CPUs for Workstations and Small Business Servers

By Johan De Gelas – April 2002

Workstation applications such as 3D animation, 3D rendering, finite element analysis, photo editing, and so forth are known to be very CPU intensive. These are exactly the applications that many small businesses run. Those small business do not have loads of cash and are faced with the same problems as you and me: what system offers the best performance for a decent price? In this article we investigate the following several CPU intensive applications on four different CPU configurations:

- Athlon XP 2100+
- Dual-Athlon MP 2000+
- 2 GHz Northwood Pentium 4
- 2.4 GHz Northwood Pentium 4

The applications and benchmarks we test with are:

Application	Benchmark
2D Photo Editing	Photoshop 6.01
3D Rendering	3DSMax 4.26 & Maya 4.0
Finite Element Analysis	Phase² 5.0
Realtime Software 3D Rendering	Kribi
Video Encoding	Windows Media Encoder 7.1
Database Servers (OLTP)	ZD Serverbench 4.1

The AMD-760MPX motherboards are quite a bit less expensive than the first AMD-760MP motherboard: (Tyan's Thunder K7). Add together two Athlon MP 2000+ chips and you are in about the same price-range as a Pentium 4 2.4 GHz system. So, if you run a very CPU intensive application on the Windows platform, these are the fastest CPU configurations you can buy at a relatively reasonable price. We were not able to include any SMT-capable Xeons, but for now, these are quite a bit more expensive than the tested configurations (around $700). Two 2.2 GHz Xeons and a dual-Xeon motherboard can cost over $2000, while our dual-Athlon system and 2.4 GHz Pentium 4 system both cost around $750. At a later date when Hypertreading has matured and the price of the Xeons is lower, we will include them as well.

The single Athlon 2100+ and 2 GHz Pentium 4 system are interesting in the case that our application does not benefit enough from extra CPU power. In that case it might be interesting to lower the amount the money spent on the CPU and invest in more memory, a better disk subsystem, etc.

This review has taken weeks of benchmarking and testing, and we are happy to finally be able to share our experiences with you. We did not include CAD benchmarks, as many are very video card intensive. We will focus on OpenGL video performance on a later date, but you can check out our previous OpenGL card review for more information on the subject. Our previous workstation reviews can be found here and here (Compaq, Dell, Fujitsu-Siemens, Tyan).

Configuration Details

All systems were tested with NVIDIA's Detonator 28.32 drivers. The Wildcat 5000 driver was Version 05.04.02.20 (Athlon and P4 Optimized). The desktop was set at a resolution of 1024x768x32bpp (85 Hz refresh rate). Our testbed still runs Windows 2000 SP2, as most professionals and hardcore hardware enthusiasts prefer the matured Windows 2000 SP2.

Dual Athlon MP 2000+

- ASUS A7M266-D Bios version 1.005
- MSI K7 Master-L Bios version 1.1
- Tyan Tiger MPX Bios version 1.01
- Gigabyte GA-7DPXDW Bios version 2/21/2002
- 512 MB Samsung ECC buffered PC2100 DDR SDRAM CAS2

Duron 1300, Athlon 1400, XP 2000+, XP 2100+

- ASUS A7V266-E (VIA KT266A) bios version 1.006
- 512 MB Corsair PC2700 XMS (DDR-RAM) CAS 2 (2x256 MB)

2 GHz Pentium 4 (Willamette), 2A GHz Pentium 4 (Northwood) i845 DDR, and 2.4 GHz Pentium 4 (Northwood)

- ASUS P4T-E (i850 RDRAM chipset) bios version 1.005
- 512 MB Corsair RDRAM-45 (4x128 MB)

Common Hardware

- IBM Deskstar DPTA-373420 ATA/IDE 34.2 GB (7200 RPM, ATA-100)
- ASUS Geforce 3 Ti500 64 MB
- AT 2700 10/100 NIC
- Sound Blaster Live!
- Power Supply: ANTEC PP-412X 400W

Software

- Via 4 in 1 Drivers 4.37A
- Intel chipset inf update 3.10.1011
- AMD AGP Driver 5.53, AMD PM Drivers, AMD IDE Drivers (Ver 1.43)
- Windows 2000 Service Pack 2
- DirectX 8.1

We'd like to thank the following helpful people for their support and crucial contributions to this review:

- Damon muzny (AMD) made sure we were able to test the Athlon XP 2100+ and dual-Athlon MP 2000+.
- We thank for Nicole Chia (Gigabyte) for the speedy delivery of the Gigabyte GA-7DPXDW
- Thanks to Augustine Chen (ASUS), Carol Chang (ASUS) and Sharon Tan (BAS computers Netherlands) for the ASUS motherboards.
- Mat Vernon of Tyan sent us the Tyan Tiger MPX.
- Angelique Berden of MSI was provided us with the MSI K7-Master-L.
- Jurgen Eijmberts (Intel) and Marieke Leenhouts (MCS) made sure we could test the 2A GHz and 2.4 GHz Pentium 4 CPUs.
- Robert Pearce of Corsair, provided us with Corsair's PC2700 XMS and Corsair RDRAM

The Pentium 4 Solution

As the i850 is faster than the i845, and more mature than the SiS645DX, we went with the excellent ASUS P4T-E board. This board comes with an AGP Pro slot, making it more suitable for a workstation. We have discussed the Pentium 4 and the Pentium 4 boards in great detail in previous articles. For a detailed discussion of the Pentium 4 Northwood, read our review. For comparison of the different Pentium 4 boards, see our Pentium 4 Platform Shootout.

The Dual Athlon Solution

The three market leaders (MSI, ASUS, Gigabyte) have joined Tyan and created a motherboard around the AMD760 MPX chipset. This may indicate how strong industry support for the Athlon is, as the Athlon MP has only a very small part of the workstation market.

AMD's previous chipset, the AMD-760MP could only be found on the Tyan Thunder K7 and the Tyan Tiger K7. You can find more info about those boards here. The main difference between the 760MP and the 760MPX is the fact that the new southbridge (768 instead of 766) of the MPX chipset can run at 66 MHz. The AMD-762 northbridge was already capable of that, and the result is that the new AMD760MPX chipset has a 533 MB/s (64 bit PCI at 66 MHz) connection between the southbridge (768) and the northbridge (762). In contrast, the AMD-760MP northbridge was connected to the AMD-766 Southbridge via the PCI bus (133 MB/s), which was a bit bandwidth limited with today's ATA100 hard disks (up to 20-50 MB/s), USB 2.0 PCI cards (up to 60 MB/s) and 100 MBit ethernet (10 MB/s).

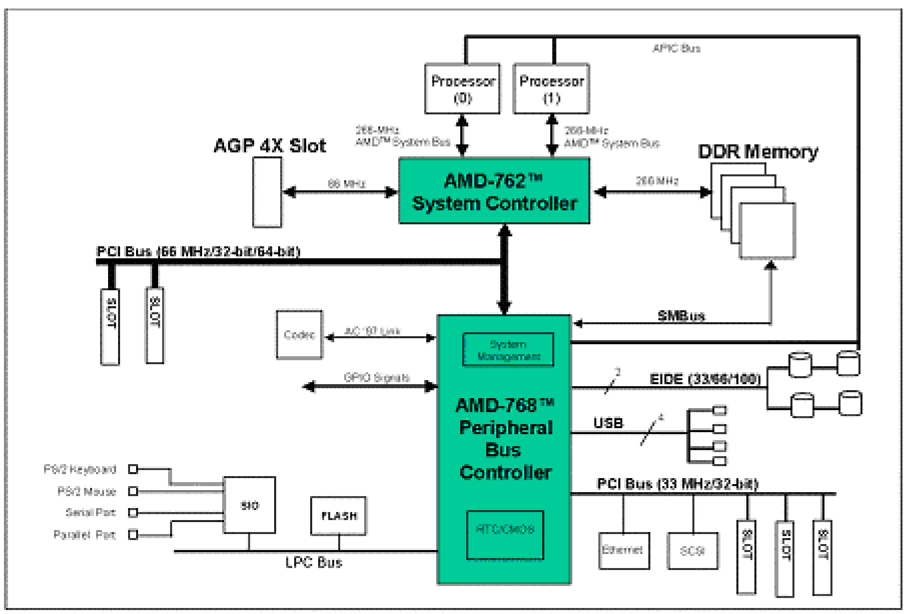

The main improvement is, of course, the fact that the AMD-760MPX chipset can support two 64 bit/66 MHz PCI cards. High-end SCSI and Fibre-Channel RAID cards, as well as Gigabit Ethernet cards are typical components in high-end servers. Such devices generally support 64-bit/66 MHz PCI.

However, 66 MHz PCI cards can run at 5V and 64-bit/66 MHz PCI slots are not compatible with standard 32-bit/33 MHz PCI cards. So for most users, those 66 MHz PCI slots are lost. If you need more than 4 PCI slots, your best option is still the Tyan Tiger MP, because it comes with two "normal" 32-bit/33 MHz PCI slots and four 64-bit/33 MHz slots which can also be used for 33 MHz PCI cards. See below:

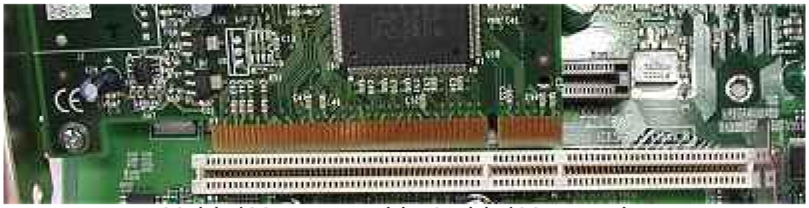

32-bit/33 MHz card in 64-bit/33 MHz slot

The 66 MHz PCI slots can not accept 33 MHz 5V PCI cards and have a different layout. See the picture below. Please note that the back of the motherboard on the picture below is on right side (left side in the previous picture).

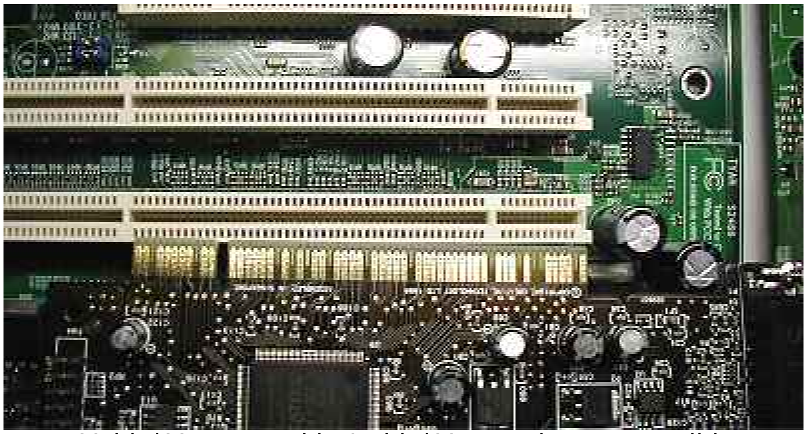

32-bit/33 MHz card in 64-bit/33 MHz slot: not possible

As you can see, it is not possible to insert a 33 MHz PCI card in the 64-bit/66 MHz slots of the MPX boards. Also, be aware that some (RAID cards for example) are universal cards, they can work at 5 and 3.3V. As a result they have an extra notch and can be inserted in the 3.3V 64-bit slots.

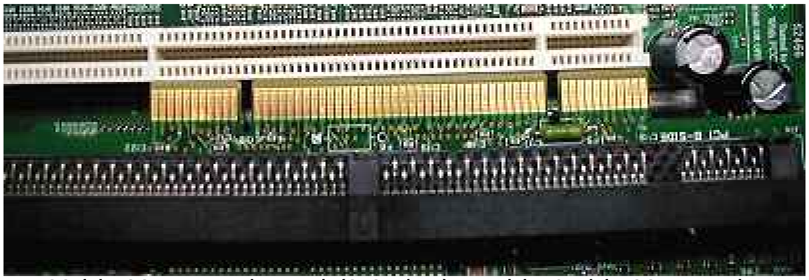

32 bit 33 MHz universal (3.3V/5V) card in 64 bit 33 MHz slot

However, you must verify whether or not these universal boards can work at 66 MHz. We tried it out with a 32-bit/33 MHz universal SCSI RAID controller, and our card didn't fry, but it didn't work either.

It should be noted that the 768 chipset has a minor issue: the USB interface might malfunction in some rare cases. AMD has produced a revised southbridge, and you may expect that the motherboards manufacturers will start using this southbridge on their motherboards in the coming weeks. Nevertheless, all boards come with a separate USB card. Let us take a look at the different AMD760 MPX solutions.

The Motherboards

ASUS A7M266-D

The ASUS A7M266-D features some impressive power circuitry (dual tri-phase). Note that it needs an ATX 12V power supply to function properly.

A huge passively-cooled heatsink is put on top of the northbridge. The heatsink does not get hot, and we like the fact that it doesn't need a fan which can break down. Also thumbs up for the LED that indicates whether the board is getting the 5V "power good" signal of the ATX power supply.

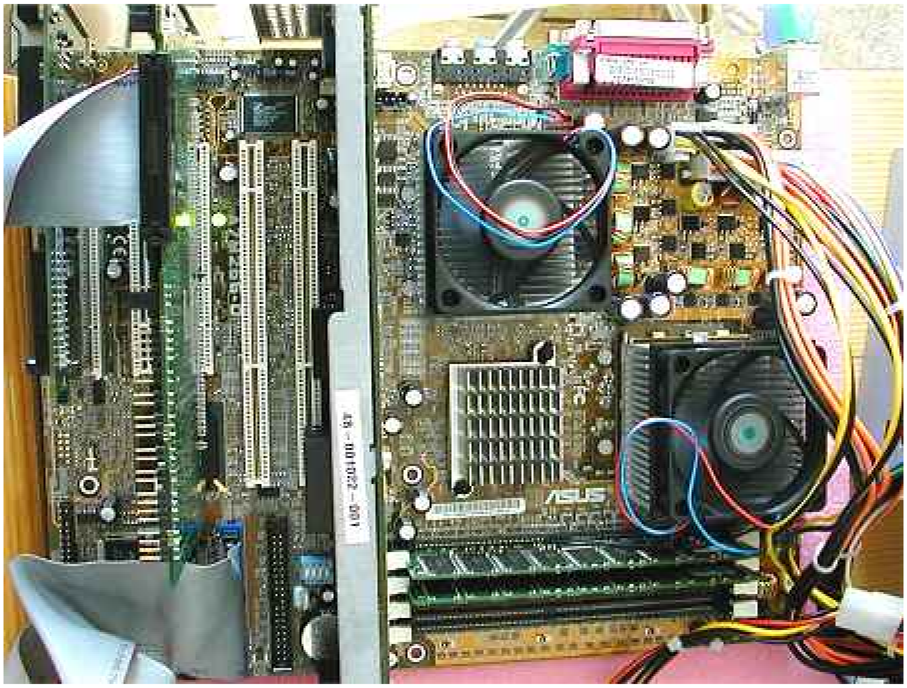

The ASUS board has provided by far the best overclocking possibilities, allowing you to specify the voltage of both the DDR SDRAM and the CPU. The CPUs can be overclocked by steps of 1 MHz and the ASUS board was able to overclock further than any of the other solutions.

Our experience with the ASUS A7M266-D was good, but not as excellent as we are accustomed to seeing from ASUS' products. First of all, we are really missing an integrated NIC. With only three 32-bit PCI slots, onboard LAN is a necessity, particularly considering that any workstation or server is going to be connected to a LAN. To ASUS' defense, they integrated a relatively good sound chip, which may eliminate the need for a PCI soundcard.

The most important disadvantage of the ASUS board was that our Wildcat 5000 - a AGP PRO 50 videocard - could not run totally stable. The SpecViewperf benchmark crashed rarely, but 3DSMax 4.26 refused to start an OpenGL accelerated viewport regardless of what BIOS settings we tried. In other words, 3DSMax froze at the splash screen. We reported the issue to ASUS, and ASUS has started working with 3DLabs to resolve this bug.

We also tried a Geforce 3 Ti 500 and Geforce 4MX 440 on this board. As Lost Circuits reported, it is best to set AGP aperture to 64 MB-128 MB and the AGP mode to *Uncached Speculative Write Combining (USWC)*. The board worked fine after these changes, but still crashed occasionally in SpecViewperf. The fact that all AMD760MPX boards have trouble running SpecViewPerf with NVIDIA GeForce 3 & 4 based cards made it clear to us that this issue is caused by some weird AMD760 MPX/NVIDIA 28.32 driver problem. This first assumption was confirmed, when we noticed that - not surprisingly - the ASUS P4T-E (i850 chipset) had no problem at all with the 28.32 driver.

In case of the ASUS motherboard, the 28.32 driver issue was less severe. In fact, SpecViewPerf crashed occasionally on the ASUS, while it crashed almost always on MSI's, Gigabyte's and Tyan's motherboards. The AWadvs-04 and DX-06 tests caused problems in particular. Not really a big deal, as AWadvs simulates 3D Rendering and animation packages, and Maya and 3DSMax ran very well on all motherboards. Fortunately, MedMCAD-01 and ProCDRS-03, which simulate CAD use, ran happily on all boards as well.

We were unable to detect any problems with NVIDIA's 21.85 driver. It seems that the trouble starts with the 23.xx driver, which in turn seem to have made many OpenGL benchmarks unstable.

Gigabyte GA-7DPXDW

Gigabyte's dual Athlon board comes with very attractive specs. It has nice power circuitry, similar to the ASUS board, and a huge array of "anti-ripple" capacitors. Gigabyte offers the richest feature set: integrated LAN, AC97 sound, and a Promise ATA-100 RAID controller, which can be set to function as a RAID controller or "normal" ATA-100 controller.

The RAID controller can be turned off if you don't use it, and as such, you can reduce the boot time significantly. Good points so far, but Gigabyte's board contends with the same problems as ASUS. With our Wildcat II 5000, OpenGL based benchmarks tend to choke from time to time. Also, we noticed that the CPU core voltage ranged from 1.7 to 1.74V, while our Athlon MP's 2000+ require 1.75V. Contrary to ASUS and MSI there is no BIOS option to remedy this, so we suspect that the relative low voltage to the CPU may cause some of the trouble we encountered. Of course, we tested with a very early board and BIOS, so it is likely that Gigabyte will be able to correct this with a BIOS upgrade.

MSI K7 Master-L

MSI's K7 Master-L comes with a rather modest two-phase power circuitry, but the MOSFETS present in large numbers are covered with large heatsinks. These heatsinks are large enough: even after intensive benchmarking, they were not hot.

It is quite funny to see an Intel chip on a Dual Athlon board. MSI's board comes with a modern Intel 82559ER LAN Controller, which works very well and is recognized immediately by Windows 2000.

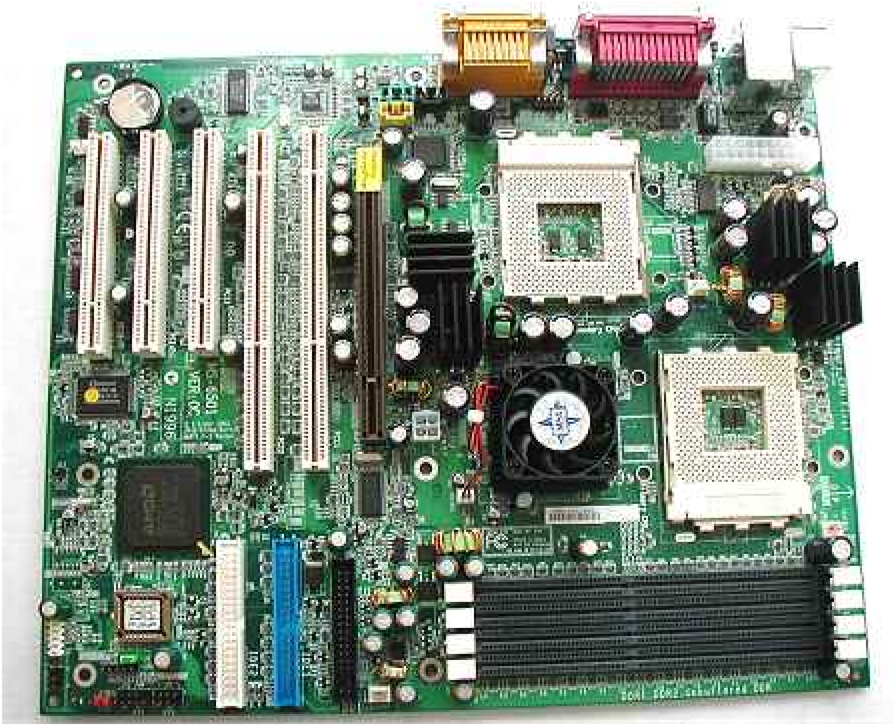

Our first impression of MSI's offering was rather dim, as quite a few benchmarks locked up. Upon further investigation, we found out that the voltage was hovering around 1.7-1.72V. Once we set the voltage to 1.8V, this resulted in a CPU core voltage of 1.76-1.78V and the board was running exceptionally well. It even proved to be the only AMD-760MPX board able to feed our 3DLabs Wildcat II 5000 very well. With this AGP PRO 50 board and the voltage set at 1.8V, we didn't encounter even a single lock-up, even after weeks of testing with the most intensive OpenGL and non-OpenGL applications.

One reason why MSI's board seems to do better is the placement of the 4-pin 12V connector: very close to the AGP Pro extension. As the AGP Pro specifications indicate, the AGP Pro slot delivers extra amperage from the 12V lines. Placing the 12V lines close the AGP Pro extension may have been a very good idea of MSI's engineers (more information about AGP Pro).

Tyan Tiger MPX

Tyan's Tiger MPX is the only board that does not require an ATX power supply. Tyan allows you to attach a normal HD power connector instead of the the typical 4-pin ATX12V connector.

While the other boards require a 400W supply, the Tyan board should work fine with a 300W supply -- as long as the 5V rail can deliver 30A. We didn't try with a 300W supply, but the Tyan board also worked fine with our 350W Sparkle. The most likely reason is that Tyan's board has no AGP Pro slot.

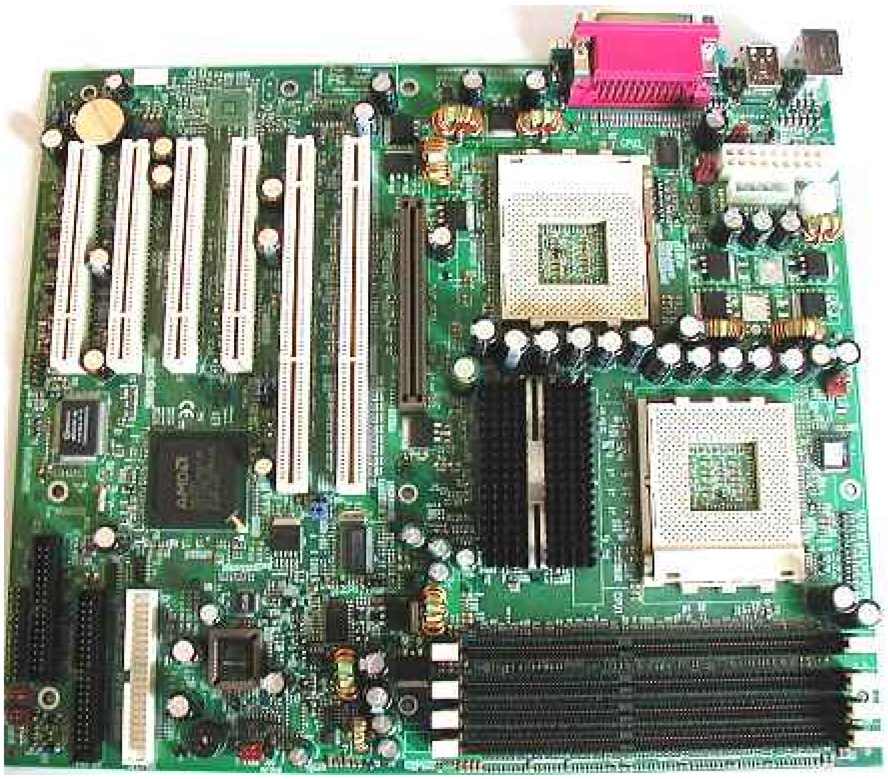

Our early sample came without any onboard LAN, but the version you can buy (S2466N) includes a 3COM NIC chip onboard. Apart from the Nvidia/AMD driver issue, the board ran flawlessly. The relatively low power supply requirements and the 4 PCI slots make Tyan's board one of the easiest to upgrade to. The layout is a bit weird though, with the floppy connector pretty far to the left end of the board.

As mentioned earlier, AMD has produced a new revision of the AMD768 southbridge and AMDZone has reported that you should see the S2466N-4M very soon. The S2466N-4M includes functioning onboard USB 1.1 ports, thanks to the updated AMD-768 southbridge. It might be worth waiting for this version particularly if you need to use USB devices.

If you need AGP Pro, Tyan will point you to the Thunder K7X which also comes with two LAN connectors and optionally a dual SCSI UW160 interface and server management features.

Motherboard Comparison Summary

As a table can present this information to you for comparison much more efficiently, you can find all the features of the different boards below listed below:

Feature	ASUS A7M266D	Gigabyte GA-7DPXDW	MSI K7 Master-L	Tyan Tiger MPX	ASUS P4T-E (Pentium 4)
AGP Pro	Yes	Yes	Yes	No	Yes
32-Bit PCI slots	3	3	3	4	5
64-Bit PCI slots	2	2	2	2	0
On-Board USB	No	Yes	Yes	Disabled (*)	2
USB PCI card	4 Port 2.0	4 Port 2.0	4 Port 2.0	4 Port 1.1	No
On-board LAN	No	Intel 82550	Intel® 82559ER LAN Controller	3Com 3C920 ASIC**	No
On-board Sound	CMEDIA 8738 6-channel	AC '97	AC '97	No	No
Raid controller	No	Promise	No	No	No
Overclocking	Up to 150 MHz	up to 149 MHz	Up to 150 MHz	Up to 150 MHz	Up to 160 MHz
Max. overclock	142 MHz	138 MHz	138 MHz	138 MHz	150 MHz
Increment 1 MHz	Yes	No	No	No	Yes
Voltage Adjustments CPU	Yes	No	Yes	No	Yes
Voltage Adjustments RAM	Yes	No	No	No	No

* Works on the S2466-4M
** Not on our sample board S2466, but available on the S2466N

Furthermore, you can also find a compressed view of our experiences with each board below:

Video Card Stability	ASUS A7M266-D (Dual Athlon)	Gigabyte GA-7DPXDW (Dual Athlon)	MSI K7 Master-L (Dual Athlon)	Tyan Tiger MPX (Dual Athlon)	ASUS P4T-E (Pentium 4)
ATI Rage Pro (3.3V AGP)	Good	Good	Good	Good	Not possible only 1.5V
Geforce 3 Ti 500	Good, but (*)	Good, but (*)	Good, but (*)	Good, but (*)	Good
3D Labs 5000 (AGP Pro)	Mediocre	Mediocre	Good	Not Possible, No AGP Pro	Good
Memory	**ASUS A7M266-D**	**Gigabyte GA-7DPXDW**	**MSI K7 Master-L**	**Tyan Tiger MPX**	**ASUS P4T-E (Pentium 4)**
Unregistered DIMMs?	Yes, 2	No	Yes, 2	Yes, 2	Yes, 4

* Problem in two tests of SpecViewPerf (OpenGL) with Nvidia video driver /AMD AGP driver

To sum it up: The MSI K7 Master-L made a very good impression once we pushed the voltage of the CPU up a bit higher. It was also the only board that worked flawlessly with our Wildcat II 5000 AGP Pro card. It allows a lot of tweaks (Memory, CPU voltage) and it comes with a very good integrated LAN chip.

The Tyan Tiger MPX seems to be the easiest upgrade for those who do not need a AGP Pro card: no heavy power supply requirements, 4 PCI slots and an integrated LAN connector. AC97 sound would have been nice though. Typically workstation and server people don't want to invest in a PCI sound card (buggy drivers, possible IRQ conflicts, another point of failure), but like to have some basic sound (alert sounds of management software, playing MP3s to make the hard work more pleasant, etc.)

Both ASUS and Gigabyte's offering are stable boards apart from the Wildcat II 5000 issue. The NVIDIA driver issue is not really important as we have not seen any instabilities in Maya and 3DSMax, only in SpecViewPerf. As a result, the number of people that will be affected by this Nvidia/AMD driver issue will be relatively small.

Benchmarks

As it is impossible to simulate the way every workstation or server buyer uses his/her software, we'll try to formulate a few basic rules of thumb, which describe the typical behavior of a certain application. To understand this better, take a look at the table below. We measured memory performance with PCMark 2002 in Windows 2002 SP2. The results are similar to our results under DOS (Cachemem and Stream).

Configuration	Read Bandwidth (MB/s)	Write Bandwidth (MB/s)
Dual Athlon - AMD760MPX, buffered DDR266	870	570
Single Athlon - AMD760MP, buffered DDR266	990	590
Single Athlon XP - KT266A, unbuffered DDR266	1290	610
Pentium 4 - i850, PC800 RDRAM	1793	1037

A dual Athlon system offers a lot of processing power but it's memory subsystem is no match for the memory subsystem of the Athlon XP running on a VIA KT266A platform, or the Pentium 4 on the i850 platform. As a result, a dual Athlon system will only be interesting for applications that are very processor intensive, but which do not depend much on the memory bandwidth or latency. The Athlon XP 2100+ in our test has the advantage of an aggressively tuned northbridge, which can deliver up to 48% more bandwidth than a dual CPU AMD-760MPX system.

So, even if an application is nicely multithreaded and demands a lot of CPU power, if it makes too much use of the memory subsystem, a dual Athlon 2000MP+ workstation might fail to perform much better than a higher-clocked single processor Athlon, like the upcoming Athlon XP 2200+. Nothing really new for many of our loyal readers, but we will see further that it is good to keep this in mind when looking at the benchmarks.

3D Animation and Modeling

We'll start with a classic workstation application, 3D animation and modeling. 3D Studio Max is the most popular tool for game developers and probably the most wide spread 3D animation and rendering package in use today. Specifically, we test with the 4.2.6 version, which is optimized for the Pentium 4 and Pentium III, and should provide up to 30% better performance for Pentium 4 and Pentium III configurations. More information here.

It is important to note that the efficiency of a second CPU will vary a lot with the sort of scene you are working on. The second CPU will have the best effect in scenes where you use a lot of lights and raytracing on a medium amount polygons (+/- 100-200k). If you need more polygons, then the importance of the memory subsystem grows.

Our first test is the architecture scene from the SPEC PAC 3DS MAX R3.1 benchmark. This test has a moving camera that shows a complicated building, a virtual tour of a scale model. This incredibly complex scene has no less than 600,000 polygons, 7 lights, with raytracing and fog enabled. Frames 20 to 22 were rendered at 500x300 to the virtual frame buffer (memory).

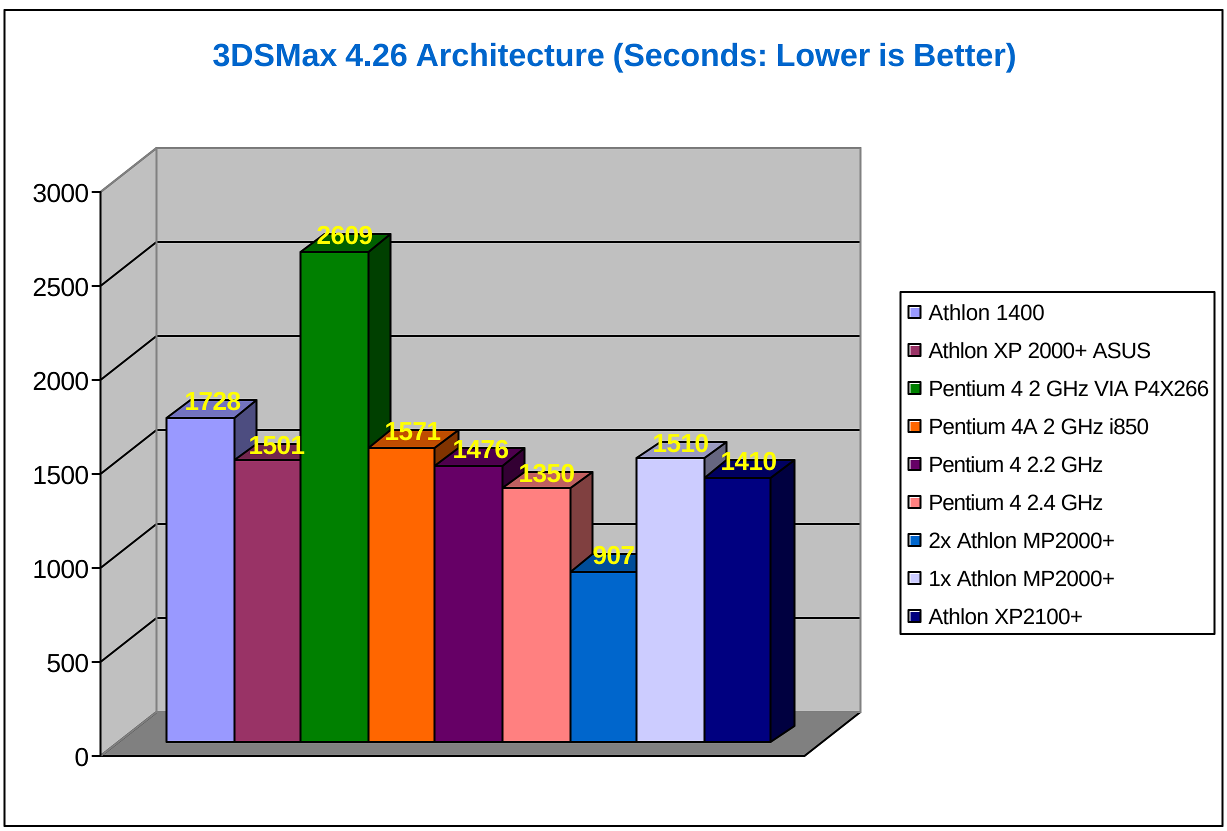

3D Studio Max Continued

The second CPU accelerates the rendering by 66%, which is excellent. Even with huge amounts of polygons, a second CPU pays off much more than a strongly improved memory subsystem (Single Athlon MP 2000+ vs Athlon XP 2000+). One of the reasons for this is the inclusion of 7 lights, and the use of raytracing in the Architecture scene.

The Ape animation, a typical game character rendering makes heavy use of light: no less than 44 different light sources. The scene also features complicated inverse kinematics: bone manipulation to control the facial animation and parameter wiring to move the fingers. Maxscript (macro language) is used to control certain movements. The polygon count is relatively low, only 26,000 polygons. Motion blur, which is one of the functions optimized for SSE-2, is achieved by rendering the scene in six passes. We rendered frames 20 to 25 at 320x240 to the virtual frame buffer.

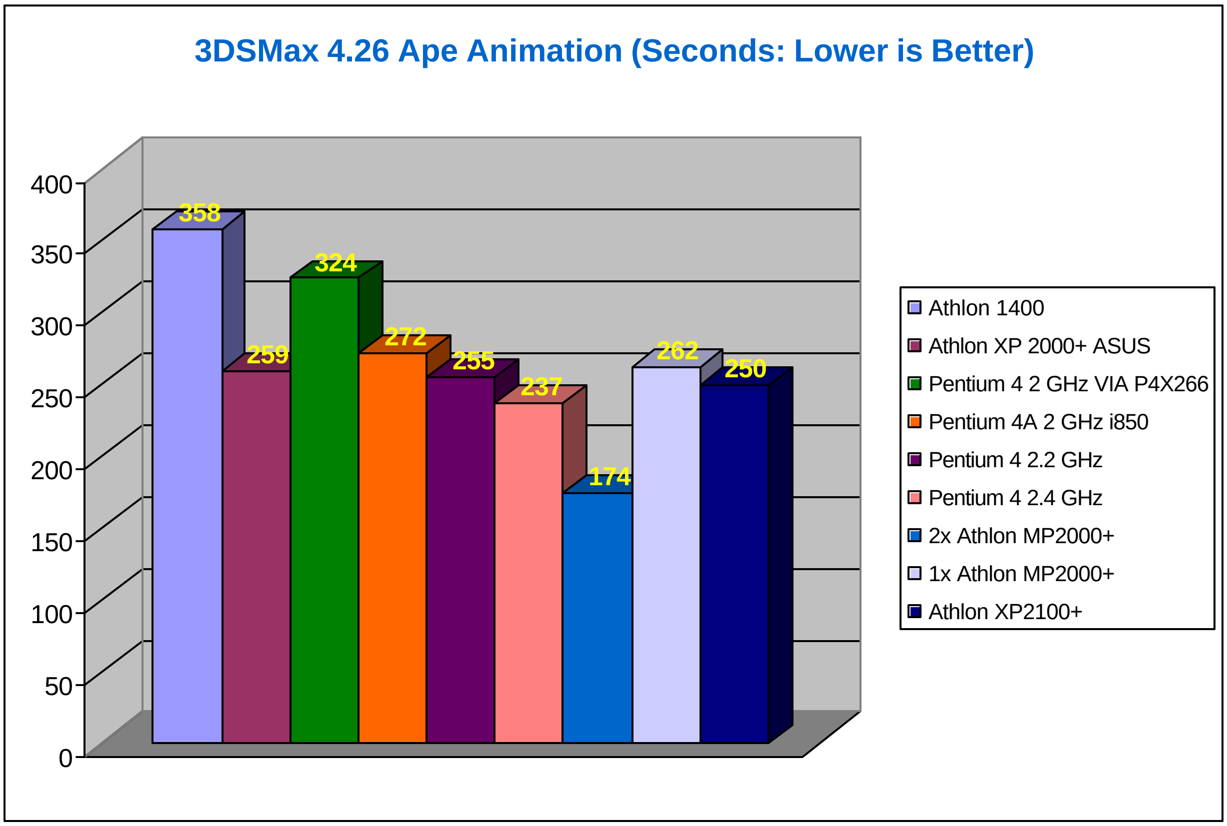

While the Pentium 4 performs well, there is no substitute for dual CPU power in 3D rendering. Dual Athlon MP systems are by far the most cost effective choice for the 3D artist.

Maya

Maya, developed by Alias|Wavefront, is one of the high-end 3D animation packages often used to produce art for commercials, movies, and games. Maya is available for Windows NT/2000, IRIX, Red Hat Linux, and Mac OS. We used the Maya 3 benchmark from the Maya-testcenter for these results. This benchmark renders one 512x512 image, based on a scene that features shadow maps, raytracing, displacement mapping, 3D textures, and post effects such as glow.

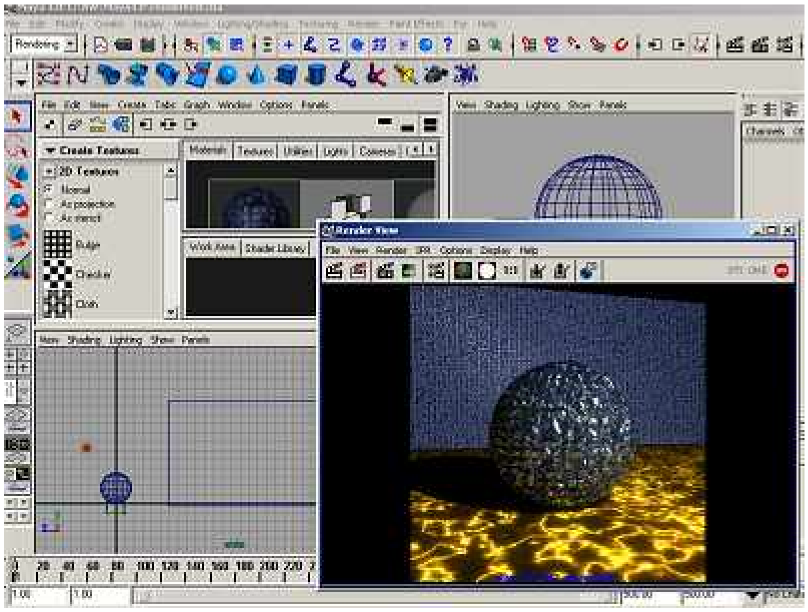

Please note that the rendered scene contains few polygons. Therefore, it only gives a partial image on how Maya performs on the different CPUs. Unfortunately, we ran out of time while trying to create a new Maya benchmark. With huge amount of polygons, it is likely that the memory subsystem becomes more important. But as long as you use raytracing, Maya's performance will mostly depend on CPU performance.

As we have been benchmarking this scene for quite a while, we inserted a few results from older reviews for comparison (marked with (*)).

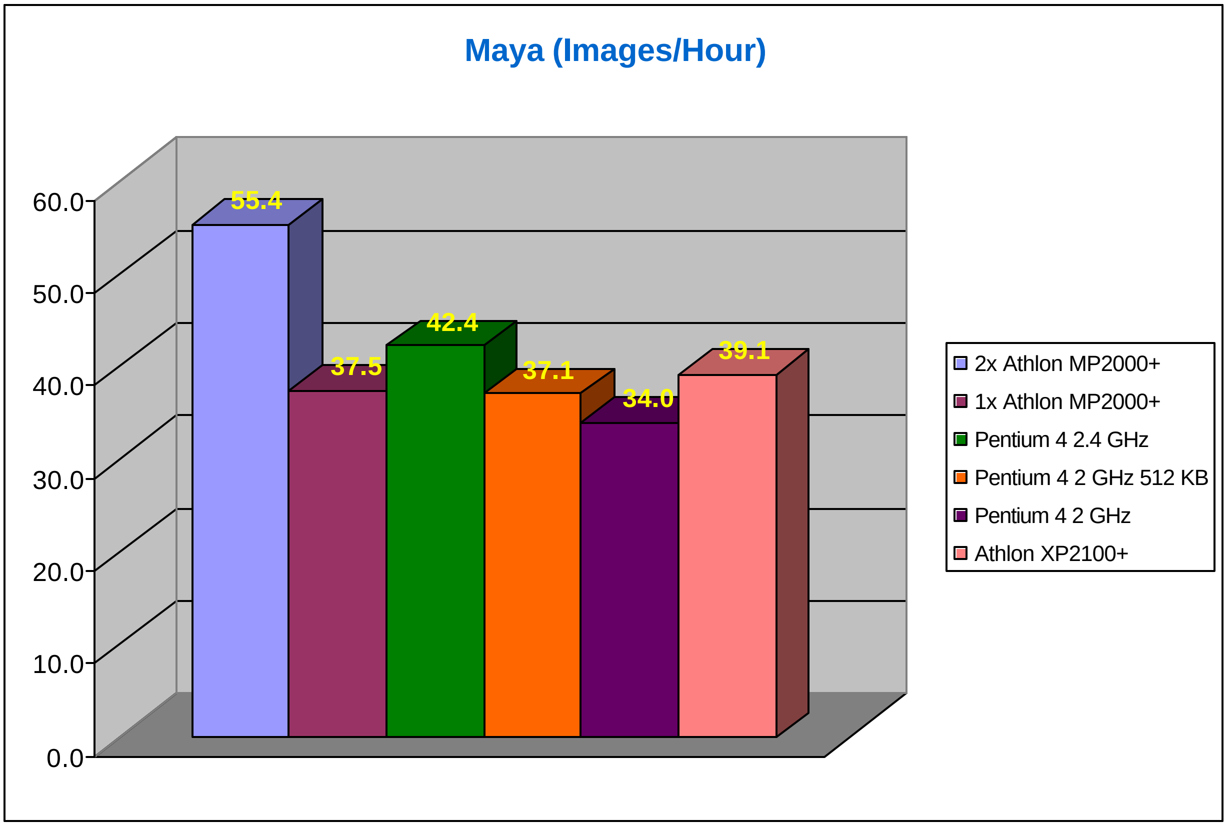

(*) Tests run on boards with the AMD-760MP chipset. See here

Notice that the dual Duron doesn't make much sense, even though is a very cheap option. Even if a dual Duron 1300 would scale as well as a dual Athlon MP (it won't), it would perform like an Athlon 1800 XP. The efficiency of a Dual Athlon is already limited by the memory bandwidth and the size of the caches.

The Tyan motherboard still accept Dual Durons, but the other boards now refuse to accept a non-Athlon MP as a second CPU:

Frankly, it is probably to be better safe than sorry, and pay a little more for an Athlon MP if you want a dual CPU configuration.

Real-Time Software Rendering: Kribi

Kribi, a product of Adept Development, is quite a remarkable 3D engine. Originally developed by Eric Bron, a regular visitor of Ace's Hardware, is is designed to handle up to 10 billion polygons and enable real-time photorealistic rendering. The Kribi engine uses 100% software rendering (a pure CPU benchmark) and cannot work without SSE instructions. The benchmark can therefore provide a bit of insight as to how powerful the SSE implementation of a certain CPU is.

The engine is marketed to developers who wish to develop full-fledged 3D-applications. We tested with the final version of KribiBench, and you will soon be able to download it for yourself at Adept Development.

KribiBench: City Ultra

The first scene - City Ultra - was quite spectacular: no less than 16.7 billion polygons in total.

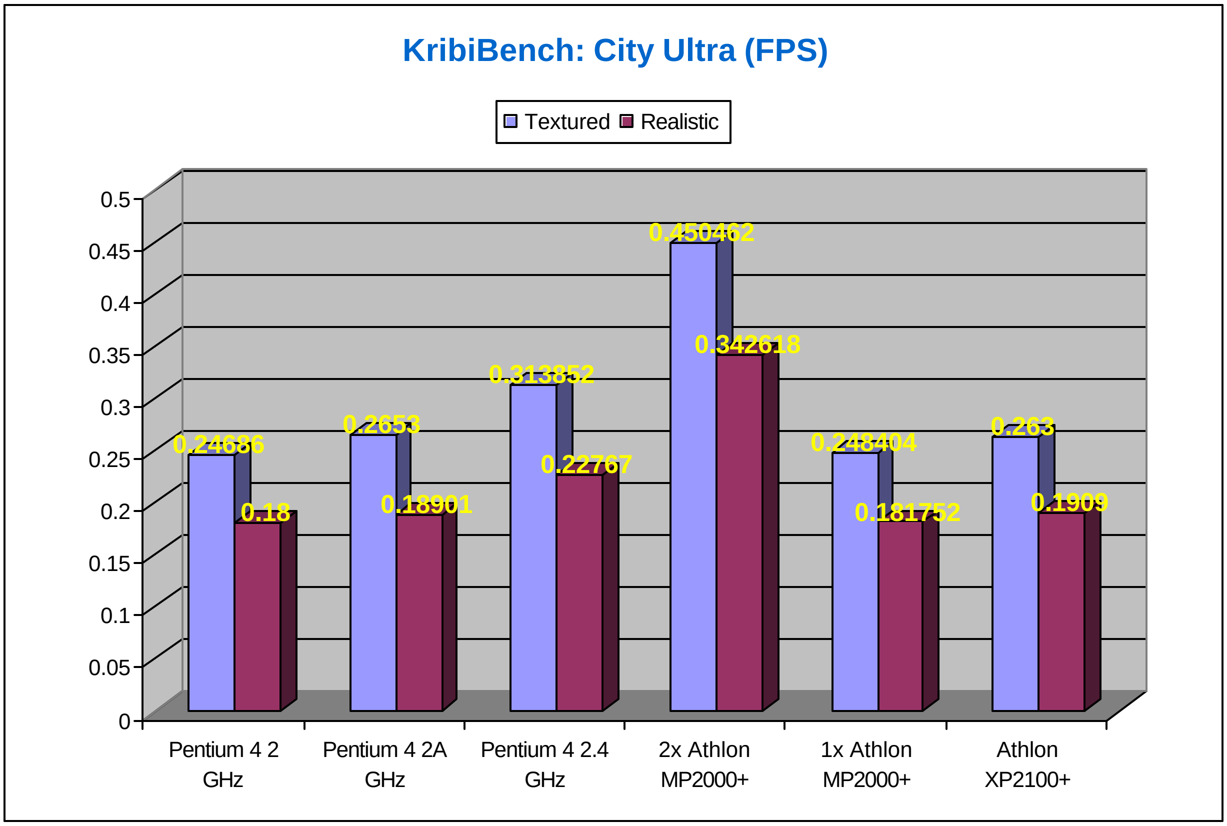

With 16.6 billion polygons, you can not expect miracles, of course, and the scene crawls on all CPU setups. The polygon counts are, however, the total number of polys in the model which is not the same as the number polygons which are actually rendered per frame. The Kribi Engine tries to cull the polygons which are not in view as quickly and efficiently as possible (occlusion culling). The Athlon 2000+ and 2 GHz Pentium 4 both perform very similarly, so even the combined power of the SSE engine and FPU of the Athlon lives up to the rating system...

As we have benchmarked with the City scene (106 million polygons) in the past, we include it once more. The Jet Shadow contains only 16,000 polygons, but features very realistic shadows and lighting. As SSE code tends to contain very few branches, the higher-clocked SSE unit and larger cache of the Pentium 4 gives it a slight edge on the Athlon 2000+ (1667 MHz).

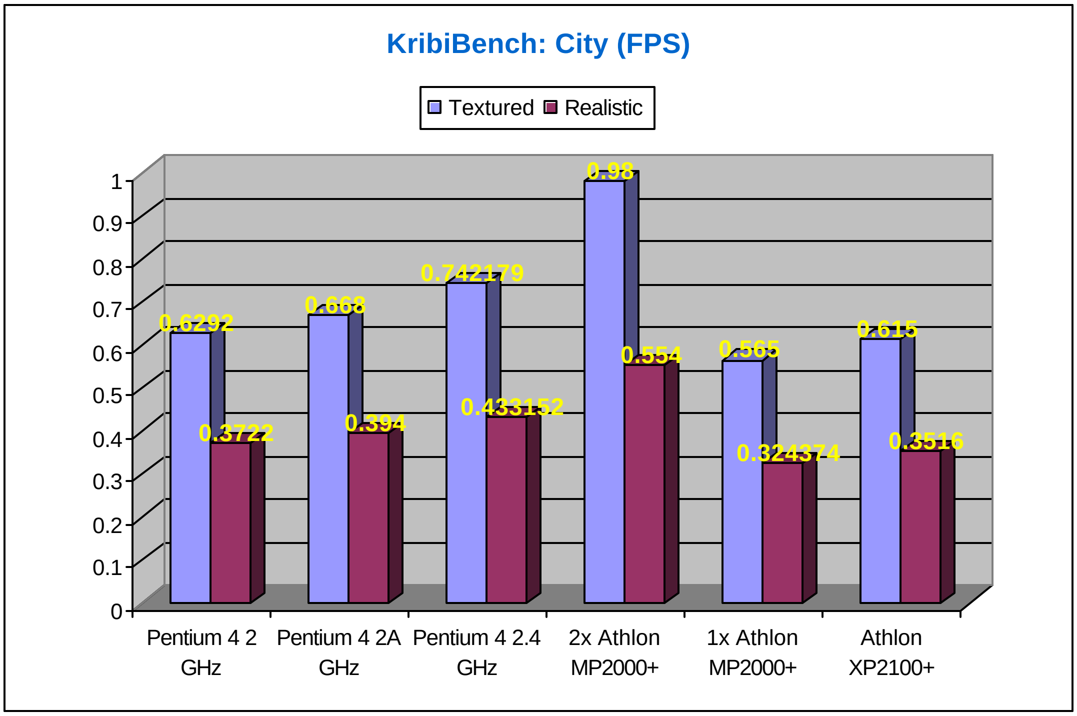

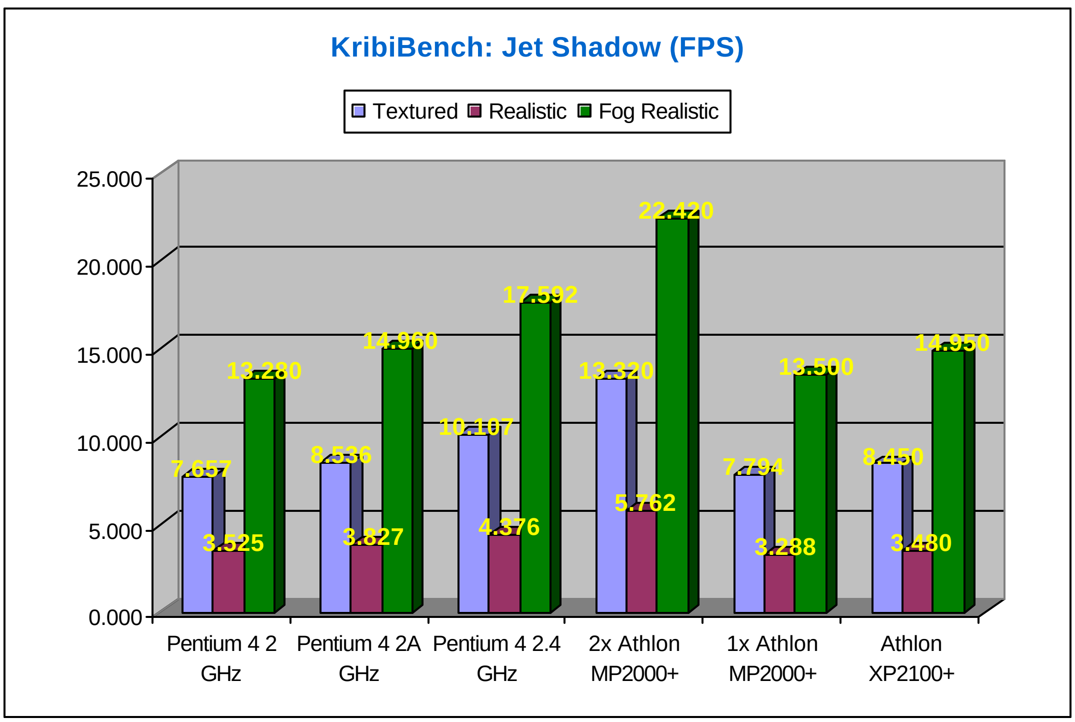

Windows Media Encoder

We used Windows Media Encoder 7.1 to encode a 24 MB AVI file to a streaming 1 MB WM8 video with a bitrate of 242 kbps (more info can be found here). Please note that due to the fact that we use Windows 2000, we were limited to WME 7.1. The Athlons should perform better with WME 8.0, which is available for Windows XP.

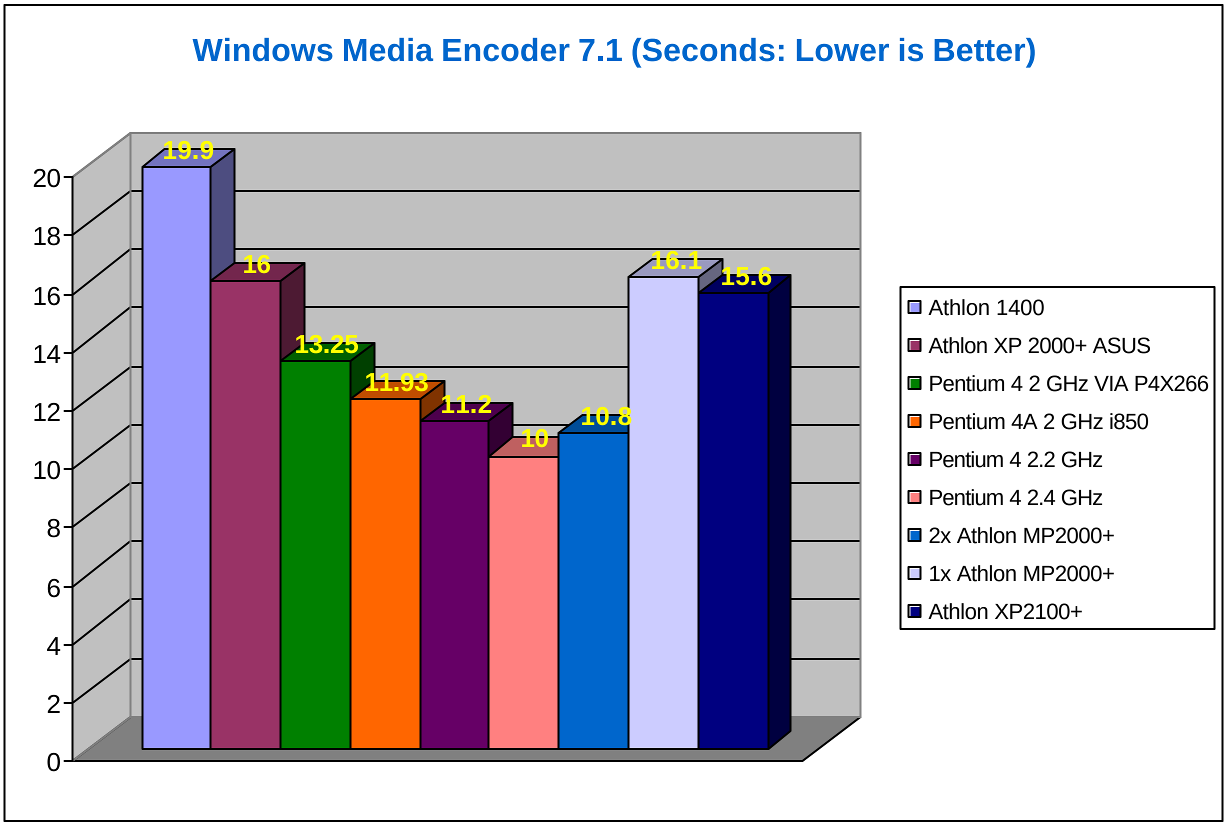

Although 24 MB of data has been processed, the benchmark is still CPU bound. Hardware and software prefetch seem to do their task very well, and the Pentium 4 2.4 GHz outperforms even the Dual Athlon MP.

Finite Elements

Our readers told us on our Message Board that "not everyone uses a workstation to produce pretty pictures," and requested that we tried out a Finite Element Analysis (FEA) package.

We agreed because Finite Element Analysis is an extremely useful mathematical method that has numerous applications in the engineering and scientific world. The basic idea behind FEA is reconstructing complicated objects as an assembly of small, simple pieces. In engineering, FEA is used to simulate and calculate the forces that appear inside the structure that is put under stress by external of internal forces. One of the objectives is to reinforce a structure where necessary, and design the lightest and strongest structure possible.

We used Phase² 5.0 from Roc Science as it can be downloaded for free, and most well known FEA software is extremely expensive. A short description follows:

> Phase² is a two-dimensional plastic finite element program for calculating stresses and displacements around underground openings, and can be used to solve a wide range of mining and civil engineering problems. Progressive failure, support interaction and a variety of other problems can be analyzed.
>
> Phase² provides an integrated graphical environment for data entry and visualization. A CAD based modeler allows for point and click geometry input and editing.

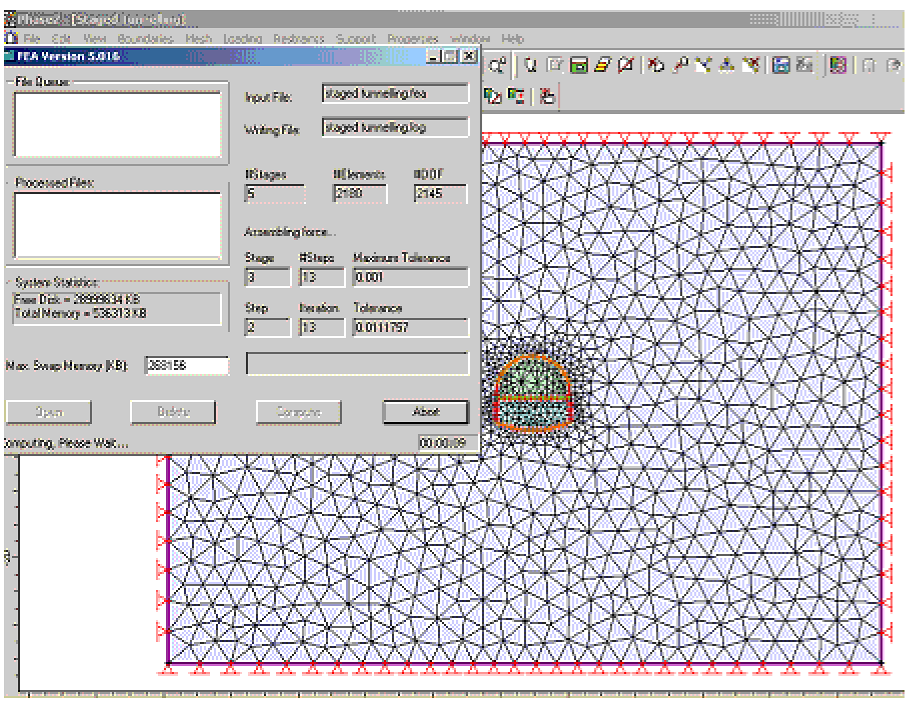

Unfortunately, Phase² 5.0 does not support SMP yet. In practice, most popular (expensive) FEA packages run very well on multiprocessor systems. FEA is very CPU intensive and many threads can run in parallel.

Phase² 5.0 Benchmark Results

We used the sample "Staged Tunneling" project for our benchmarks. The elements where 3 noded triangles, and to compute the forces by Gaussian elimination precisely, 500 iterations were used. Field stress was, of course, gravity.

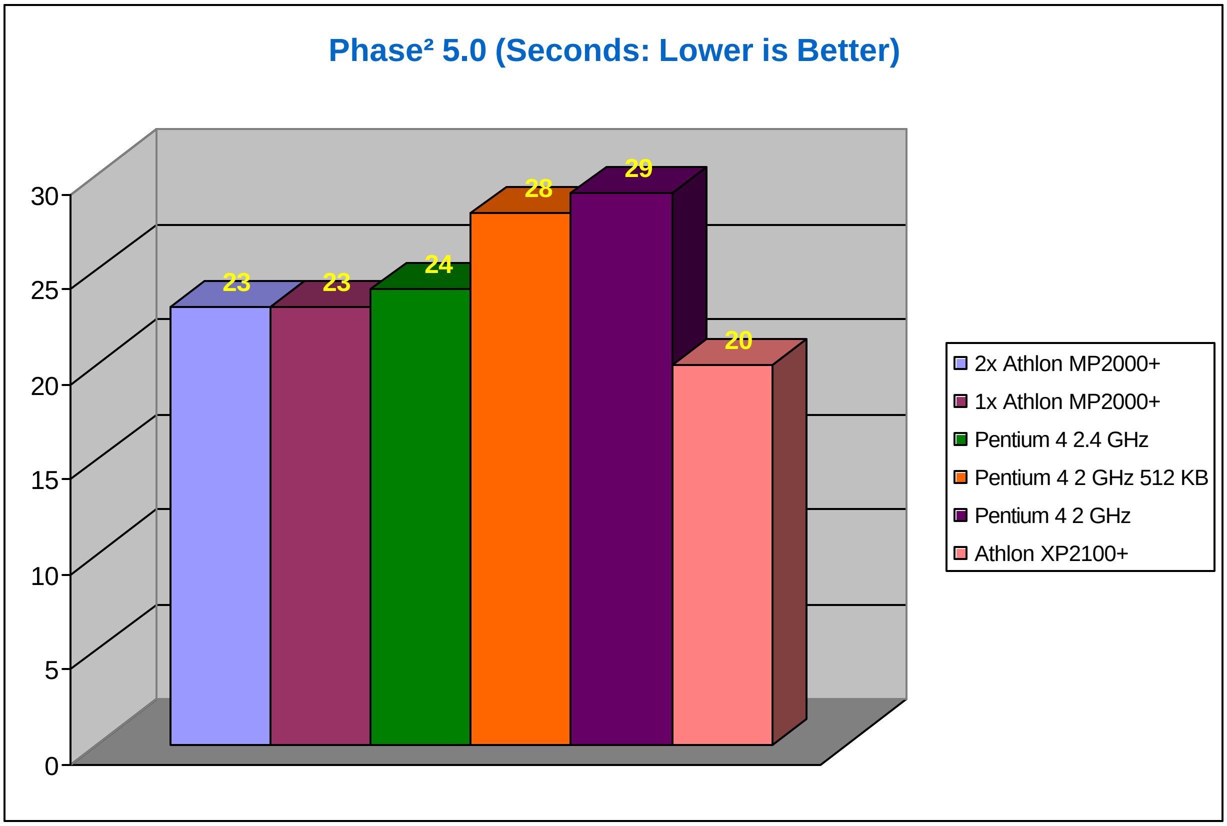

FEA is a total triumph for the Athlon's triple-pipelined FPU. FEA is extremely floating-point intensive, and you can never throw enough FP power at it. So we can assume that it is very likely that the high-end FEA packages will run very well on the dual Athlon.

Photoshop 6.01

Photoshop needs no introduction, as it is probably the most popular and widely used photo editor ever. Version 6.01 has optimizations for all MMX capable CPUs as well as special optimizations for the Pentium 4.

To test Photoshop, we used PS6Bench 1.11, a Photoshop action that runs 21 different operations on Photoshop 6.01 and measures the run time with Adobe's timing feature. We used the PS6Bench "Advanced" benchmark which manipulates a 50 MB image. The red scores are the best scores. With Photoshop 6.01 running, 300-340 MB physical RAM memory (out of 512 MB) was used, so disk swapping did occur on occasion, though not often.

Each filter ran 3 times, and the results you see here are the average of these 3 iterations. Still, we have run each action twice and we must say that the typical error margin is quite high, between 2 and 10 percent. This is particularly true for the Rotate 90 and reduce size filters (very high error margins: 10%). Tests 5, 8, 9, 13, 15 to 21 have low error margins (2-4%), all other tests are in between.

Photoshop Results

Test Nr.	Photoshop 6.01 Filter	Pentium 4 2 GHz - 256 KB	Pentium 4 2 GHz - 512 KB	Pentium 4 2.4 GHz	Athlon XP 2100+	Single Athlon MP 2000+	Dual Athlon MP 2000+	Dual Speedup
1	Rotate 90	8.8	10.9	7.4	9.2	12.6	12.1	4.7%
2	Rotate 9	17.9	17.3	16.4	13.2	16.7	15.9	5.5%
3	Rotate .9	17.6	17.3	15.3	14.9	16.9	15.9	6.5%
4	Gaussian Blur 1	8.6	8.6	6.5	7.8	10.7	9.4	13.9%
5	Gaussian Blur 3.7	13.2	15.3	12.9	10.7	17.6	15.4	14.5%
6	Gaussian Blur 85	14.6	14.3	13.7	14.0	18.5	16.2	14.2%
7	Unsharp 50/1/0	7.1	7.2	6.1	5.5	7.1	6.9	1.9%
8	Unsharp 50/3/7/0	13.8	13.8	13.5	13.9	17.9	15.7	14.0%
9	Unsharp 50/10/5	14.5	13.8	12.8	12.9	18.4	15.6	18.2%
10	Despeckle	10.7	10.4	9.4	6.7	7.7	7.0	10.5%
11	RGB-CMYK	23.1	23.5	20.0	14.4	17.2	16.9	1.8%
12	Reduce Size 60%	4.3	4.9	3.8	3.0	3.9	3.4	16.8%
13	Lens Flare	17.5	16.7	14.8	14.2	20.1	19.1	5.4%
14	Color Halftone	18.9	17.9	14.6	12.5	16.7	16.5	1.0%
15	NTSC Colors	7.4	7.4	6.4	6.6	7.2	7.2	0.0%
16	Accented Edges	21.9	21.2	18.2	16.7	19.3	19.5	-0.8%
17	Pointillize	38.6	38.5	32.8	35.2	37.9	37.1	2.2%
18	Water Color	46.1	46.2	38.2	32.4	33.9	34.6	-2.0%
19	Polar Coordinates	32.4	31.9	28.0	18.7	19.2	19.2	0.0%
20	Radial Blur	85.0	84.2	73.6	81.1	91.7	92.1	-0.4%
21	Lighting Effects	6.5	6.2	5.6	5.8	6.5	7.2	-9.7%

This is a lot of information, so let's look for some general trends. First of all, the larger cache of the Northwood does not seem to improve performance by much, but it is no wonder with 50 MB of data. Much more disturbing is the fact that a dual Athlon MP2000+ CPU machine is beaten by the Athlon XP 2100+ in every single benchmark. The latter is from 4 to 44% (!) faster than the dual machine. The second CPU offers little or no gain in most Photoshop operations, and it is clear that the Athlon with best memory subsystem wins. To understand this better, take a look at the following screen capture:

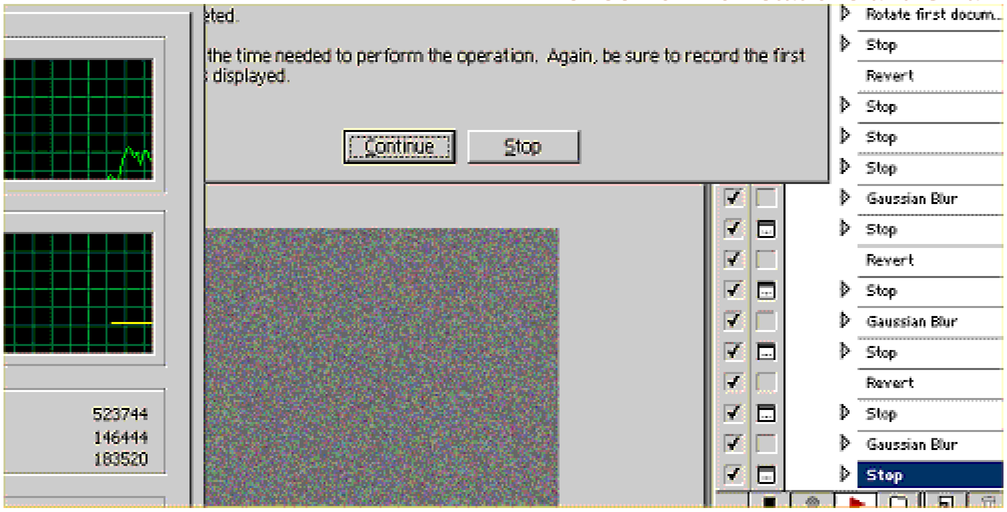

It shows the CPU usage graph of the Windows 2000 Task Manager just after finishing a Gaussian blur filter on the Athlon XP 2100+ machine. You can see that CPU usage never peaked higher than 20%, and this explains a lo t. The powerful 1733 MHz Athlon is waiting on the memory subsystem, and you can imagine that throwing in even more CPU power provides little or no gain.

The Radial Blur filter is another example. This time, the CPU usage of the single CPU machines was pushed to 100% but the filter doesn't seem to be multithreaded, as the dual machine was never pushed above 50%. Even the best multithreaded filters were not able to push the dual machine beyond 70% for more than a fraction of a second.

The "PS6Bench Intermediate" action creates a 20 MB image. Applying filters on a 20 MB picture showed even worse results for the dual Athlon compared to the single Athlon MP. The second processor decreased performance in almost every filter.

AMD's best CPU seems to be slightly ahead of Intel's 2.4 GHz showpiece: it wins 12 out of 21 tests. Considering that the Athlon XP 2100+ is much cheaper, a highly-clocked Athlon XP is the best choice for a Photoshop workstation -- a much better choice than a dual Athlon machine.

Database Servers

Database servers can be very CPU intensive, as well. SQL Databases typically produce several processes and/or threads, and run thus very well on multi-CPU configurations. From Chris' upcoming article:

> Databases are very complicated to setup, Oracle for example has more options than most operating systems and it has most of the functions of an one too. For SQL database systems, there is little in the way of "pre-packaged solutions", and a huge consultancy industry available to help. There are some large off-the-shelf applications that use SQL databases though (like SAP), but each has different characteristics.

A handful of different database packages (Oracle, Sybase, DB2, Informix...), thousands of options, makes benchmarking SQL databases a daunting and time-consuming task. Given the focus of this article on CPUs for workstations and small servers, we decided to make use of ZD's ServerBench 4.1 benchmark. This benchmark is not a true real-world database performance analysis, but instead a synthetic benchmark designed to simulate a multi-user workload similar to that seen with OLTP database workloads.

To help provide the resources for this test, I got the cooperation of my employer, the PIH West-Vlaanderen, a university in West Flanders (Belgium) which offers a Multimedia & Communication Technology Bachelors degree. Within the framework of the "Multimedia & Communication" department, I got the assistance of one my students from last year, Stijn Verbauwhede. I'd like to thank Dean Lode De Geyter for his kind support of this project and Stijn verbauwhede for all his hard work and assistance.

ZD ServerBench: What it Does

Getting back to ServerBench, this is what ZD Labs (now called E-Testing Labs) has to say about the benchmark:

> ServerBench is a ZD benchmark program that lets you measure the performance of an application server in a client/server environment. When we talk about an application server, we're referring to a setup where most of your data and your applications execute on the server. The client PCs send requests for work to the applications on the server. Once the server finishes the work, the application sends a response to the client. The client then performs some work related to the response.

ServerBench simulates a two-tier model where the database and application servers are located on the same machine. Serverbench consists of 7 tests:

- 1 Processor Test
- 4 Disk Tests
- 2 Network Tests

Sizing a complete (Disk, CPU, memory, and network) database server is beyond the scope of this article, as this article focuses only the CPU specific subtests (the "Processor Test"). The other tests were part of our database performance project in cooperation with the PIH, but will be discussed later. We limited ourselves to a mix of CPU intensive tests from ZD Serverbench (p_60.tst) because we wanted to compare the different CPUs. The disk tests - which are not included - clearly indicated the differences between, for instance, a disk system consisting of a three disk setup in RAID 5 instead one SCSI disk.

The processor test mimics the processor intensive portions of typical database servers, and measures the server's ability to perform data searches, string manipulation, integer arithmetic, memory copying etc. It uses common programming constructs, such as if-else, looping, and procedure calls.

More About ServerBench

The manual explains how ZD Serverbench simulates a Database server and application server:

> The processor test works by manipulating a 125-byte record. This record contains the following fields:
> * A Row ID, a two-byte integer
> * A Customer ID, a two-byte integer
> * A Salesperson ID, a two-byte integer
> * Available credit, a four-byte integer
> * Amount bought, a four-byte integer
> * Outstanding balance, a four-byte integer
> * Order pending, a one-byte flag field
> * Business type, a one-byte enumeration 1 to 10
> * The customer's name, a 20-byte character string
> * The customer's address, a 75-byte character string
> * The customer's phone number, a 10-byte character string
> The data set is 2800 records of 125 bytes for each client for a total of 350,000 bytes of data for each client.
>
> ServerBench manipulates the data file by:
> 1. Totaling the outstanding balances of each salesperson's customers.
> 2. Counting the number of customers who have exhausted their credit.
> 3. Dividing the total amount sold by the number of people to produce an average.
> 4. Searching the data for business type "1" and copying that data into the temporary buffer.
> 5. Searching the data for customers with outstanding orders, sorting the data, and copying that data into the second temporary buffer.
> 6. Looking at the amount of goods purchased for both of the subsets described above and, for the complete set of records, computing the minimum, maximum, and average amount of goods purchased in each group, and then totaling the amount of goods purchased in each group.
> 7. Giving a total amount due the company.
> Because ServerBench doesn't display the final results of this test, it doesn't format the data.

When ServerBench runs the processor test, the server allocates 350,000 bytes of memory for each client and reads the data into the client's workarea. In other words, the data set is memory resident and there is no interaction with the disk at all. So, the results are purely synthetic and not reflective of real database server performance, but they do give us a look at the role the CPU plays as a component in overall database server performance.

Benchmark Setup

We used a relatively simple setup: six clients, one controller, and one server were connected via 100 Mbit network cards on a 100 MBit switch. The network used TCP/IP as the network protocol. The server, running Windows 2000 Advanced Server (Service Pack 2), was setup to mimic a small business server, and as such it was the Domain controller and the database server at the same time. Up to two ServerBench clients were run on each client.

All clients had the following configuration:

- Windows 2000 Professional SP2
- PIII 733 or Athlon 800
- 256 MB of PC133 CAS 3 SDRAM
- VIA KT133 chipset or VIA apollo 133
- Seagate Barracuda ATA III ST320414A Model ST320414A (7200 RPM, ATA-100)
- Realtek RTL8139(A) PCI Fast Ethernet Adapter (100 MBit)

The configuration of server was identical to our configuration info page, with one exception: we used one Maxtor Atlas II 9 GB 10,000 RPM SCSI disk instead of the IDE disk. Below you can find a partial view of our test setup:

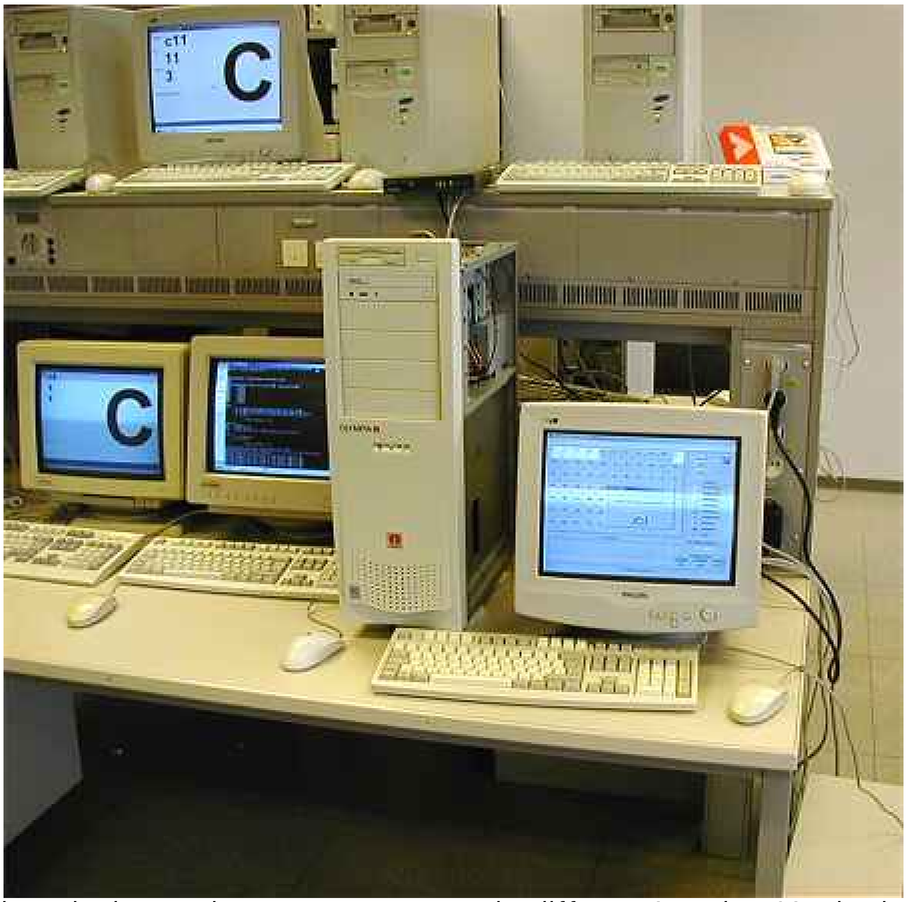

The small LAN in the PIH that was set up to test the different CPU's in a SQL database server

Results of ZD ServerBench 4.1 are expressed in "Transactions Per Seconds" or TPS. From ServerBench manual:

> A transaction consists of the request a client sends to the server, the response it gets back, and the time it takes from the moment the client sends the request until it receives a reply from the server. ServerBench includes the time the transaction spends:
>
> - Traveling along the network to and from the server.
> - Waiting in a queue on the server to receive a service.
> - Receiving the service. For example, if the transaction requires disk service, this is the amount of time the disk took to provide the service
>
> To determine the overall server TPS scores it produces, ServerBench:
>
> - Tracks the amount of time each transaction takes to complete.
> - Tallies the number of completed transactions. ServerBench does not count incomplete transactions or transactions that began during the Ramp up or Ramp down periods.
> - Creates a total TPS score for each transaction by adding together the TPS score for each client.
> - Uses a weighted harmonic mean to turn the total TPS scores into an overall score

Let's take a look at our results. We ran up to 12 processes simultaneously. That might seem a bit low, but you must remember that those 12 processes run without any pause. In other words, our results are not only valid for a LAN with only 12 clients running one process. Typically, unless every user of every client does heavy "datamining," periods of intensive database queries are interleaved with "non-active" periods. These results might be valid for a LAN with maybe 3 times more clients.

ZD Serverbench 4.1 (TPS)	1 process	4 processes	8 processes	12 processes
1x Pentium III 733 MHz	53250	52133	52189	52323
Dual Pentium III 733 MHz	50920	87815	89264	87811
1x Athlon MP 2000+	121367	120722	121134	121349
Dual Athlon MP 2000+	118453	224108	213489	231333
Pentium 4 256k 1.6 GHz	70763	71521	71303	71569
Pentium 4 256k 2 GHz	80902	81739	81439	81748
Pentium 4 512k 2 GHz	127629	125458	127241	126994
Pentium 4 512k 2.4 GHz	150709	151488	151285	151919

We included a Dual Pentium III machine for reference, as many small companies use this kind of configuration for their small business servers. Database workloads can often be very branch and memory intensive, and therefore can be very sensitive to cache size/performance. As a result, the Pentium 4 Willamette, with it's huge branch misprediction penalty and small L2-cache, is a weak database server CPU. An extra 256 KB of cache in the Northwood boosts the performance of the Pentium 4 enormously, we noted no less 55% better performance! The larger cache allows the Pentium 4 Northwood to scale much better. A 20% (from 2 GHz to 2.4 GHz) boost in clockspeed results in 19.6% increase in performance, while the Pentium 4 Willamette sees only a 14% boost from a 25% (2 GHz versus 1.6 GHz) higher clockspeed. Nevertheless, it must be noted that the 1667 MHz Athlon MP 2000+ MP, with its exclusive 128 KB L1 and 256 KB L2 caches, can still compete with the new Intel Pentium 4. Additionally, the CPU intensive subtests of ServerBench clearly benefit from a dual CPU setup. The second Athlon offers up to 90% better performance with more than one process active.

Looking at the enormous performance increase the Pentium 4 received from doubling the cache, AMD should outfit the Athlon MP with a larger cache to ensure success in the server market. And, of course, a better memory subsystem would help too. Looks like we have just found out what AMD's upcoming Hammer CPU is all about...

Conclusion

Right now the Athlon MP performs just as well as the Pentium 4 Northwood, which has 128 KB more cache. Considering the very high prices of a Xeon motherboard and dual Xeon configuration, the Athlon MP is a very cost effective, high performance solution for a workstation or low-end server. The Athlon MP lacks some cache to be a great server CPU and outperform Intel's offering by a wider margin, but the Athlon architecture has great potential. Therefore, if AMD can live up to their promises, the Hammer family could - with proper OEM support - will prove to be a fierce competitor in the server market. The reason why we say this, is that the Hammer CPU strongly improves the weak points of the Athlon MP: the performance of memory subsystem, the branch predictor, and most importantly, the features that improve manageability and reliability (Chipkill, ECC, Heatspreader, Thermal diode support, etc.).

Another interesting observation is that there is no "best" workstation processor overall. The price is too high compared to the Athlon XP 2100+, but Intel's 2.4 GHz Pentium 4 is the fastest single processor for 3D rendering (SSE-2 optimized software like 3DS Max, Maya, and Lightwave), video streaming, and software rendering. The high clockspeed and extra cache cannot hide the relatively slow FP unit, as shown by the FEA benchmark, where the Athlon XP 2100+ is by far the fastest processor.

The Photoshop benchmarks showed us surprising results as well: a dual CPU system is of little use compared to a faster single processor system in this application. A fast clocked single processor with the best memory subsystem performs much better, between 4 and 44%!

The old dual Xeon (Foster, 256 KB L2) doesn't make much sense anymore, as a higher-clocked Pentium 4 will perform just as well. Neither does a dual Duron, which is outperformed by fast-paced Athlon XPs.