AMD, Cyrix Lag on FP, MMX Performance

Tests Show Strength on Business Applications, Weaknesses Elsewhere

by Michael Slater

An extensive series of tests recently performed by PC Magazine Labs provides the most detailed view yet of the complex performance characteristics of Intel's processors and the alternatives from AMD and Cyrix. The test results confirm our initial expectations that the alternative chips fall significantly short of Pentium II, and in some cases even behind Pentium/MMX, on code that makes extensive use of floating-point or MMX functions (see MPR 6/2/97, p. 32). (Note: Ziff-Davis, which publishes *PC Magazine* and created the benchmarks used in these tests, also owns MDR, which publishes *Microprocessor Report*.)

As Figure 1 shows, the performance of the chips varies significantly from one application and operating system to another. In this figure, the results for all benchmarks have been normalized so that Pentium/MMX-233 scores 1.0, and all the processors are running at 233 MHz (except for Cyrix's 6x86MX, which is rated PR233 but runs at 187.5 MHz). All the systems had the same type and size of main memory, the same graphics system, and the same disk drive (see sidebar "Systems and Benchmarks," page 3).

Because these benchmark results are based on a small set of motherboard designs, specific Photoshop images and tasks, and particular scripts for all the applications, the results are not entirely conclusive. But it is apparent that the relatively weak MMX and FP performance of the Cyrix and AMD chips is a significant factor for some applications.

To highlight the differences among the processors, the suite included several tests designed to stress MMX and floating-point capability—and they stress these functions far more than most applications today. The results nevertheless provide an early warning that the weaknesses of these processors could become more important as 3D games become more advanced and applications such as photo editing become more pervasive.

The saving grace for AMD and Cyrix is that few of today's applications make extensive use of FP and MMX functions. On most of today's applications, including many 3D games, the AMD and Cyrix parts perform well. By the time applications that are more demanding of FP and MMX performance are widespread, both companies are likely to have enhanced versions of their chips that address these weaknesses. (The first such parts will be disclosed at October's Microprocessor Forum.) Buyers seeking a long life for their system should carefully consider the weaknesses of the current chips, however.

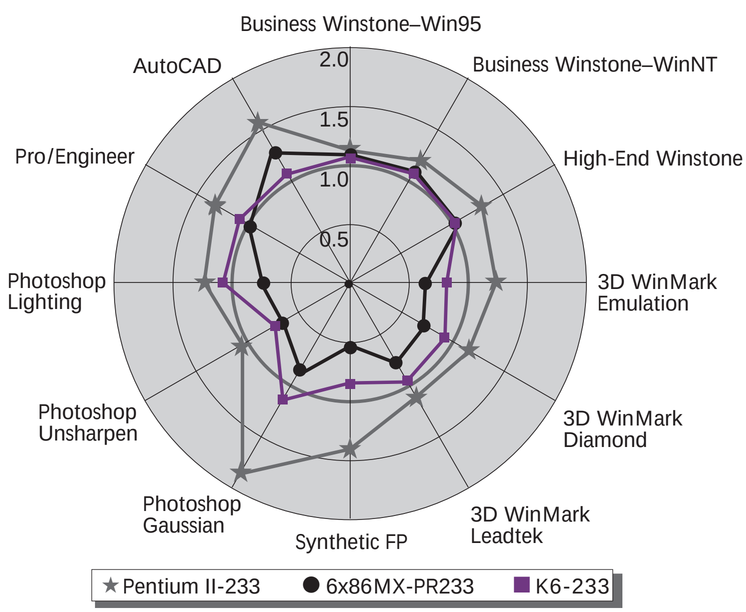

Figure 1. The AMD and Cyrix processors performed well on Winstone but fell short on Photoshop and 3D WinBench. Ratings are scaled so that Pentium/MMX-233 scores 1.0. Other than Winstone, these tests stress FP and MMX performance to give insight into future application performance. (Data source: PC Magazine Labs)

Performance Derives From Microarchitecture

The sources for the performance variations are not hard to find in the microarchitectures. Both Pentium/MMX (see MPR 10/28/96, p. 20) and Pentium II (see MPR 2/17/97, p. 1) have dual-issue MMX units, while neither the AMD K6 (see MPR 10/28/96, p. 26) nor the Cyrix 6x86MX (see MPR 10/28/96, p. 23) does. Both the AMD and Cyrix parts have shorter latencies for some MMX operations, but all the chips are fully pipelined, so the MMX throughput is the same for a series of multiply or multiply-add operations. The big advantage of the Intel chips is their dual-issue design, which doubles the throughput (except for multiply or multiply-accumulate instructions, which cannot be issued in pairs because there is only one multiplier).

For FP operations, Intel's processors—both Pentium and Pentium II—are fully pipelined, while the K6 and 6x86MX are not. AMD's and Cyrix's chips therefore deliver only half the peak throughput of Intel's CPUs on FP operations. AMD's design has shorter latencies than Intel's, but throughput is more important in most applications. Cyrix's chip has longer latencies than Intel's.

The 6x86MX is further handicapped by its clock speed. Although it is impressive that the chip achieves its PR233 rating at only 187.5 MHz, for single-cycle opera-

	ZD Winstone 97			ZD 3D WinBench 97			Synthetic FP	Photoshop			Pro/ Engineer	Auto-CAD
	Business		High End	Emu-lation	Dia-mond	Lead-tek		Gaussian blur	Unsharp mask	Lighting effects		
System memory	32M	64M	64M	32M	32M	32M	32M	128M	128M	128M	128M	128M
Intel Pentium-200	39	48	19	10	20	55	734	21.3	20.4	118	638	640
Intel Pentium/MMX-200	44	54	22	23	24	58	759	18.5	5.4	104	443	576
Intel Pentium/MMX-233	45	56	23	26	26	62	872	15.0	4.9	91	411	486
AMD K6-200	47	59	23	19	22	53	616	15.1	7.3	100	403	477
AMD K6-233	49	62	24	21	24	58	708	13.6	6.8	86	377	446
Cyrix 6x86MX-PR233	50	63	24	16	19	45	425	17.6	7.6	121	409	371
Intel Pentium Pro-200	42	63	28	11	22	63	960	11.6	12.7	84	349	364
Intel Pentium II-233	51	67	30	32	30	70	1,217	8.1	4.6	73	312	311
Intel Pentium II-266	53	70	32	35	31	74	1,329	7.5	4.1	63	291	284
Intel Pentium II-300	54	71	33	38	32	76	1,552	6.9	3.8	57	279	270
Units	Rating: larger is better							Seconds: smaller is better				

Table 1. The two Business Winstone configurations show the benefit of 64M of DRAM. Note that the 64M configuration was run under Windows NT, which further boosts the Pentium Pro and Pentium II scores. Numbers in purple are for Windows 95; all others are for Windows NT. (Source: PC Magazine Labs)

tions its peak performance is cut by 20%, all else being equal. On the other hand, this chip benefits from its 75-MHz bus. The K6, Pentium/MMX, and Pentium II all have 66-MHz system buses. But unlike the other chips, Pentium II does not depend on its system bus for access to its level-two cache. In a 266-MHz Pentium II, the level-two cache interface runs at 133 MHz. The increased L2 cache bandwidth is one factor that boosts the performance of this processor.

There are countless other subtle differences among the processors, from variations in instruction timing to differences in the number of instructions that can be issued at the same time and restrictions on instruction grouping. Sorting out just how valuable each feature is can't be done with any accuracy without extensive simulation data, making it hard for outsiders to evaluate the microarchitectures. In the end, delivered performance is what matters.

The challenge in evaluating processors on delivered performance shifts from understanding the effect of each microarchitectural feature to choosing relevant benchmarks, applications, and tasks, in addition to using hardware that holds the peripherals and memory systems constant. To show the variations in performance on business applications, as well as on applications that make extensive use of floating-point and MMX functions, PC Magazine Labs used Ziff-Davis's Winstone 97, 3D WinBench 97, a synthetic floating-point benchmark, several Photoshop tasks, and Pro/Engineer and AutoCAD driven by scripts that perform a series of functions. Table 1 shows the actual scores.

AMD, Cyrix Fare Best on Business Winstone

The AMD and Cyrix processors fare best on the Business Winstone 97 benchmark, which consists of eight 32-bit business applications. These applications use few floating-point instructions, and they contain no MMX code. The applications are driven by scripts that exercise the applications in ways believed to be typical of actual use.

On Business Winstone under either Windows 95 or Windows NT, the AMD and Cyrix processors delivered very similar scores. The 6x86MX was about 2% faster than the K6; both chips were 9–13% faster than Pentium/MMX at the same clock speed (or PR rating, in Cyrix's case), making them attractive alternatives to that chip for users of typical business applications. Compared with Pentium II, they came close under Windows 95—only 4% short for the K6 and 2% short for the 6x86MX—but both lagged about 6% behind under Windows NT.

The High-End Winstone 97 benchmark is designed to represent the workload of more technical users. This suite includes a C++ compiler, two photo-editing programs, a CAD package, and two data analysis/visualization programs. It has some floating-point code, but no MMX software. On this benchmark suite, run only under Windows NT, the AMD and Cyrix processors remained very close. Both parts delivered 4% better performance than Pentium/MMX but 20% less than Pentium II.

The Winstone suites tend to compress the differences among processors—floating-point and MMX issues aside. The scripts perform a long series of functions, including all the tasks a user would typically perform. Many of these tasks involve updating the display, and some involve accessing the disk; these functions tend to be limited by the I/O system, not the processor. The small variations in performance from one processor to another on these I/O tasks dilute the larger differences on processor-intensive tasks when they are all combined to produce the final score. These results therefore understate the performance differences users will experience on processor-intensive tasks.

CAD and Photoshop Show Pentium II's Strength

To address this limitation, as well as to provide workstation-class applications that use more floating-point and MMX functions, PC Magazine Labs developed scripts to exercise two CAD programs—Pro/Engineer and AutoCAD—as well as Photoshop. The CAD tests involve a mix of tasks, so they

Systems and Benchmarks

All the systems tested had 60-ns EDO memory (size varied by test, as shown in Table 1), a 5G IDE disk drive from IBM, and a 512K L2 cache (except Pentium Pro, which had 256K). The display card for all measurements except the 3D WinBench/Leadtek tests was a Diamond Stealth 3D 2000 with 4M DRAM and an S3 Virge controller. The Leadtek test used a Leadtek WinFast 3D L2200 board with 4M of SGRAM and a 3Dlabs Permedia NT controller. Display resolution was 1024 x 768 for all tests except 3D WinBench, with 16.7 million colors for High-End Winstone and 65,536 colors for all other tests. The 3D WinBench tests were run with a resolution of 640 x 480 and 65,536 colors.

The Pentium and Pentium/MMX systems used a Tyan Titan 1572 motherboard with an Intel 430TX chip set. The Pentium II system was an HP Vectra VL 6/266 with an Intel 440FX chip set; the Pentium Pro system was an HP Vectra VA 6/200 with an Intel 440FX chip set. The K6 system used a Polywell 500TX motherboard with an Intel 430TX chip set, and the 6x86MX system used an M Technology Mustang R-534F motherboard with an SiS 5571 chip set.

For more information on Winstone 97 and 3D WinBench, see *www.zdbop.com*. These benchmarks are available for a nominal charge.

The synthetic FP test, developed by Ed Henning of *PC Magazine UK*, consists of a weighted average of five algorithms: Poisson's equation, FFT, planetary orbitals, complex polygon area, and linear equations.

The Photoshop tests used a 15M image. The Gaussian blur used a radius of 10, and the unsharpen mask used a radius of 1.0.

For additional details on the Pro/Engineer, AutoCAD, and Photoshop tests run by PC Magazine Labs, as well as additional test results, see the September 23 issue of the magazine or *www.mdronline.com/pc_processors*.

exhibit some of the same compression due to I/O tasks as Winstone. The Photoshop tasks are processor-intensive, however, and involve minimal I/O. These tests were run with a typical workstation configuration, using Windows NT and 128M of DRAM.

The AMD and Cyrix processors outperformed Pentium/MMX—but fell well short of Pentium II—on both the AutoCAD and Pro/Engineer suites. Interestingly, Cyrix's processor was especially strong on AutoCAD: 31% faster than Pentium/MMX and 20% faster than the K6 (but still only 84% of Pentium II's performance). On Pro/Engineer, on the other hand, the K6 was 8% faster than the 6x86MX.

On the Photoshop tests, both the AMD and Cyrix parts dropped far below the performance of Pentium II, with Cyrix suffering the worst. These tests use both MMX and floating-point instructions. The "lighting" test is the most floating-point intensive of the Photoshop tests. The Gaussian blur benefits greatly from MMX when used with a small radius value, but not with larger ones; PC Magazine Labs used a radius of 10, eliminating most of the MMX benefit and making this a test of basic integer performance. The stellar performance of Pentium II on this test may have been due to the chip's large L2 cache bandwidth.

The K6 was 12–41% faster than the 6x86MX on the Photoshop tests—but it delivered only 60–85% of the performance of Pentium II. The K6 fell to 72% of Pentium/MMX performance on the MMX-intensive "unsharpen" filter but outperformed that chip by 6–10% on the other tests.

The 6x86MX delivered a lackluster 64–85% of Pentium/MMX performance and only 46–61% of Pentium II performance on the Photoshop tests.

The synthetic FP test implements a variety of FP algorithms (see sidebar). These results correlated reasonably well with the Photoshop lighting test, the most FP-intensive of the application tests.

Cyrix, AMD Weak on 3D WinBench

The ZD 3D WinBench tests showed big differences among the chips. As with the Photoshop tests, the K6 did significantly better than the 6x86MX, but even the K6 was well below the performance of Pentium/MMX and far below that of Pentium II.

The 3D WinBench suite renders a variety of images. Most of these images are considerably more complex than those in today's 3D games. Unlike Winstone, which is based on today's most widespread applications, 3D WinBench is forward looking: its goal is to evaluate performance on 3D games that will ship during the coming year. This approach was dictated by the rapid advances in 3D accelerators and the corresponding evolution of 3D games; while next year's business applications probably won't be too different from last year's, next year's 3D games will be quite different from today's.

Because 3D WinBench uses relatively complex scenes, it stresses the processor's floating-point unit more than most of today's games—and therefore shows bigger differences among the processors. The benchmark also uses Direct3D's lighting and geometry functions, whereas most games today perform these functions internally and avoid the higher-quality, FP-intensive lighting algorithms. Cyrix and AMD are quick to point out that their processors perform well on most of today's games.

The 3D WinBench test was designed to evaluate graphics cards, not microprocessors. But since the geometry portion of the graphics pipeline exercises the processor's FPU, it provides a good test of FP performance using algorithms that will become increasingly important.

If no 3D accelerator is present, all rendering functions are performed in software. This emulation software, which is part of Microsoft's Direct3D package, makes extensive use of MMX instructions. Thus, to provide a test that emphasizes MMX performance, the software was first configured to force all rendering functions to be emulated. Because of the single-issue MMX units in the Cyrix and AMD processors, testing in this mode showed the biggest differences. The K6 was 32% faster than the 6x86 but still delivered only 82% of the performance of Pentium/MMX and 67% of Pentium II's.

The emulated 3D benchmark provides an interesting illustration of the variations in MMX performance. Realistically, though, anyone who cares about 3D performance will have a 3D accelerator, making the MMX rendering code irrelevant. With a low-end accelerator (S3's Virge chip), the performance gap was cut in half, but the K6 still delivered only 91% of Pentium/MMX performance and 80% of Pentium II's. With a faster accelerator (3Dlabs' Permedia NT chip set), the gap between the processors shrank slightly but didn't disappear. In some ways, the faster rendering engine stresses the processor more, since more FP operations must be performed to keep up with it.

Whether 3D WinBench turns out to be an accurate predictor of performance on future games depends on just how the software and hardware evolves. Many companies are designing future 3D accelerators to include geometry engines, and if this function is moved to the accelerator, then the processor's FP capability will become irrelevant. Taking advantage of such accelerators, however, will require changes in the way the games are written; most games don't use Direct3D to perform geometry and lighting functions, handling these tasks internally. Changes in Direct3D itself are also required to enable geometry acceleration.

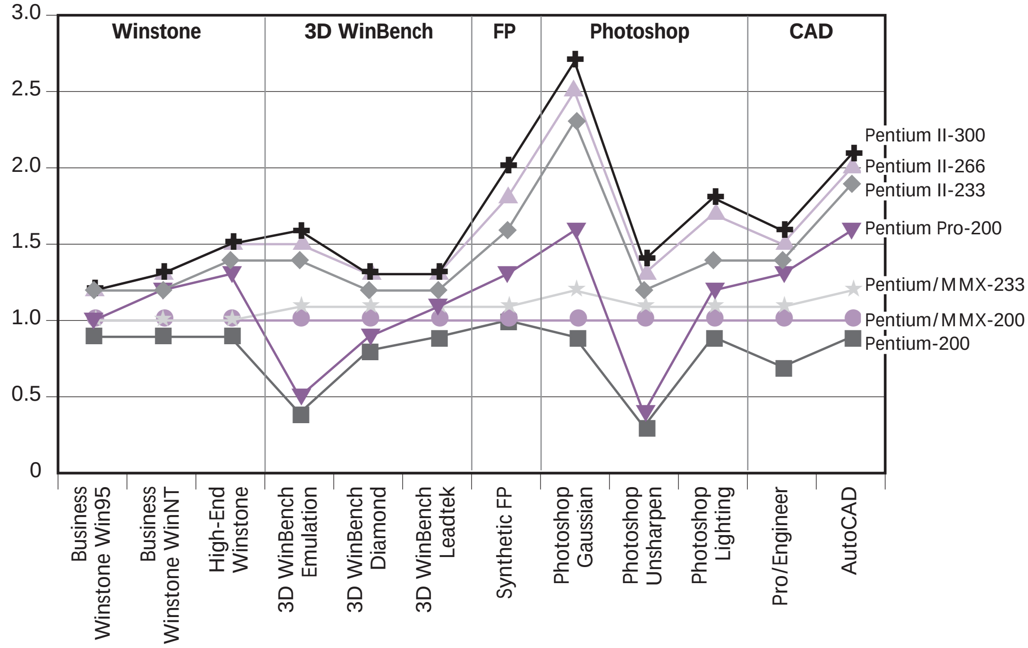

Figure 2. Pentium II provides a modest benefit over Pentium/MMX on Business Winstone but delivers a big boost on some other tests. Ratings are scaled so Pentium/MMX-200 scores 1.0. (Data source: PC Magazine Labs)

Comparing Within Intel's Lineup

In addition to illuminating the strengths and weaknesses of the AMD and Cyrix processors, the benchmark results also demonstrate the benefits of MMX and highlight where Pentium II's advantage over Pentium/MMX is most significant.

Figure 2 shows the relative results for the Intel processors, scaled so the 200-MHz Pentium/MMX scores 1.0 on all tests. At the same clock speed, Pentium/MMX delivered 11–16% better performance than the "classic" Pentium on most of the non-MMX tests, showing the benefit of the larger cache and pipeline tweaks. Curiously, Pro/Engineer—which does not use MMX—got a 44% boost from Pentium/MMX. This test presumably fits much better into Pentium/MMX's 32K cache than Pentium's 16K cache.

These results make it clear which tests are the most MMX-intensive: 3D WinBench in emulation mode and the Photoshop unsharpen filter jump out as relative low points for the non-MMX Pentium and Pentium Pro. Note that Pentium Pro was significantly faster than Pentium/MMX on Gaussian blur, indicating that MMX doesn't give much of a boost on this operation.

The advantage of Pentium II over Pentium/MMX varies greatly from test to test. Running Business Winstone under Windows 95, the Pentium II benefit is only 13% (with both processors at 233 MHz). With the same benchmark under Windows NT, the Pentium II advantage grows to 20%—and for High-End Winstone, it jumps to 30%. Of course, Pentium II is available at higher clock speeds than Pentium/MMX (or the alternative processors), increasing its lead. The fact that the Pentium II advantage over Pentium/MMX is noticeably greater under Windows NT than under Windows 95 indicates that Pentium II, despite some improvements over the Pentium Pro design, is dragged down by some of the 16-bit code remaining in Windows 95.

The Pro/Engineer and AutoCAD benchmarks show an even more significant gain of 32–56% for Pentium II over Pentium/MMX—more than enough to justify the higher system price for users of applications such as these. These programs benefit from the greater cache bandwidth of Pentium II, as well as from the higher core performance.

The Photoshop results illustrate Pentium II's advantage in both integer and floating-point performance, as well as the faster L2 cache. The Gaussian blur filter shows Pentium II in its best light—an impressive 85% better performance than Pentium/MMX at the same clock speed. Since this task is neither MMX- nor FP-intensive, it is apparently the cache bandwidth that leads to the big boost. On the MMX-intensive unsharpen filter, on the other hand, the Pentium II benefit was a mere 7%.

The 3D WinBench results show a modest speedup for Pentium II vs. Pentium/MMX. When using software for all rendering, Pentium II is 23% faster, but the benefit drops to 13–14% with a 3D accelerator. The low-cost Virge-based 3D card actually reduced Pentium II system performance, compared with software emulation. With a Pentium II processor, the CPU apparently is able to render more quickly than the low-cost accelerator card. The Permedia-based card more than doubled the 3D rating of the Pentium II systems, compared with the same system using software emulation or the Virge-based card. (Note, however, that the 3D WinBench score also assigns extra points for implementation of additional quality features, so a score that is twice as high does not necessarily mean images are rendered twice as quickly.)

Are Higher Speeds Worth It?

The benchmark results also give insight into the benefit of higher clock rates. The 233-MHz Pentium/MMX shows how small this boost can be: only 2% over the 200-MHz chip for Business Winstone under Windows 95. Even High-End Winstone delivers only a 5% boost. At 233 MHz, the processor is beginning to saturate the bus, and the I/O system limits the performance increase. The Photoshop and CAD tests showed larger increases, due to their more CPU-intensive nature; tasks that fit well within the on-chip cache get a much bigger boost.

For Pentium II, even though it has much greater L2 cache bandwidth—which, unlike Pentium/MMX's L2 cache interface, scales with the CPU speed—the situation is only slightly better. The 266-MHz Pentium II delivers only 4% better performance than the 233-MHz version on Business Winstone under Windows 95. The boost increased to 7% on High-End Winstone. In both cases, the rating gain was apparently dragged down by the relatively slow I/O systems.

This does not mean users won't see larger performance gains, however, even on the same applications that are in the Winstone suite; when performing a CPU-intensive function, such as a spreadsheet recalculation, the gain will be greater. The Pro/Engineer, AutoCAD, and Photoshop tests all showed performance gains of 7–16%. On 3D WinBench, performance went up 11% with software rendering but only 4–5% when a 3D accelerator was used.

At the chip level, paying an extra $139—26% more—for the faster chip may seem out of proportion for a 5–10% performance gain. A premium of $150 to $200 for the same performance gain on a $2,500 system, however, is not such a bad deal.

The 300-MHz part, however, is another story. Even after the recent price cuts, this chip sells for $851—$182, or 27%, more than the 266-MHz version. Yet the performance improvement was negligible: 1–3% on Winstone and on 3D WinBench with a good 3D card. It did provide more improvement, 4–11%, on the CAD and Photoshop tasks, so it might be worthwhile for some users. The big gap between the gain on Winstone and the gain in Photoshop suggests that the drag of the I/O operations in Winstone is becoming overwhelming at higher clock speeds.

Performance Picture More Complex All the Time

These benchmark results demonstrate how complex the performance picture has become for PC processors. The AMD and Cyrix processors are positioned by their makers solely on the basis of Business Winstone—which represents the most common PC workload, and for which the processor designs were tuned. For the corporate market, this is a reasonable choice; business applications (outside of the CAD and authoring realms) are moving very slowly toward FP- or MMX-intensive tasks.

For the consumer market, however, which is where most of the AMD and Cyrix processors are being sold, the Business Winstone rating overstates the performance users will see on advanced 3D games or photo editing. Software-based modems and video compression and decompression programs, which often are MMX-intensive, will consume more processor bandwidth on AMD and Cyrix systems.

Cyrix's 6x86MX delivers slightly better results than AMD's K6 on the Winstone suites—but it falls far behind on the Photoshop and 3D WinBench tests. Both processors fall short of Pentium/MMX, and far short of Pentium II, on some of these tests. The differences are irrelevant for most of today's applications, but users concerned with getting the longest life from their processor, and who want to be able to run the hottest new 3D games, should think twice about straying from the Intel lineup—at least until AMD and Cyrix release their next-generation chips.

On the other hand, the cost savings that result from selecting an AMD or Cyrix processor can enable users to afford more robust system configurations. Buyers limited to a particular price point might have to settle for 32M of DRAM with an Intel processor but could afford 64M with an AMD or Cyrix processor; the additional memory can make up the performance difference in the processors. For 3D graphics, a higher-performance card can make a bigger difference than a faster processor. And when Cyrix and AMD release enhanced versions of their processors, CPU upgrades may be possible.

For users of mainstream business applications, the Cyrix and AMD parts are comparable: both are very good values, delivering better-than-Pentium/MMX performance at low prices. For users concerned about advanced 3D game or photo-editing speed—and who are not especially price-sensitive—Pentium II is the clear choice. M