PC Makers Serve Pentium Before Its Time

Cache, Local Bus Likely Targets for Improvement

By Mike Feibus and Dean McCarron

After a stutter-step delay orchestrated by Intel, computer makers earlier this month finally unveiled their first systems based on the Pentium processor. Despite the unexpected delay in announcing their systems, however, it appears that most PC vendors still need more time to optimize their designs to take full advantage of the features Pentium brings to the x86 world. Some of the systems previewed by *Microprocessor Report* implemented a 64-bit data bus, some featured a write-back, second-level cache to complement Pentium's internal cache write policy, and some featured wide, 128-bit memory systems. But no system put it all together to take full advantage of Intel's latest offering.

To be sure, there are other features in Pentium that are new to x86s—such as dual data and instruction caches, branch prediction, and dual pipelines—but they are internal to the chip and require no system redesign to exploit them.

For this analysis, *Microprocessor Report* examined a half-dozen Pentium-based systems announced since Intel's embargo was lifted mid-May. These systems are listed in Table 1 below. Most of the systems contained 60-MHz Pentium processors, although all were designed to run at 66 MHz when faster chips are available.

Although we have examined each of these systems' designs, it is difficult to provide a direct price/performance comparison of the machines, for several reasons. In particular, the systems vary widely in configuration—and subsequently price—from high-end servers with gigabyte and larger disk arrays to desktop machines featuring local-bus graphics. Also, prices are still in flux for some of the machines, since many of the vendors are still awaiting volume shipments of Pentium chips.

Servers Dominate

Four of the six systems—from Acer, ALR, Digital, and Hewlett-Packard—were presented as server configurations. It is hard for most computer buyers to justify paying the stiff premiums demanded by Pentium servers today. That reluctance is magnified for information-system managers looking for file servers—in this instance the money required to buy these systems clearly would be better spent on 486-based servers with faster storage and connectivity, as file servers generally are not CPU-bound.

However, it's not surprising that computer manufacturers would try to push Pentium-based servers—at least early on. Pentium processor deliveries from Intel are expected to be few and far between this year, so most PC manufacturers are seeking to garner the most revenue from each chip.

Signs are that most PC manufacturers are content to sell 486-based systems until Pentium processors are more plentiful. Few of the systems announced during Spring Comdex in Atlanta recently were actually on display on the show floor. Most manufacturers opted for showcasing their latest 486-based servers and desktop systems. Expect systems manufacturers to be actively selling Pentium-based machines at Fall Comdex; Intel confirms that its schedule calls for more than 100,000 processors to ship in the fourth quarter alone.

Interestingly, the two purely desktop models announced so far come from Compaq and NEC, two of Intel's largest customers. ALR also announced a Pentium-based desktop machine, although the design is not targeted purely at single users. ALR's Pentium design is doing double-duty in both desktop and server configurations. The desktop configuration, labeled the Evolution V, is aggressively priced at $3600 *(see **0707MSB.PDF**)*.

One could surmise that Compaq and NEC have fared better than their competitors in securing Pentium deliveries from Intel for the remainder of this year, and are thus more willing than their counterparts to allocate at least some of their Pentium chip allocation to desktop units. Both Compaq and NEC expect to price their desktop units at $4500–$5000, about $1500 more than those

	AcerFrame 3000MP	ALR Evolution VQ	Compaq Deskpro 5/66M	DECpc 560ST	Hewlett-Packard NetServer 5/60 LM	NCR 3360 (dual Pentiums)	NEC Image p60
Estimated Price	NA	$4500	$5000	$6500	$6500	$18,000	$4500
Min RAM (Mbytes)	256	8	8	8	16	32	8
Cache Architecture	2-Way Set Assoc.	Direct Mapped	Direct Mapped	2-Way Set Assoc.	2-Way Set Assoc.	2-Way Set Assoc.	Direct-mapped
Cache Write Policy	Write-Back	Write-Back	Write-Through	Write-Back	Write-Back	Write-Back	Write-Back
RAM Data Path Width	64 bits	128 bits	128 bits	32 bits	32 bits	64 bits	64 bits

Table 1. Three of the Pentium-based systems examined were configured as servers, two as desktop machines, and one that is being offered both ways.

vendors' comparably equipped 486DX2 systems.

Memory Differentiates

Among the six Pentium-based systems we analyzed, the area of greatest differentiation lies within the memory subsystem. In fact, the systems were similar in all respects with the exception of the cache and memory architectures. This holds true for both desktops and servers.

Among the six vendors surveyed, practically every possible second-level cache architecture, write policy, datapath width, and memory size was used. The cache architecture of preference—at least for the servers—appears to be two-way set associative; Acer, DEC, and HP all implement the caches in their Pentium systems this way.

Both desktop units analyzed—from Compaq and NEC—implement direct-mapped cache controllers. ALR's offering was the only server previewed with a direct-mapped cache implementation.

It's not surprising that most of the manufacturers designing servers opted to make their second-level caches two-way set associative. In general, most OEMs don't try to cut corners in their server designs, going with higher performance whenever possible.

The most likely reason that two-way caches are so prevalent in the roundup is that there was only one off-the-shelf cache controller available in time for the first wave of Pentium systems—and it is a two-way set-associative controller. OEMs that wanted systems to announce at Spring Comdex could either use Intel's 82496 controller, or they could roll their own, as did ALR.

The high-performance 82496 chip is actually designed for systems with multiple processors. The chip's multiprocessing benefits are clearly going unused in the single-Pentium designs. However, the time-to-market advantages obviously outweighed the cost implications of using this part. (For systems based on a $1000 microprocessor, the premium-priced cache controller probably has minimal impact on the overall price.) The '496 is clearly a temporary fix, and ultimately will settle into its multiprocessing Pentium niche over time as other Pentium chip sets become more widely available.

In fact, the '496 wasn't the only way that Intel helped manufacturers get to market quickly with their initial Pentium-based systems. DEC and HP both relied on Intel for their entire motherboard design. Both companies' systems are based on Intel's 8235x EISA chip set, which explains why the DRAM datapath is only 32 bits wide.

Direct-Mapped Alternatives

As mentioned, ALR, Compaq, and NEC didn't follow the Intel cache controller crowd. Most desktop PC manufacturers include direct-mapped caches rather than opting for the higher-performance set-associative alternatives in order to shave overall system costs.

ALR's machine is noteworthy because it is the only server examined which opted for the lower-cost, direct-mapped alternative. As mentioned, ALR's choice may be explained by the fact that it offers the Evolution in a desktop configuration as well. Several other design choices indicated that ALR may be at least as interested in enticing desktop users with the Evolution. For example, the Evolution VQ was the only server to offer local-bus graphics, and it was the only one to supply lower-cost IDE hard drives instead of SCSI.

NEC's Image P60, the only PCI-based system in the group, is direct-mapped because it utilizes Intel's 82430 PCI chip set *(see **070403.PDF**)*, which includes a direct-mapped cache controller. The 82430 chip set enhances performance by combining cache tags on-chip. This chip set allows the P60 to achieve burst performance of 3-1-1-1 (three cycles for the first data transfer, and one cycle for subsequent transfers) by using synchronous SRAMs.

Compaq more than offsets the use of the lower-performance direct-mapped cache controller with its wide, 128-bit DRAM data path. At the heart of Compaq's mem-

Pentium Delay Doesn't Help PCI Proliferation

Despite a recent flurry of hardware announcements featuring PCI, *Microprocessor Report* has been able to locate only one PCI-based Pentium system to date.

That machine—a 60-MHz EISA desktop from NEC called the Image P60—was disclosed during Comdex last month in Atlanta, although it won't be commercially available until later this year. There were also a few motherboard manufacturers displaying PCI-based Pentium products at Comdex. Quotes on the availability of those boards stretched into next year.

Pentium PCs utilizing PCI undoubtedly will be unveiled over the next few months, and many Pentium machines available should feature that bus by Fall Comdex in November. The dearth of PCI-based Pentium systems at this juncture, however, is interesting because it casts some doubt on the theory that Intel delayed announcements of Pentium-based PCs to May from March to give some computer makers a chance to build systems showcasing PCI.

The irony is that if Intel was trying to help PCI along, the mission was foiled by delays in the company's own PCI system-logic chip set. Even with the extra two months, OEMs still didn't have enough of Intel's so-called "Mercury" chip set—the 82430 PCI chip set—in hand to warrant announcing systems using the parts. And no OEMs have the A1 step of the chip set, which fixes a host of bugs that impact everything from the secondary cache to graphics. The A1 revision is scheduled to be in OEMs' hands this July.

NCR's Multiprocessing Architecture Takes the High Road

Taking the high road of all the Pentium-based system announcements recently is NCR: its System 3000 family of Pentium-based systems feature the company's symmetrical multiprocessing architecture. The high-end System 3000 model can be configured with up to 16 Pentium processors.

Available now is the NCR 3360 workstation, a single- or dual-Pentium computer that carries a price tag of $10,000 to $18,000 (for single- or dual-CPU versions). NCR also unveiled several other computers in the 3000 Family, although they are not yet available. They are: the 3430, a deskside, single-Pentium server upgradable to two Pentiums; the 3455, which features up to six processors; and the 3555, which sports up to 16 Pentiums.

The 3360 is intended primarily as a server. The expansion bus is an NCR-enhanced Micro Channel Architecture. NCR's MCA bus operates at an 80-Mbytes/s data transfer rate. It can handle multiple SCSI controllers.

Rounding out its server-based features, the 3360 can interface up to 41 SCSI devices to a system. It also offers tremendous RAM capacity, and it features error-detection-and-correction (EDAC) for reliability. Although intended as a server, buyers can add an internally-designed graphics coprocessor that supports $1280 \times 1024 \times 24$ displays.

The heart of the System 3000 family is its multiprocessing architecture. The architecture features two proprietary MP buses that connect the processors with system memory and the Micro Channel interface, as shown in the figure below. The architecture carries the family up to 16 processors.

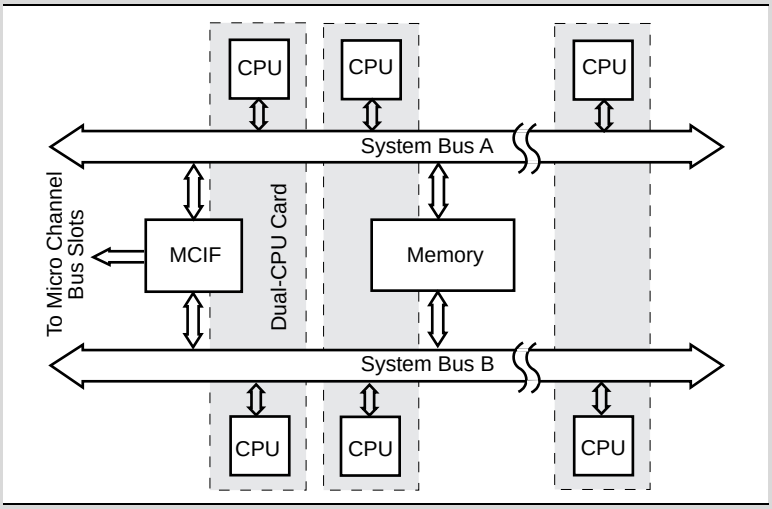

The primary memory architecture features a dual-port, two-way interleaved structure. There is a central cache directory to allow snooping across all processors. Cache coherency is handled by the memory controller using the cache directory.

Each processor has its own two-way set-associative cache—256K per processor. Because of the distributed architecture, where each processor has its own local cache using Intel's 82496 controller, the system board is relatively uncomplicated. Playing the role of traffic cop, one ASIC is used to handle the cache and two ASICs interface to the system bus.

ory subsystem—which it calls TriFlex—is a three-way memory controller that communicates via a 64-bit bus to the processor. The controller provides a 64-bit path to the second-level cache, a 32-bit path to the EISA bus, and the 128-bit path to main memory.

In fact, the wide memory bus may be a bigger factor in determining performance than the cache size, architecture, or write policy. The choices Compaq made in designing its Deskpro 5/66M fuel this argument. The Compaq and ALR systems were the only machines analyzed to implement a 128-bit DRAM data path. On the flip side, Compaq was also the only manufacturer to implement a lower-performance, write-through cache write policy. And as already discussed, Compaq opted for direct-mapped cache as well.

As a result, early benchmark tests on the new systems from computer publications such as *PC Magazine* consistently show the Deskpro leading the pack in performance. At this time, no BAPCo or SPECmark measurements have been published for any of these systems.

In contrast to Compaq, Acer chose a 64-bit data path from cache to DRAM, mirroring the 64-bit data bus that all the machines have in common.

DEC and HP both selected a 32-bit cache-to-DRAM data path. While this choice was undoubtably made for expedience—rather than designing their own motherboards, both companies chose to take advantage of Intel's Xpress design that features a 32-bit EISA bus with processor/cache daughter cards—a cache miss on these systems could potentially be a significant drag on performance.

As mentioned, the cache architecture in ALR's system is direct-mapped, making it the only server in the roundup without a set-associative cache subsystem. ALR also made a few other performance compromises to help keep the system price low. For example, the company is using comparatively slower (15 ns) asynchronous SRAMs for cache memory.

ALR's system logic is a proprietary design called Quadflex. It features a 128-bit memory bus with two banks of 256K cache, each with a 64-bit path. Unlike Compaq, ALR implements its design primarily with PLDs and low-density ASICs. The ASICs are basically buffers, comprising two 74F543-equivalents.

I/O Limited

Although these are the first x86-based systems to break the 50-MHz external clock barrier, little has been done to ease I/O bottlenecks. Surprisingly, only two of the systems offer standardized local-bus graphics: NEC's Image P60 features Weitek's new PCI accelerator, and ALR includes VL-Bus graphics from ATI. HP includes Western Digital's 90C30 SVGA controller on the motherboard. DEC, the other Intel Xpress design among the systems, uses an add-in graphics board from Cardinal based on S3's 86C924. There was a vacant, 116-pin outline on the DEC motherboard for WD's 'C30.

The servers largely feature SCSI controllers on the EISA bus, although several vendors said they have plans to move storage control onto the motherboard using the PCI local bus by year's end.

Interestingly, both desktop machines included Analog Devices' 1848 chip for basic audio I/O capability.

Conclusion

The first wave of Pentium-based systems has now been introduced—although the systems clearly are not yet widely available. If the systems we analyzed are any indication, then designs for systems built around Intel's superscalar processor are still evolving from the designs of older 486-based systems.

Computer manufacturers likely will be evolving their system-logic implementations for Pentium over the course of the year; no doubt they will continue to focus on the memory subsystem for improvement.

Industry-standard local buses like PCI surely will help contribute to overall system performance by easing bus traffic to peripherals. The servers we analyzed will benefit from local bus—today their disk arrays and network controllers are effectively I/O-limited because they connect to the EISA expansion bus. While it's faster and wider than ISA, EISA's capabilities pale in the face of Pentium's 64-bit data path and base frequency of 60 MHz.

With the availability of a functional PCI chip set from Intel scheduled for July, Pentium systems featuring that bus should be commonplace by Fall Comdex. As a result, Pentium systems as a class should be boasting higher performance by November.♦

