 3D graphics chips. Similar to the MediaGX and MXi, the M3 implements a unified memory architecture (UMA), which locates the frame buffer in main memory. This organization could require up to 300 Mbytes/s of memory bandwidth just to refresh the display, but Cyrix avoids this requirement by maintaining a shadow copy of the frame buffer in compressed form (see MPR 3/10/97, p. 1). The refresh controller reads only the compressed frame buffer, reducing refresh bandwidth requirements by a factor of 20 and essentially eliminating this source as a significant consumer of main-memory bandwidth.

3D Performance Rivals Dedicated Graphics Chips

Although the M3's 3D-rendering performance may not match that of high-end discrete graphics chips in 2000, it should provide enough 3D performance to satisfy the

Operation	Latency/Throughput			
	Integer	x87 FP	MMX	3DNow
Add/Sub/…	1 / 1	4 / 1	1 / 1	3 / 1
Shifts	1 / 1	–	2 / 1	–
Rotate/Bit Scan	2 / 1	–	–	–
Multiply	5 / 1	5 / 1	4 / 1	5 / 1
Divide (single)	32 / 32	23 / 23	–	23 / 23
Divide (double)	–	33 / 33	–	–
Est (1/x, 1/√x)	–	–	–	5 / 1
Load, L1 Hit	4 / 1	4 / 1	4 / 1	4 / 1
Load, L2 Hit	11 / 1	11 / 1	11 / 1	11 / 1

Table 1. All of Jalapeno's execution units (except divides and transcendentals) are fully pipelined, giving them a throughput of one instruction per clock. (Source: Cyrix)

appetite of a large number of PC users. While discrete 3D-graphics chips will have transistor budgets that the M3 cannot match, its architecture has efficiency advantages that will give it competitive 3D performance with fewer transistors.

Cyrix's philosophy on 3D is different than that of other 3D-graphics vendors. These vendors typically use a standard-cell design, build their parts in a previous-generation process technology, clock their pipelines at a slow rate, and compensate with a large die and many millions of transistors. In contrast, Cyrix will custom design the M3's 3D unit, employ the latest process technology, and run much of the pipeline at the core CPU frequency. With this approach, as Table 2 shows, the M3's 3D unit will use only 3.3 million transistors (not including the L2 texture cache)—significantly fewer than the 8 million found in ATI's Rage 128.

One advantage of putting 3D rendering on chip with the CPU is that it will scale with process technology. CPU processes have historically stayed between one-half and one full generation ahead of other logic devices. This is likely to remain true, as CPUs are now the official process driver for National's fabs, as well as for many others. This ensures that the M3's 3D unit will always get leading-edge process technology and that its geometry and rendering throughput will stay in balance as technology advances.

At this time, Cyrix is unwilling to release details on the M3's 3D unit, citing proprietary technology. The company did say, however, that the unit includes triangle-setup and rasterization hardware that will exceed 3 Mtriangles/s and 266 Mpixels/s, putting it in the same category as the Rage 128. This would make it a high-end 3D chip today, but probably only an entry-level chip by 1Q00.

The M3's 3D unit uses industry-standard APIs, including a virtual AGP interface, and provides modern 3D features, such as independent fog and alpha blending, anisotropic texture filtering, and antialiasing. The M3's unit also includes MPEG-2 and DVD decoders from Mediamatics (a National acquisition) and a VGA-compatible display and video controller capable of driving an external RAMDAC at speeds up to 230 MHz (dot clock). Notably absent is an LCD controller for notebook displays.

Feature	Cyrix MXi	Cyrix M3	AMD K7	Intel † Coppermine
x86 Decode	2 complex	2 cpx (1 mem)	3 complex	1 cpx + 2
Issue Rate	2 x86	6 ROPs	9 ROPs	5 ROPs
Reorder Depth	1 x86	32 x86	72 x86	20 ROPs
Pipeline Stages	7 stages	11 stages	10 stages	12–14
BTB Entries	512 entries	1,024	2,048 × 2b	≥ 512
Return Stack	8 entries	16 entries	12 entries	4 entries
L1 Cache	64K unified	16K / 16K	64K / 64K	16K / 16K
On-chip L2	None	256K	None	256K
ITLB/DTLB	16 + 384	32 / 32	24/32+512	32 / 64
Clock Rate	350 MHz	>600 MHz	>500 MHz	>600 MHz
Transistors (-3D)	6.5 million	25 million	22 million	24 million
IC Process	0.25µ 5M	0.18µ 5M	0.25µ 6M	0.18µ 6M
Die Size	90 mm²	120 mm²	184 mm²	120 mm²
Production	1Q99	1Q00	1H99	2H99
Est Mfg Cost*	$55	$60	$105	$65
DRAM Ctl	On-chip		N/A	Camino †
DRAM Type	SDRAM	2× DRDRAM	N/A	DRDRAM
DRAM B/W	1.6 GB/s	3.2 GB/s	1.6 GB/s	1.6 GB/s
Open Pages	16 pages	32 pages	N/A	32 pages
3D	On-chip		E.g., with ATI Rage 128	
Geometry	3DNow 1.4 GFLOPS	3DNow 2.4 GFLOPS	3DNow 2.0 GFLOPS	KNI 2.4 GFLOPS
Setup	2 Mpolys/s	>3 Mpolys/s	3.2 Mpolygons/s	
Fill Rate	120 Mpx/s	>266 Mpx/s	200 Mpixels/s	
Texture Cache	None	32K	8K	
3D Transistors	2.5 million	3.3 million	8 million	

Table 2. The M3 provides on-chip DRAM interface and 3D capabilities that current AMD and Intel processors must provide externally. (Source: vendors, except MDR *estimates and †projections.)

New Port Provides Flexible I/O

As every vendor of integrated CPUs has discovered, putting I/O functions onto the processor is problematic. While saving very little cost, I/O integration restricts the flexibility of—and thus the market for—a chip. To avoid this, Cyrix's MediaGX and MXi both used PCI as an interface to a south-bridge I/O chip. But next-generation I/O, such as IEEE-1394, Gigabit Ethernet, and Fibre Channel, can tax a single PCI bus.

In response, Cyrix is developing a new I/O port for the M3. The company is not ready to disclose details on the port but indicated that it is a low-pin-count, high-speed port that will connect to an I/O chip that will provide a PCI interface and other high-speed I/O functions. Cyrix is working with third-party chip-set vendors to standardize the new port and establish infrastructure support for it.

The M3's caches support the MESI cache-coherence protocol, allowing coherent I/O through this new port. Cyrix has no ambition, however, toward multiprocessor systems, and, accordingly, the new I/O port provides no MP support.

Different Design Philosophy, Different Results

Cyrix has followed a different philosophy with the M3 and Jalapeno than Intel followed with the P6 or AMD will follow with the K7. Whereas these processors were designed to be the biggest, baddest machines on the block, Jalapeno was designed for the most bang for the buck. The signs of this philosophy are evident throughout the design.

By issuing only two instructions per clock, rather than three, Jalapeno should save considerable area (complexity generally increases with the square of the issue width). Jalapeno's limited reordering capability and the simple branch predictor may sacrifice some additional IPC relative to the K7, and probably even the P6, but should also reduce silicon area. The Jalapeno core (without the M3's L2 or 3D unit) is about 75 mm² in CMOS 9; we estimate that Intel's Mendocino core with KNI added would be about 15% larger than Jalapeno if rendered in a process similar to CMOS 9.

These choices may pay off; with 40% to 60% of a PC processor's time spent in the OS, where instruction-level parallelism is scarce, and the rest spent in applications that have only slightly more ILP, little is gained from a high IPC

processor. And with a small, simple core, Jalapeno stands a better chance of achieving its frequency objectives than it would with a more complex core.

Having saved silicon area in the core, Cyrix can invest more transistors in the memory system and in 3D. This choice may ultimately yield more performance than if those transistors had been spent on a more complex core. Even if the M3 doesn't outperform the competition, the efficiency gained through the tight integration of CPU, L2 cache, 3D graphics, DRAM memory controller, and UMA frame buffer (with compression) should give it a cost advantage.

The M3 faces a number of obstacles. One is process technology. According to the MDR FET Performance Metric (see MPR 9/14/98, p. 1), National's 0.18-micron CMOS 9 will be 30% slower than IBM's 0.18-micron CMOS-8S and 20% slower than AMD's 0.18-micron CS50; Intel's 0.18-micron P858 will likely be similar to IBM's 8S. Cyrix, however, seems to be satisfied with National's technology and, in fact, claims even better results than with IBM, because IBM wouldn't give Cyrix access to its newest technology anyway, and National has been more responsive to Cyrix's needs.

Another problem is KNI (see MPR 10/5/98, p. 1). While Jalapeno's implementation of 3DNow may provide performance similar to Katmai's implementation of KNI, KNI is destined to win the software battle and will be well on its way to that goal by 2000. With M3 still a year from silicon, Cyrix could consider adding KNI to Jalapeno. Given the speed with which the company switched Cayenne from its proprietary MMXFP to 3DNow, such a move might be possible.

For More Information

For more information check out the Jalapeno Web site at *www.cyrix.com/html/about/1998/jalapeno.htm*.

A third problem is chip-set and motherboard infrastructure. With its new I/O port, Cyrix has created a problem similar to the one that AMD created for the K7 with its new Slot A interface. Cyrix's problem may be less severe, since the M3's I/O port is not attempting to be a generalized memory and multiprocessor bus. But the problem is still daunting, as the M3 will not connect to existing chip sets or plug into existing motherboards (unless Cyrix repackages it on an interposer with a PCI-bridge chip to fit an MXi socket).

If the M3 were nine months earlier, it would be a killer part, delivering performance on a par with Katmai's and with excellent 3D-rendering capability to boot. Such a part would have fared well even in the performance-desktop segment, allowing it to collect high margins. By early 2000, however, Intel's Coppermine and AMD's K7 will force the M3 to compete in only the sub-$1,000 basic-PC segment. In this segment, which is National's focus anyway, the M3 will be quite attractive, easily outperforming a 450-MHz Celeron with Whitney graphics. If Cyrix can achieve its frequency and schedule goals, Jalapeno's simple design and M3's high level of integration could give it a home in millions of low-cost PCs. If this is the kind of processor we can expect to see in $600 PCs, you can bet they're going to be hot! M