Professional Grade Revisited

Wildcat III 6110, FireGL 8800, and Quadro 4 XGL Reviewed

By Johan De Gelas – May 2002

The battle in the professional OpenGL market is heating up and NVIDIA is no longer satisfied with delivering budget cards like those based around the Quadro 2. With the Quadro 4, however, Nvidia's burning ambition is nothing less than dethroning the excellent Wildcat series from 3DLabs. Nvidia claims that the "**Quadro 4 XGL series sets the standard** for professional graphics by delivering breakthrough application performance."

3Dlabs, the defending king of the high-end professional OpenGL market is not impressed and replies with the latest incarnation of the successful Wildcat series, the Wildcat III. And they have made a few **bold claims** of their own: "The Wildcat III is in a league above the competitors, games-based graphics technologies face huge obstacles in entering the high-end segment."

Nevertheless Nvidia's armada is impressive: it has the Quadro GXL 900 to attack the Wildcat III 6110, the Quadro GXL 750 to ensure Nvidia continuing supremacy in the mid-range OpenGL segment and the Quadro GXL 700 and 550 to ensure success in the budget market. Meanwhile the people at ATI Starnberg are not sitting still either, and they have challenged Nvidia and 3Dlabs in the midrange OpenGL accelerator market with the FireGL 8800.

The stakes are high: investments in OpenGL hardware continue to increase, as estimations vary from 9 to 15% growth each year. And as **the Register reported**, x86 workstations have been gaining ground over RISC systems. The market for OpenGL accelerators and Windows x86-based workstations is still a very attractive and growing one.

More than reason enough for us at Ace's Hardware to produce a comprehensive review of OpenGL accelerators. as they are, after all, important tools for exciting fields like game and film development, scientific visualization, and CAD. In this review, we'll show you how several different OpenGL video cards perform in 3D Studio Max 4, Maya 4, Pro-E, UGS, AutoCAD, and more. We'll start off with a list of the competitors in this review:

- **3DLabs Wildcat III 6110**
- 3DLabs Wildcat II 5000
- ATi FireGL-2
- **ATi FireGL 8800**
- Nvidia Quadro 2
- **Nvidia Quadro4 900XGL**
- **Nvidia Quadro4 750XGL**
- Nvidia GeForce 4 Ti4400
- Nvidia GeForce 4 MX460
- Nvidia GeForce 2 Ti 200

Nvidia's Quadro 2 (Elsa Gloria III) is included to show how the new generation of accelerators compares to the previous generation. This way, it should be easier for you to estimate how much return on investment you can expect to see from upgrading to one of the latest professional 3D OpenGL accelerators (i.e. Wildcat III 6xxx series, Quadro 4 series, or FireGL 8800). Just like last time, we've included commodity gaming cards like the Leadtek GeForce 2 Ti200, MSI GeForce 4 MX460 and ASUS GeForce 4 Ti4400. Many of our readers are working on tight budgets, and as such, they might consider a relatively inexpensive gaming card with T&L acceleration for their workstation. And while the current pack of gaming cards are optimized for complex multi-texturing games, they still pack some healthy triangle rates. We'll find out how much performance improves when you invest extra money in a "real" OpenGL card.

If you'd like to see a comparison of the previous generation of OpenGL accelerators, you can find a review of the **3Dlabs Wildcat II 5110 & 5000, Elsa Gloria III (Quadro 2) and ATi FireGL 2 right here**. Let's now focus on the current generation.

Wildcat III 6110 & 6210

3DLabs continues to build upon their ParaScale architecture, the architecture that was introduced with the Wildcat 4210. The Wildcat III 6110 can indeed be described as an improved version, an evolution, of the venerable Wildcat II 5110. Typical for this Wildcat ParaScale architecture is that the frame buffer and texture buffer are completely separate as opposed to all other competitors that use a unified memory architecture. It is also one of the features that makes the Wildcat architecture very scalable: performance increases very well as more pipelines are added.

For example, the Wildcat II 5000 consists of a single pipeline and has a 128-bit interface to the framebuffer and a 64-bit interface to the texture buffer. The Wildcat III 6110 and Wildcat II 5110 consist of two pipelines that still access the same 128 bit framebuffer, but each pipeline has its own texture buffer with a dedicated 64-bit path per chip, so textures can be fetched at 128-bits per clockcycle. In a way you could say that the Wildcat II 6110 has a 256-bit memory interface. This 256-bit interface is, of course, not as powerful as a full 256-bit unified memory architecture, as it is unlikely that the framebuffer and texture buffer will require exactly the same bandwidth, but offers some extra bandwidth over a "normal" 128 bit unified architecture.

A minor disadvantage is the inflexibility of the architecture. While the Wildcat III has a total of 208 MB of RAM onboard, "only" 64 MB is dedicated to the framebuffer. As a result, the high-quality SuperScene anti-aliasing of the Wildcat III 6110 is only available at resolutions below 1280x1024. The FireGL 8800 and Quadro 4 can offer slightly lesser quality anti-aliasing at higher resolutions, as they are able to use almost the full 128 MB of onboard memory as a framebuffer and swap the textures in from the main memory. If you want SuperScene anti-aliasing at high resolutions, you have to pay up to $2500 for the Wildcat III 6210, which comes with a 128 MB framebuffer and two 128 MB texture buffers (256 MB total). This Wildcat III 6210 can of course offer anti-aliasing, at very high quality, resolutions, and speed. However, full scene anti-aliasing is not very important for most OpenGL professionals as it is only applied in the final stage of the creative process. And more importantly, the image quality at high resolutions is already very good, as jaggies are very small at 1600x1200 and above.

So what makes the Wildcat III different from the Wildcat II? The picture below explains it all:

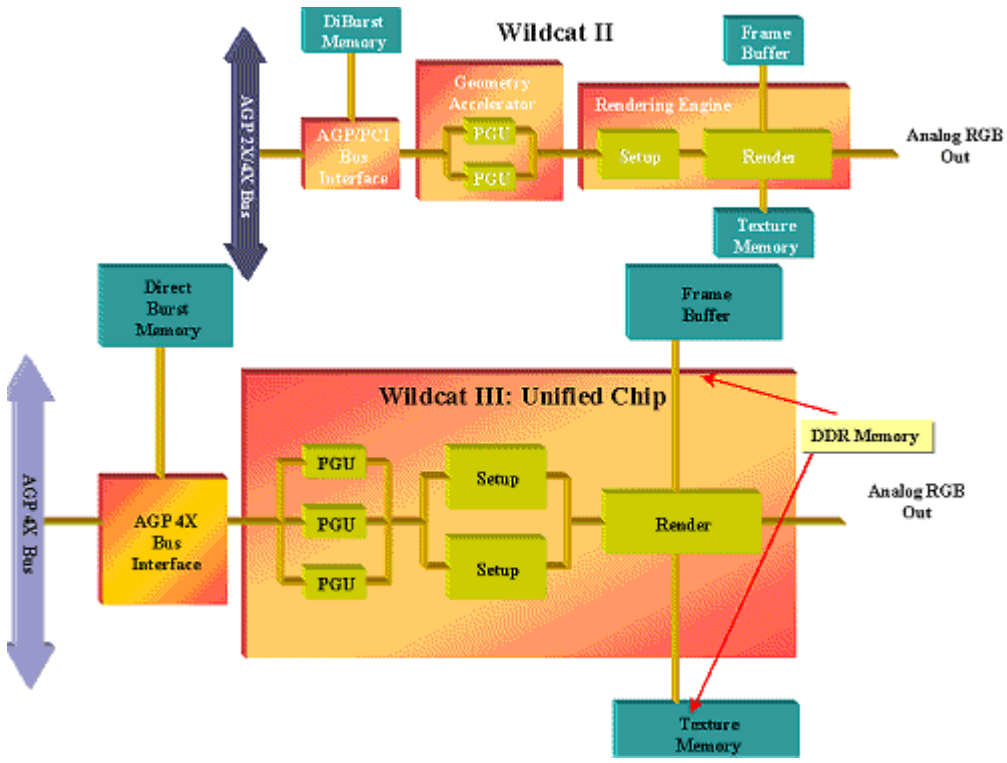

The Wildcat III consists of the same building blocks as the Wildcat II, but that doesn't mean the Wildcat III is an exact copy of the Wildcat II. The Wildcat II consisted of a separate geometry accelerator and rasterizer, whereas on the Wildcat III everything has been integrated into a single ASIC (Application Specific Integrated Circuit). However, the Wildcat III has 3 geometry units per chip instead of 2, and the dual pipeline of the Wildcat III 6110 is thus outfitted with no less than 6 PGUs. The Wildcat III also has twice as much setup engines per chip (2 instead of one) and has access to DDR SDRAM memory instead of normal SDRAM.

The rendering unit has improved gradient fill performance (important for shading) and improved multi-texturing support. Especially the DDR support and the improved rendering unit should fix a "weakness" of the 5110: shading and textured preview performance in 3D animation software. The Wildcat II 5110 reigned supreme in all wireframe tests, but the most expensive card occasionally tasted **defeat in our Maya 4 shading and texture preview** benchmarks, most likely because the card was limited by its fillrate.

For game developers, there's one minor letdown: the Wildcat III supports only DirectX 7.0. But 3DLab's first target is, of course, the CAD and high-end digital content creation market.

Nevertheless, it must be said that the Wildcat III still has some serious trump cards up its sleeve. The flagship of 3DLabs impresses with hardware support for no less than 32 lights, which is 2 to 4 times better than the competition. Also the Wildcat III is the only 3D accelerator tested here with a triangle setup engine that computes positions of edges of triangles and lines with 10-bit precision inside a pixel (10-bit sub-pixel/spatial precision). If polygons are drawn with poor accuracy and, as a result, are misaligned in the model, holes or cracks appear. These imperfections in polygon alignment not only greatly weaken the integrity of the model, but also can cost time and money further along in the design process.

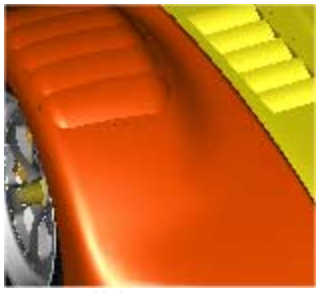

Wildcat III 6110

The image above is a full color bitmap with no dithering. The result is that the colors are not the original ones, but the pixel errors will still show. In the image above we found one pixel error, just above the wheel (black pixel).

3DLabs has **already announced** their next generation high-end 3D accelerator, the 3DLabs P10 VPU. However this new accelerator won't be available before the end of this year. The Certified application list for the Wildcat III 6110 can be found **here**.

ATI FireGL 8800

While they both target the more or less the same market, ATI's primary market for the **FireGL 8800** is slightly different from that of the 3DLabs Wildcat III. The ATI FireGL 8800 is clearly more of a digital content creation board, while the FireGL-4 card remains ATI's best CAD product. They both carry the same name, but the FireGL 8800 is based on a totally different architecture than the FireGL-2, 3 and 4. The latter accelerators are based on IBM's OpenGL accelerators, which were originally designed for UNIX workstations.

The FireGL 8800 is in fact a Radeon 8500 chip at 250 MHz with improved line anti-aliasing performance, faster memory chips (3.3 ns instead of 3.6 for Radeon 8500) and OpenGL optimized drivers.

The FireGL 8800 is especially targeted towards game developers, which can design games and then play them on the same platform, as the FireGL 8800 fully supports DirectX 8.1. What are also remarkable are the low prices at which the FireGL 8800 can be found: often below $700. The little brother of the 8800, the Fire GL 8700 can even be found for $300! This version is the workstation equivalent of ATI's Radeon 8500LE, and features a 250 MHz chip and 250 MHz DDR SDRAM.

The FireGL 8800 only features 4-bit sub-pixel accuracy, however, meaning it is very likely that small holes will appear in a shaded model. We saw some holes in the PTC car model in the Pro-E, which where not noticeably with the Wildcat III 6110.

FireGL 8800

As you can see here, the FireGL 8800 has quite a few more errors in its output. Concentrating on the red orange part of the car, you can easily spot 4 black pixels that shouldn't be there (the tire is showing through) and 2 gray pixels that also indicate an error.

Certified drivers can be found **here** and **here** is the full list of certified applications.

NVIDIA Quadro 4 900 & 750XGL

Ladies and gentlemen meet the challenger of the Wildcat III: the Quadro 4 900 XGL. Yes, according Nvidia, this is the card that will dethrone the Wildcat III. A 300 MHz Quadro 4 chip that can push up to 60 million triangles per second and render more than a billion pixels per second sure sounds impressive.

The Quadro 900 XGL can be described as a GeForce 4 Ti 4600, to which an ultra modern line anti-aliasing engine has been added. We will see further how important this improved AA line engine is for a typical OpenGL card. Like the FireGL 8800, the Quadro 4 enables game developers to design and test their games on the same platform, as the Quadro 4 fully supports all DirectX 8.1 operations.

Besides the XGL 900, The Quadro 4 XGL comes in two other forms: the 750 XGL and 700 XGL. The 750 XGL is probably the most dangerous card for the competition as the chip is clocked only 25 MHz slower than the 900 XGL. The memory is clocked at 275 MHz DDR (50 MHz slower than the 900 XGL), but this second best offering from NVIDIA comes with a very attractive price tag that hovers around $700. The 700 XGL, on the other hand, is identical to the 750 XGL but features only 64 MB of RAM.

The Quadro 4's sub-pixel accuracy is 8-bits, a significant improvement over the Quadro 2. It was well known that the Quadro 2's image quality was no match for the image quality that traditional OpenGL accelerators delivered. The image quality of the Quadro 4, however, as far as our relatively untrained eyes can see, is very good. With the naked eye the difference between the Wildcat III and Quadro 4 was almost unnoticeable.

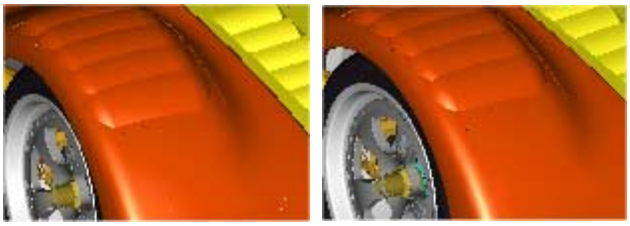

Quadro 2 *Quadro 4*

With 5 white holes and at least 2 "bad" black pixels in a very small area (red one), the Quadro 2, shown on the left, produces a picture that leaks like a sieve. Nvidia has made enormous progress, and in the Quadro 4 picture, we can only distinguish two black pixels, which are probably errors.

A list of certified applications for the Quadro 4 can be found **here**.

OpenGL Accelerators: An Overview

Here you find the specs of all the competitors of this review in one table.

Manufacturer	Nvidia	Nvidia	NVIDIA	ATI	ATI	3Dlabs	3Dlabs	3Dlabs
Product	**GeForce 4 MX460**	**Gloria III**	**Quadro 4 750-900XGL**	**FireGL-2**	**FireGL 8800**	**Wildcat II 5000**	**Wildcat II 5110**	**Wildcat III 6110**
Price	$140	$750-900	900: $1100-1250 750: $650-750	$850	$580-700	$700-750	$1400-1500	$1850
AGP Interface	AGP 4X	AGP 4X	AGP 4X	AGP 4X	AGP 4X	AGP 4X (Pro 50)	AGP 4X (Pro 50)	AGP 4X (Pro 50)
Dual Display	No	No	Yes	No	Yes	No	Yes	Yes
1st Display Out	VGA	VGA	750 XGL: VGA 900 XGL: DVI-I	VGA	VGA	VGA	VGA	DVI-I
2nd Display Out	n/a	DVI-I (analog & digital)	DVI-I (analog & digital)	DVI-D (digital)	DVI-I (analog & digital)	DVI-I (analog & digital)	DVI-I (analog & digital)	DVI-I (analog & digital)
RAMDAC	350 MHz	350 MHz	Dual 350 MHz	300 MHz	400 MHz +240 MHz external	300 MHz	300 MHz	320 MHz
Max Refresh Rate @ 1600x1200	120 Hz	120 Hz	120 Hz	85 Hz	100 Hz	90 Hz	90 Hz	90 Hz
Max Refresh Rate at 1920x1200	100 Hz	100 Hz	100 Hz	75 Hz	100 Hz	75 Hz	75 Hz	75 Hz
Geometry & Rasterizer	GeForce4 MX (300 Mhz)	Quadro2 Pro (250MHz)	Quadro4 900XL: 300 MHz 750XL: 275 MHz	IBM GT1000 (190MHz), IBM RC1000 (120 MHz)	250	1 x Wildcat II Chipset 225 MHz 5110 RE - 200 MHz	2 x Wildcat II Chipset 5110 GA - 225 MHz 5110 RE - 200 MHz	2 x Wildcat III Chipset 220 MHz

Memory Type	275 MHz DDR	250 MHz DDR	900XL: 325 MHz DDR 750XL: 275 MHz DDR	120 MHz DDR	290 MHz DDR	166 MHz SDRAM?	166 MHz SDRAM?	220 MHz DDR
Memory Size	64 MB	64 MB	128 MB	64 MB	128 MB	32 MB Texture + 32 MB Frame	64 MB Texture + 64 MB Frame	128 MB Texture + 64 MB framebuffer
Memory Bus	128 bit	128 bit	128 bit	256 bit	128 bit	128 bit	256 bit*	256 bit*
Texture Fillrate (Trilinear, Mpixel/s)	600	1000	1100/1200	200	1000	166	332	400
Polygon Rate (MTriangles/s)	25	31	60	27	45	8.5	15	33
No. of Hardware Lights	8	8	8	16	8	24	24	32
Sub-Pixel Accuracy (bits)	4	4	8	4	4	10	10	10

*128-bit framebuffer, 2 x 64-bit texture buffer

The Quadro 4 and Fire GL 8800 both feature high bandwidth memory busses, fast RAMDACs, and very high refresh rates, even at the highest resolutions. The 3DLabs board, on the other hand, features a huge texture buffer, a large number of hardware lights, and ultra-precise sub-pixel accuracy. A comparison of the different dual display possibilities is out of the scope of this article, but we may cover this in a future article.

Benchmark Configuration

All tests where performed at a 75 Hz refresh rate, V-sync was off, and the desktop was always set to 32-bit color. Our test bed still runs Windows 2000 SP2, as most professionals and hardware enthusiasts prefer the matured Windows 2000 SP2.

3D Accelerators

- 3DLabs Wildcat 5000 and 6110: 3D Labs 05.05.03.25 driver
- NVIDIA Quadro 4 and Quadro 2: NVIDAI Detonator 29.13, MaxTreme 4.00.21 (3DS Max 4.26), Powerdraft 15.00.10 (AutoCAD)
- NVIDIA's gaming cards: NVIDIA Detonator 29.13
- ATI Fire GL 8800: ATI display driver 5.12.2195.3030
- ATI Fire GL 2: ATI display driver 5.12.2195.2088

System:

- 2.4 GHz Northwood Pentium 4
- ASUS P4T-E (i850 RDRAM chipset) bios version 1.005
- 512 MB Corsair RDRAM-45 ns (4x128 MB)

Common Hardware

- Seagate Barracuda ATA III ST320414A Model ST320414A 20 GB (7200 RPM, ATA-100)
- AT 2700 10/100 NIC
- Power Supply: ANTEC PP-412X 400W

Software

- Intel chipset inf update 3.20.1008
- 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:

Augustine Chen (**ASUS**), Carol Chang (**ASUS**) and Sharon Tan (**BAS computers Netherlands**) for the ASUS P4T-E motherboard.
Jurgen Eijmberts (**Intel**) and Marieke Leenhouts (**MCS**) made sure we could test the 2.4 GHz Northwood Pentium 4.
Franz Pöller (ATI) sent us the ATI FireGL 8800.
Andrew Humber and Luciano Alibrandi sent us the Quadro 4 900 XGL
Kim Stowe and Terry Palek (**3DLabs**) sent us the 3DLabs Wildcat II 5000 and Wildcat III 6110.
Robert Pearce of **Corsair**, provided us with Corsair's **PC800 RDRAM-45**.

Drivers

Driver quality and speed is always important, but for an OpenGL accelerator it is nothing less than critical. Excellent OpenGL driver optimizations for specific workstation applications are an important way for ATI and Nvidia to distinguish between their desktop and workstation products. All three accelerators have a special panel where the workstation user can choose his or her primary OpenGL application. Once selected, the driver will then optimize the different OpenGL settings for that application.

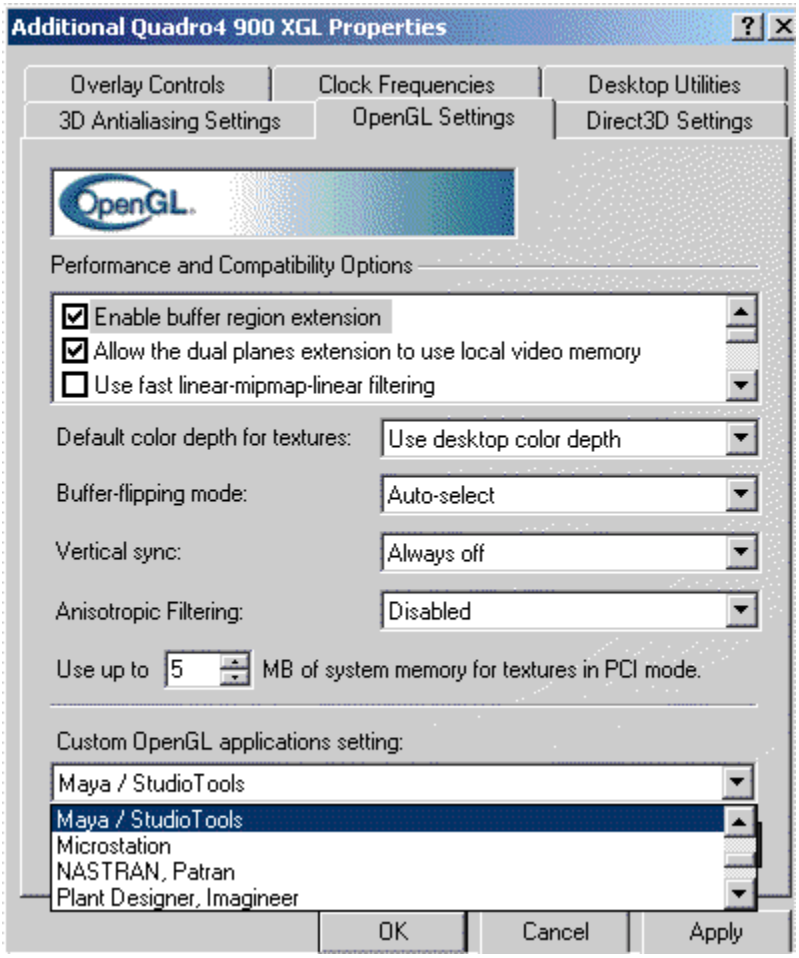

Nvidia clearly indicates which options are altered. 3DLabs and ATI go a small step further by allowing you to create a custom OpenGL profile.

In the case of ATI, however, it seems that optimizing the OpenGL API is limited to two options: enabling 8-bit double overlay planes and force copy swap.

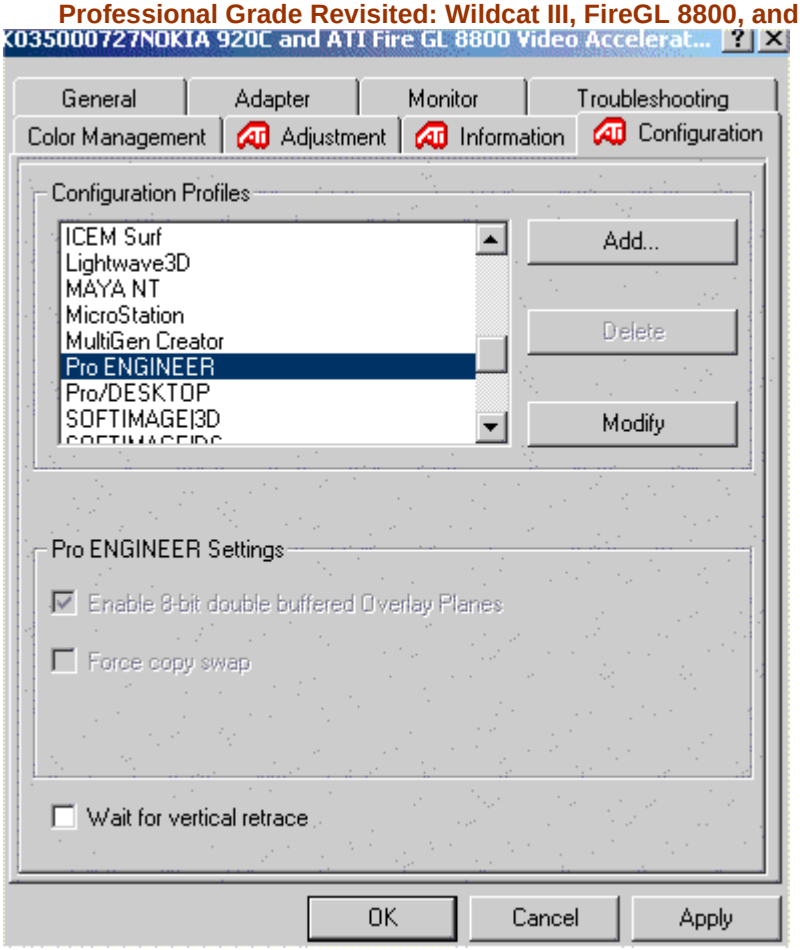

Anyway, while the right OpenGL settings can improve image quality slightly, we could not measure any significant performance boost from these optimized settings. Choosing the right setting for the application, which was benchmarked, increased performance modestly between 1% and 5%, regardless of which accelerator we used.

The real differences are made with specially optimized drivers for AutoCAD and 3DSMax. As Autodesk has 1.8 million users, it is no wonder that both 3DLabs and Nvidia feel it is important to offer a separate, specialized Heidi3D driver. ATI only offers a specialized driver for 3DSMax, which underlines our assertion that ATI is targeting the DCC market more than the CAD market with the Fire GL 8800.

3DLabs offers the anti-aliasing method with the best quality, 8/16x super sampling scene anti-aliasing. However this method needs a gigantic frame buffer as it samples each pixel 8 to 16 times. The 3DLabs 6110 is, with it's 64 MB framebuffer, limited to resolutions lower than 1280x1024 (up to 1152x864). At 1280x1024, we encountered the following dialog:

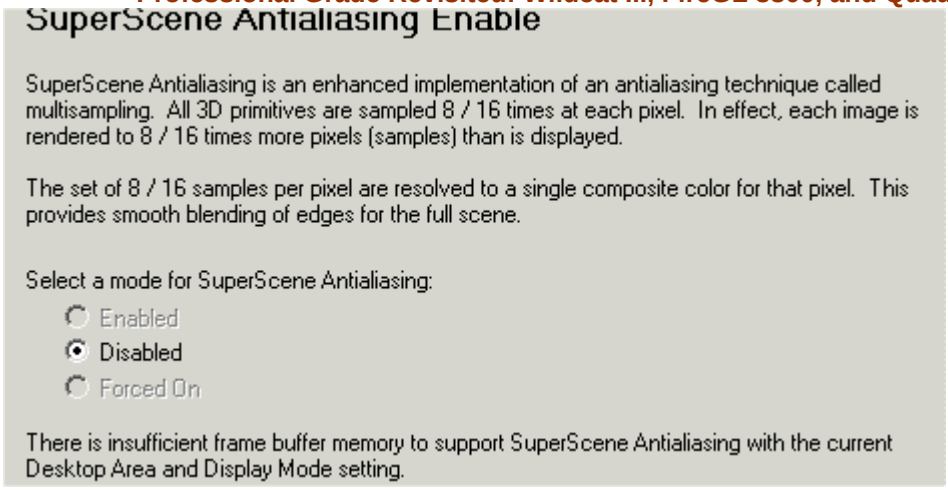

So in this particular case, the Wildcat III architecture is at a disadvantage. ATI and NVIDIA can use almost the whole 128 MB of onboard memory as a framebuffer and the 4 sample AA methods are still available at resolutions as high as 1600x1200. If what you need is the highest quality scene AA at high resolutions from 3DLabs, your only option is the Wildcat III 6210, which comes with a 128 MB framebuffer, a 256 MB texture buffer, and a staggering price tag of $2500. 3DLabs' top model offers 16 sample AA at resolutions up to 1920x1080.

However, we doubt most OpenGL users use such scene anti-aliasing frequently. At 1600x1200, the jaggies are small, and with complex scenes, speed and high quality line AA are still the primary concerns.

2D Performance

The powerful polygon beasts that we tested in the lab are built from the ground up to accelerate the 3D graphics pipeline. Nevertheless, they must have some decent 2D-acceleration too, as much work is still done in 2D. So how do they compare in 2D? We tested with ZD's High-End Graphics Winmark to find out. High-End Graphics always runs at 1024x768.

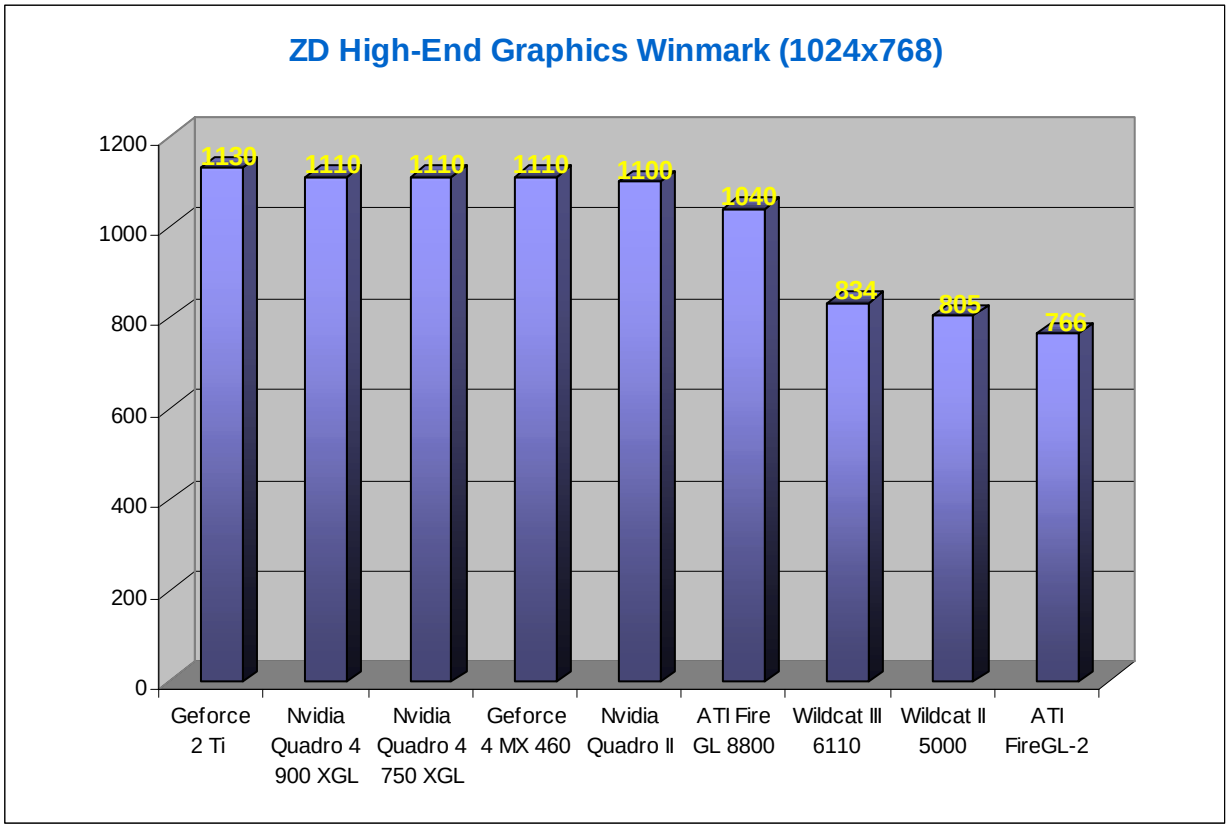

As you can see here, all the Nvidia cores perform more or less identically, with the FireGL 8800 trailing slightly behind them, followed by the Wildcat boards. These results are mirrored for the most part by the Photoshop and Microstation results reported on the following page.

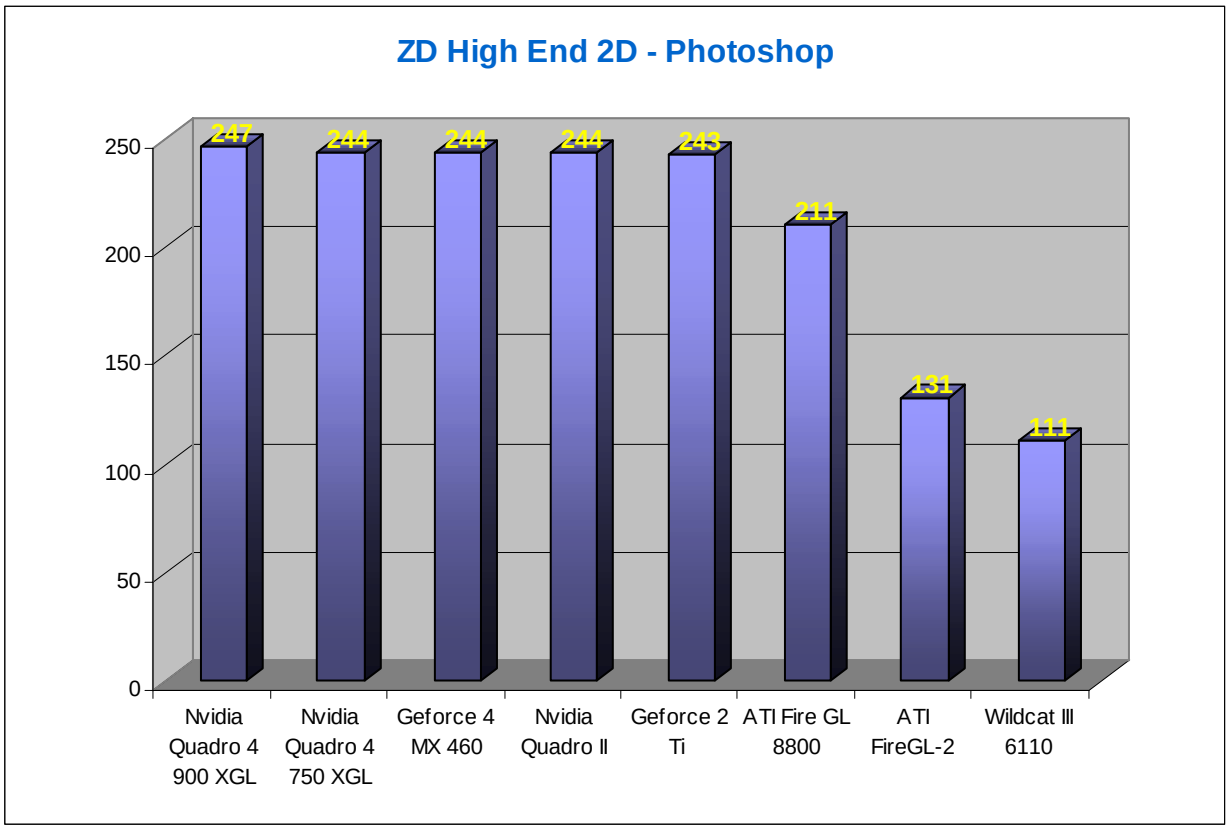

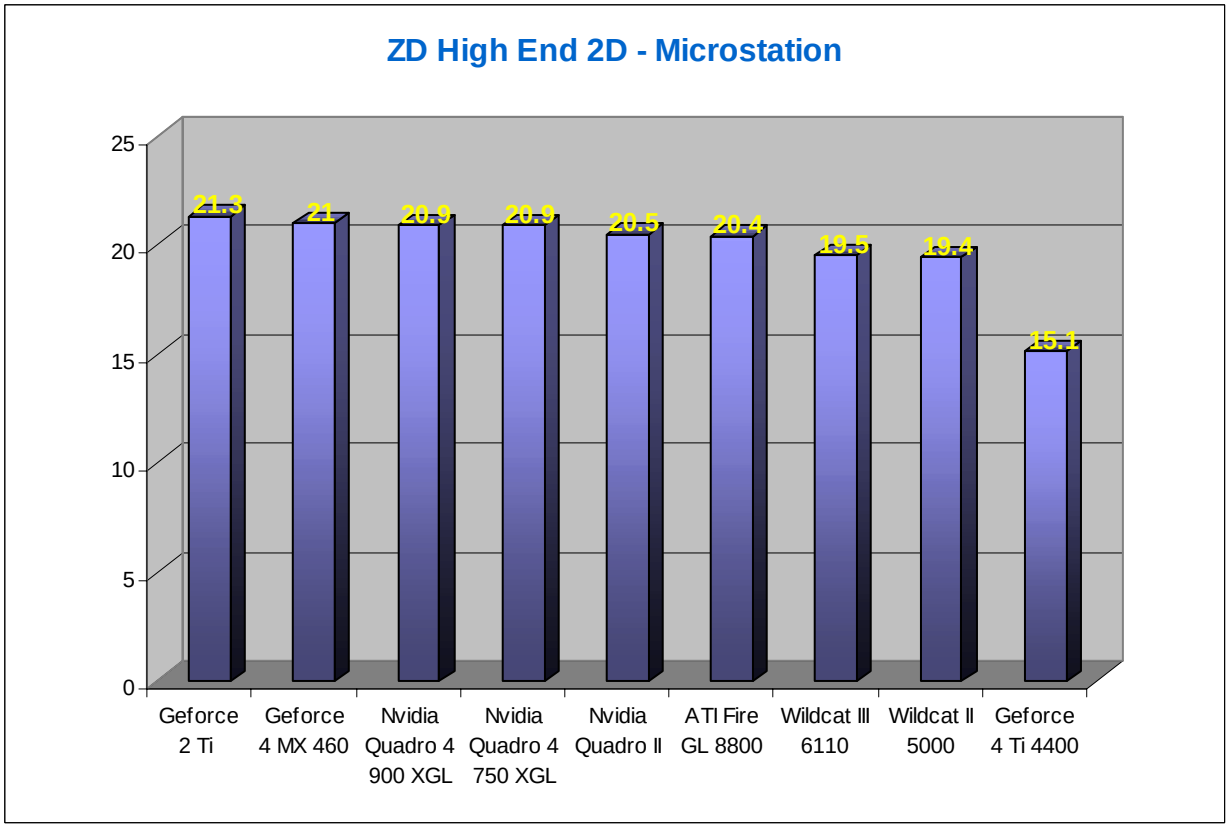

Nvidia has the strongest 2D Engine, followed closely by ATI. 3DLabs scores 33% slower. When we zoomed in on a few very complex 2D AutoCAD drawings, the Quadro 4's 2D engine proved to be superior and the difference in 2D speed was noticeable. 2D-AutoCAD runs about between 1.5 and 3 times faster on the Quadro 4 than on the 3DLabs Wildcat III. Zoom and pan functions, in particular, are much faster on the Quadro 4. For 2D CAD, the Quadro 4 is by far the champion in AutoCAD. This is no real surprise, however, as we found out about Nvidia's 2D performance awhile back in our **Matrox G550 review**.

Raw 3D Performance

Before investigating real world performance, we decided to take a look at the basic parameters that determine 3D performance: fillrate, geometry transformation and setup, and line AA wireframe speed. The results were surprising to say the least...

As typical CAD or 3D animation scenes have a great deal of geometric complexity, and limited use of multitexturing, fillrate is only a secondary concern for an OpenGL accelerator. According to ATI, a workstation user spends on average 50% of his or her time in wireframe mode, 25% in a gouraud shaded mode, and 20% in textured mode. The remaining 5% is spent playing back the final rendering result. Nevertheless, the fillrate must be sufficient enough to render pixels as fast as the polygons of each frame are transformed, lit and set up. Otherwise, fillrate becomes a bottleneck.

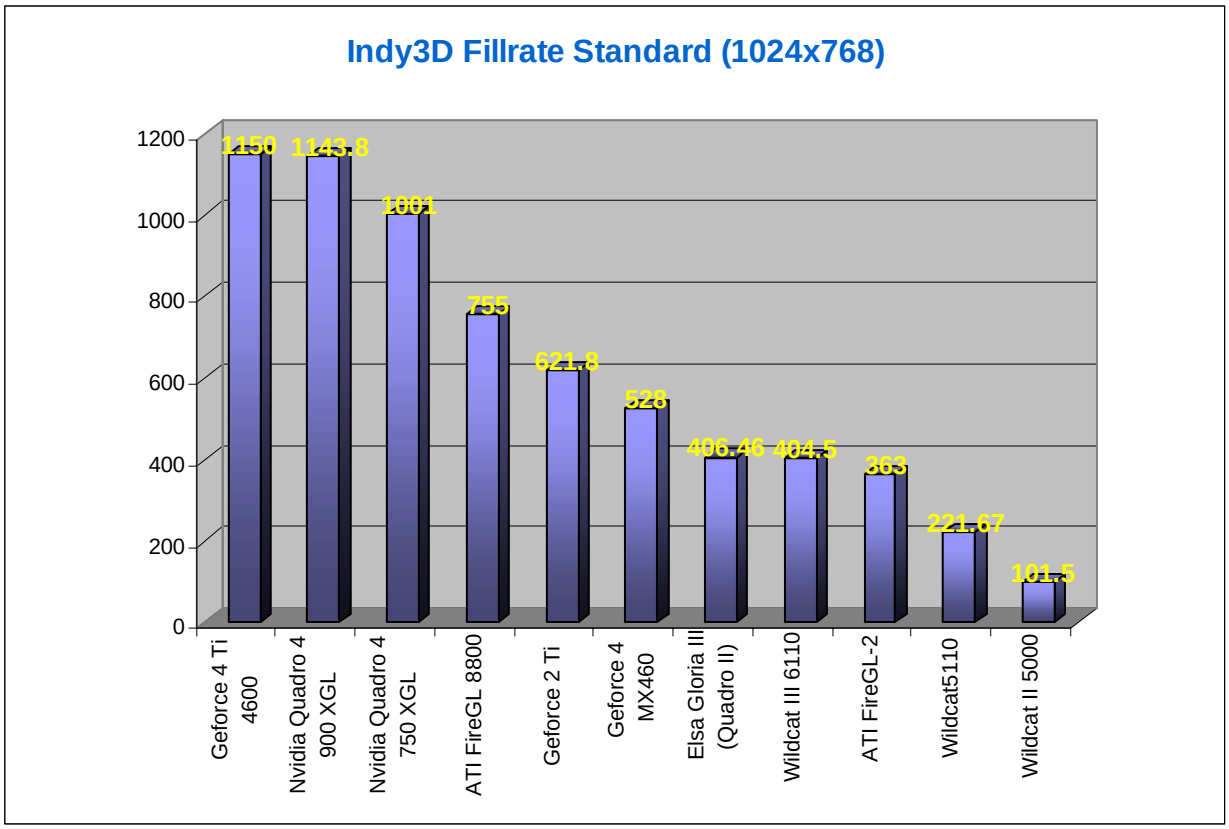

The GeForce 4 engine introduced a crossbar memory controller, Z-occlusion and lossless Z-buffer compression, and the result is that all GeForce 4-based cards are incredibly fillrate efficient. Theoretically, the Quadro 4 can deliver 1.1 billion trilinear filtered pixels per second, In practice, it is possible to deliver up to 1 billion: an efficiency rate of 90%. Compare this to its older brother, the Quadro 2, which is supposed to deliver 1 billion pixels, but which is only 40% efficient.

Thanks to hardware occlusion, an improved rendering engine, and DDR SDRAM memory, the Wildcat III is able to deliver twice the fillrate of it predecessors. This is an important improvement, as it makes the Wildcat III family more balanced, particularly in **textured mode** (Maya, 3DS Max) where the Wildcat II 5110 was limited by its low fillrate.

Our next benchmark is the standard Indy3D benchmark, a cross-platform industry standard benchmark for OpenGL and CAD. The CAD150 subtest rotates an engine model comprised of 150,000 polygons in wireframe mode and applies smooth shading to the model at the end of the test.

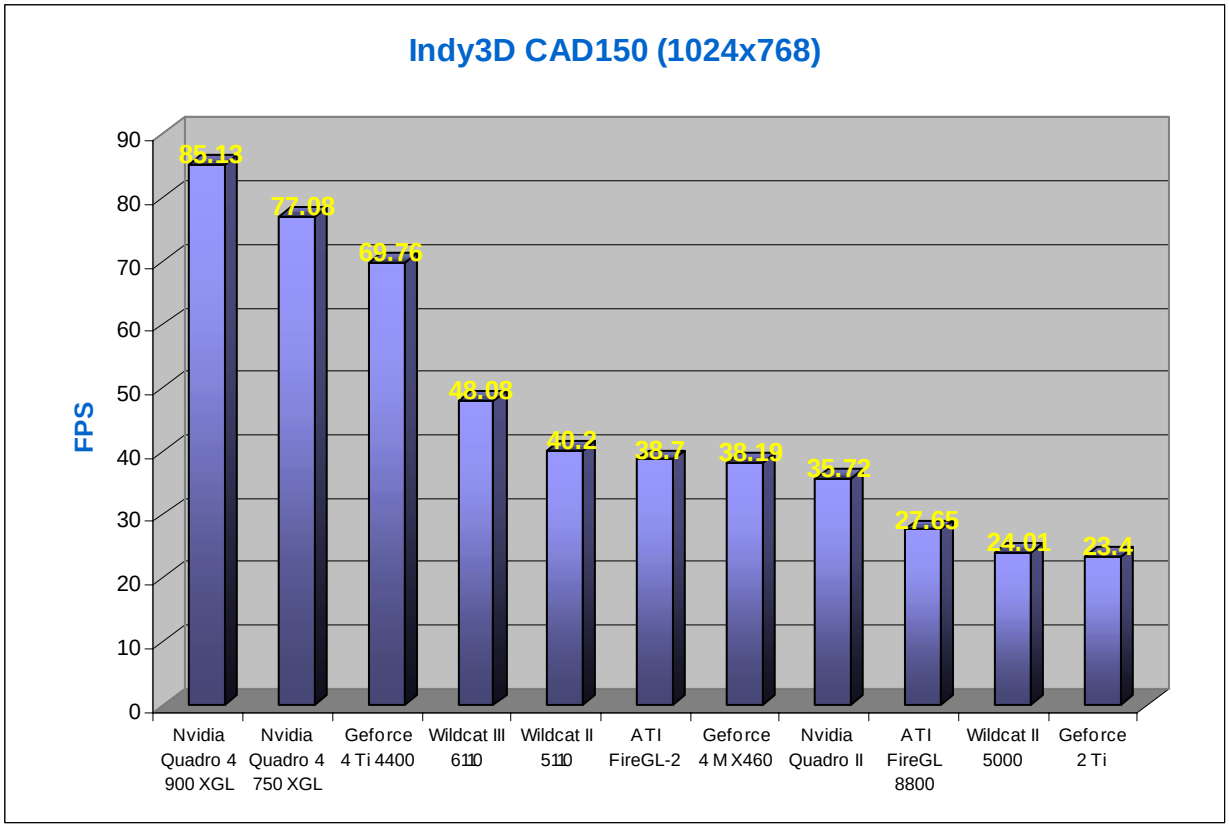

Nvidia simply crushes the competition in this test. As Indy3D is a somewhat synthetic test, however, we have to put the results in perspective somewhat. First of all, how much CPU power is used to achieve such performance?

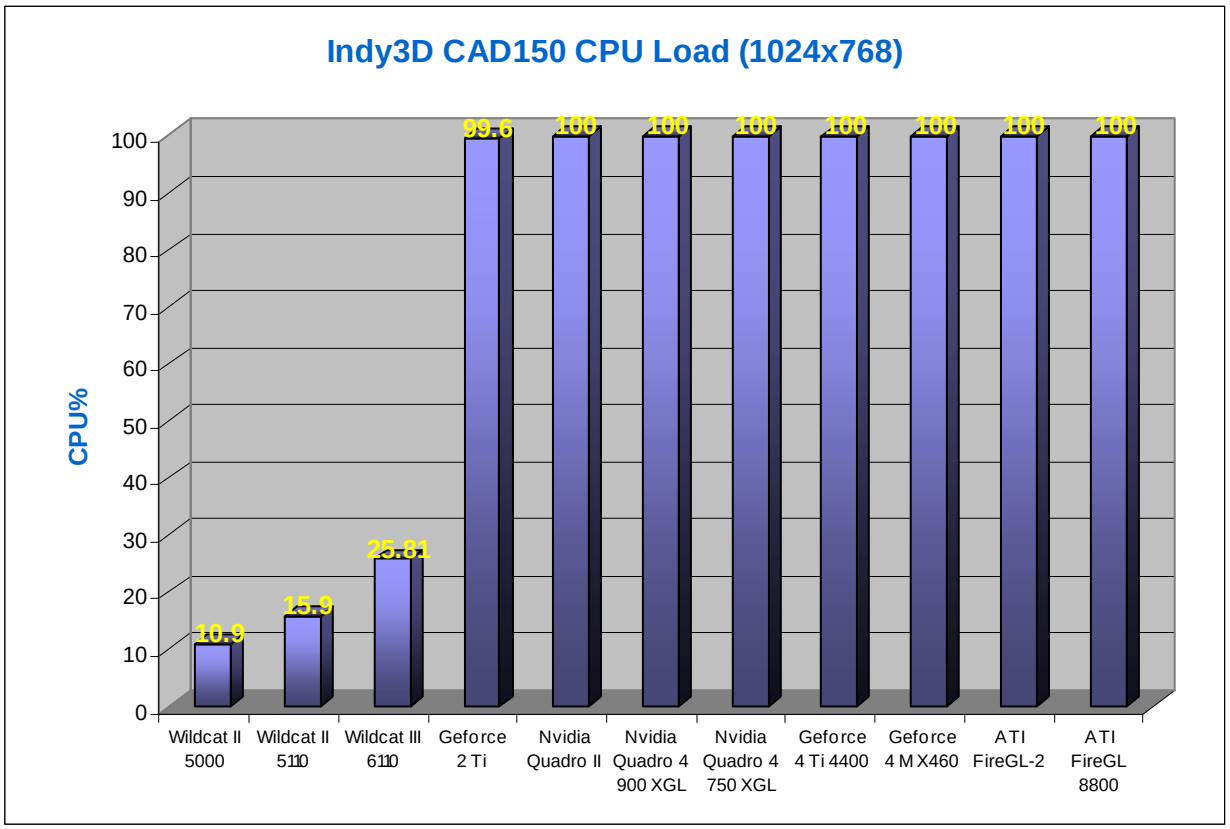

The Wildcats continue to be the OpenGL accelerators that offload the host CPU the most. While such ultra low CPU usage is not necessary for good performance, it is clear that the actual real world performance of the accelerators that fully use the available power will be somewhat lower. For example it is important for GUI thread of a CAD (or 3D animation) application always has CPU cycles available to it. Otherwise, it would be impossible to handle interactive 3D operations such as zooming (the zoom depends on how much you roll the wheel of the mouse), rotating, etc. Windows 2000 will assign resources to the GUI thread of the OpenGL application as necessary.

So what? It is indeed very rare that the GUI thread requires a lot of CPU power. So why do we still have some reservations with the Indy3D results? Well, when we watched the benchmark closely, we saw that the wireframe part ran more or less at the same speed - roughly estimated - on all accelerators, while the smooth shaded part ran ballistically fast on all Nvidia cards. In fact, all Quadros and GeForces ran the second part incredibly fast, at framerates up to 113 fps (indicated by the gauges)! It seems that the second part was, as such, fillrate limited, giving the Quadro cards a very considerable advantage over the other OpenGL cards like the ATI Fire GL-2 and the Wildcat family. No problem there, but our AutoCAD 3D rotate scripts contradicted that a 100k+ polygon model could run at such high framerates in shaded mode, even with flat shading.

So, we had some doubts about how representative the standard Indy3D benchmark was. Our primary goal of running Indy3D was determining the raw 3D power of a certain OpenGL accelerator, so we decided to make the test more real world. We tested at 1600x1200, in wireframe mode, once with line AA enabled and once with line AA disabled.

Line anti-aliasing techniques minimize the jaggedness of the wireframes drawn. Line AA is very important as CAD and 3D animation professionals spend a lot of time in wireframe mode and line anti-aliasing improves productivity. Without line AA, it is not always clear which line is in front of the other, and the whole design looks like spaghetti. Look at the pictures below for a comparison:

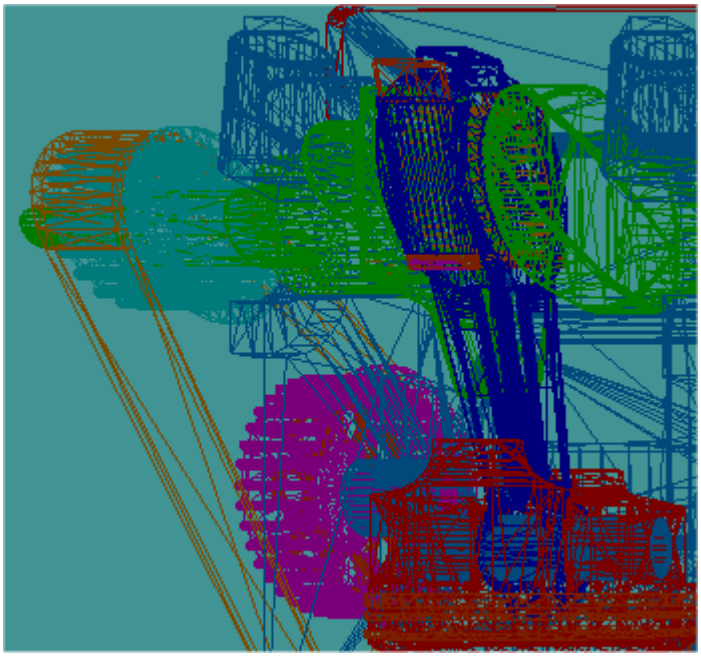

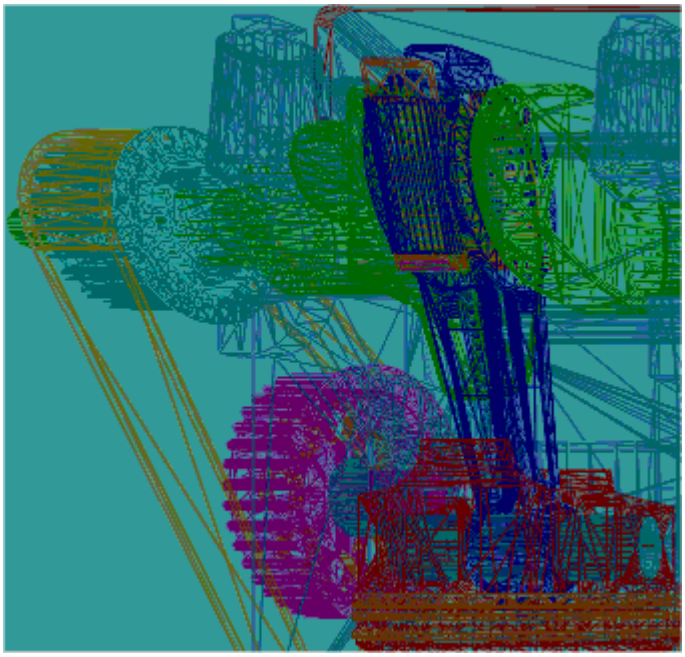

In the first picture, it is pretty hard to make out the wheel at the end of the camshaft. In the second picture, which uses line AA, you can clearly see the gearwheel at the end of camshaft. Also notice that many lines seems to touch or even entangle each other, while parallel lines are much more distinctly visible in the line AA picture. The second advantage of line AA **can be seen here**: by smoothing the jagged edges of lines, line AA eliminates the flickering and crawling artifacts that appear when you rotate a model.

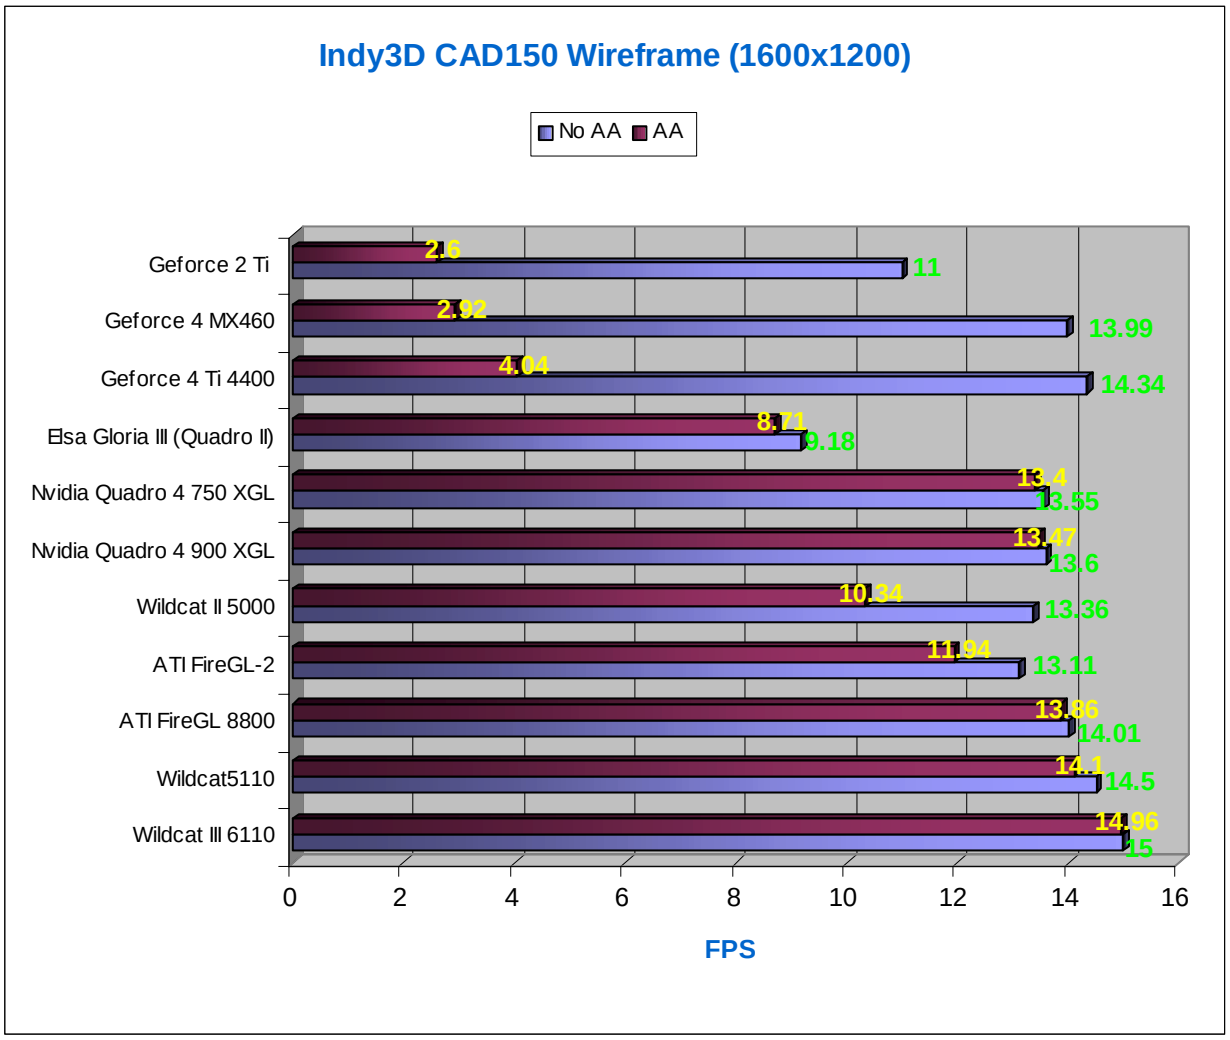

This Indy3D benchmark proves to be a totally different story. The gamers' 3D accelerators can compete, as long as we do not enable line AA. Once line AA is enabled, even the 300 MHz quad-pipeline GeForce 4 Ti4400 is defeated by the 200 MHz single-pipeline Wildcat II 5000. This is where the OpenGL accelerators feel at home. The Quadro 4 does not suffer from the terrible performance loss like the gaming cards, however. We'll see further that Indy3D's benchmark gives us some indication of real world performance, but not a very accurate one.

AutoCAD

AutoCAD is by far the most popular CAD program in use. However, support for 3D has only been present since the release of AutoCAD 2000, so it is not the most advanced 3D CAD package on earth. To enable OpenGL acceleration you must set the "Acceleration" option to "Hardware" in the "System" tab of the options dialog. In case of Nvidia and 3DLabs, a special driver for the Heidi® 3D graphics pipeline of AutoCAD is provided.

In case of 3DLabs' driver, only AutoCAD 2000 can be supplied automatically with the specialized Heidi® 3D driver. The dialog box does not have an option for AutoCAD 2002 and if you use the automatic driver installation of AutoCAD 2000, the AutoCAD 2002 incompatible i3dhdd6.hdi driver will be installed.

As we used AutoCAD 2002, we had to install the "i3dhdd7.hdi" driver manually into the "drv" directory. A trouble report has been filed with 3DLabs, and the installer for the Heidi® 3D driver will be updated to support AutoCAD 2002.

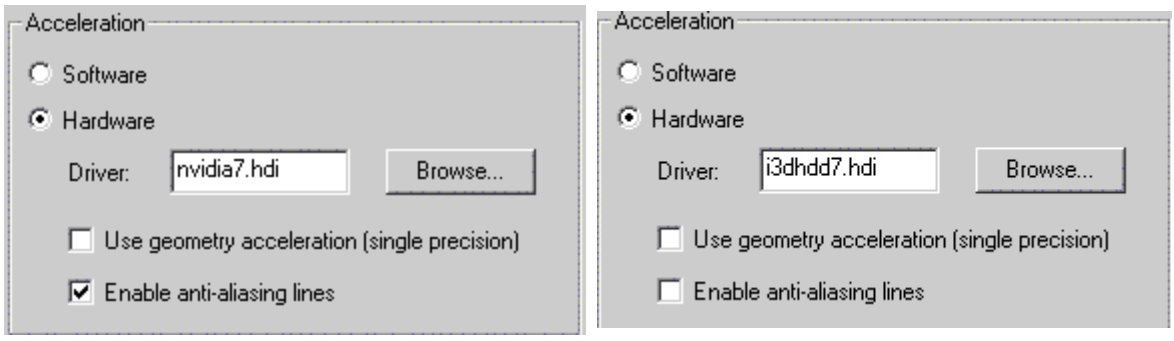

ATI does not supply a special Heidi® 3D driver, and therefore we used the standard WOpenGL.hdi driver.

To benchmark AutoCAD we used an AutoLISP script which uses the, Rotate3D, Zoom3D, and shademode commands to perform zooms and rotations in wireframe and gouraud-shading modes. We applied these commands on several of the included sample objects. For example, we tested with the "robotsolid" object, but copied the object 5 times to get a more complex model. We added up all results and normalized them based on the results of the Nvidia Quadro 2. This way, we are able to see the performance deltas immediately.

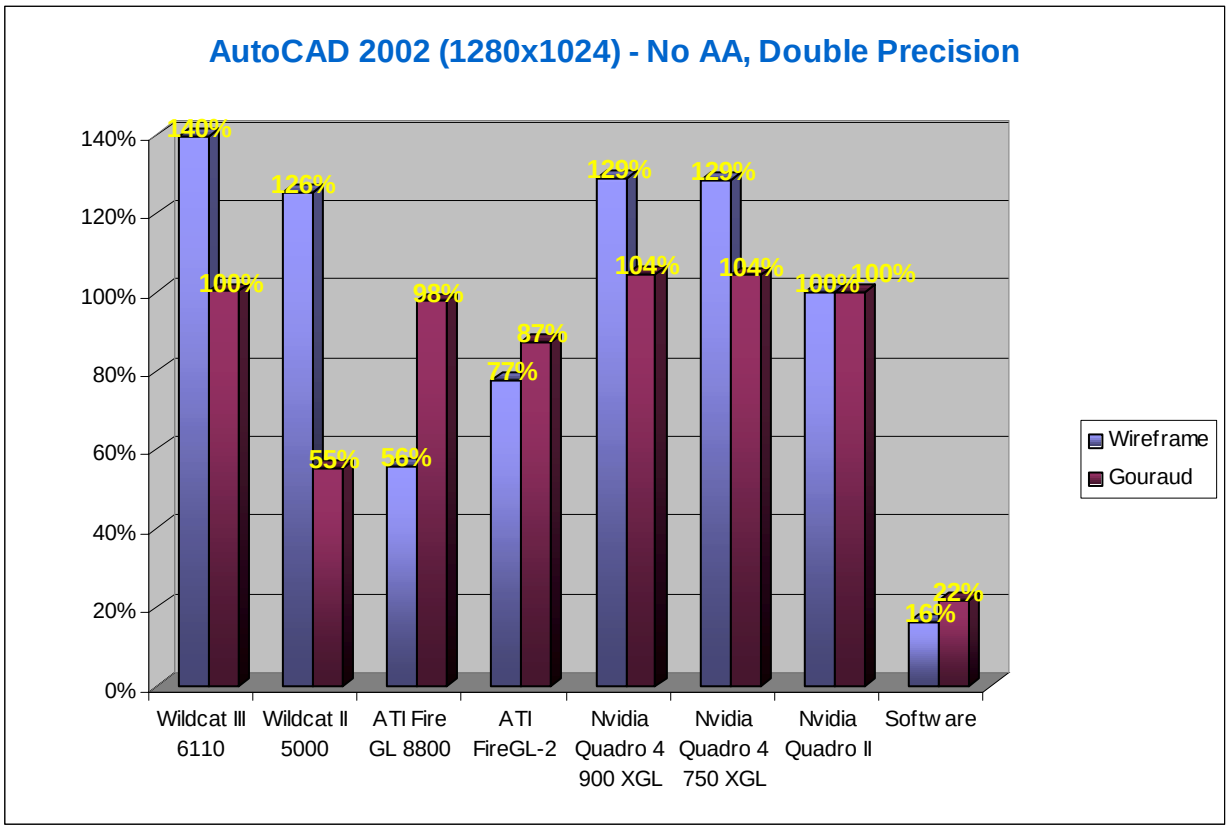

The Wildcats show their claws in wireframe mode. The "old" Wildcat II 5000 stays close to the performance levels of the newest Nvidia card, while the Wildcat III 6110 is clearly the winner of this test.

In gouraud mode, however, the differences are not so clear cut: the FireGL 8800, Wildcat III 6110, and all the Quadros perform relatively similar to one another. The Quadro 4 wins by a small margin when gouraud shading is enabled. The FireGL 8800 disappoints in wireframe and clearly lacks serious polygon power, while the Wildcat II 5000 is totally the opposite. The limited fillrate of the Wildcat II makes it a mediocre gouraud shading engine.

As we have outlined a few moments ago, an OpenGL card should offer excellent line AA performance. It must be said that not one of the competitors lost more than 3% of performance when enabling Line AA.

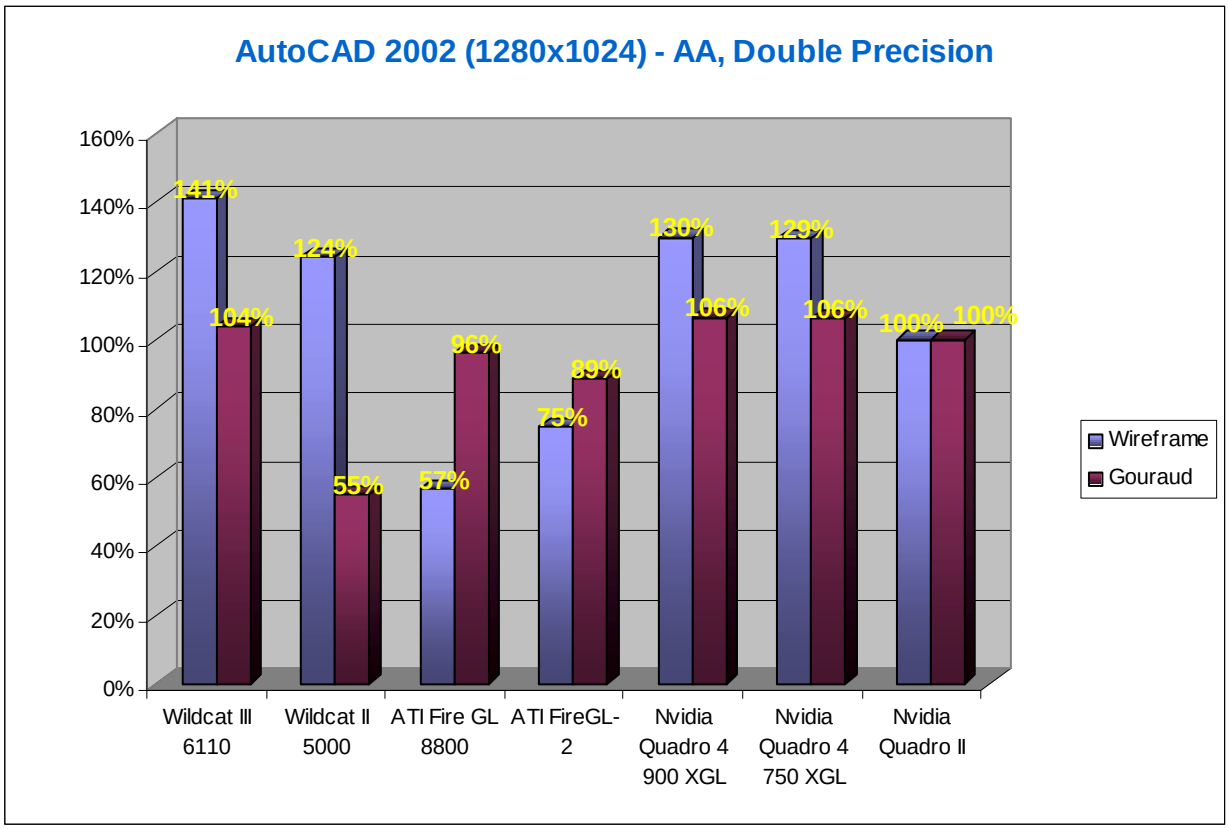

The gap between the Quadro 4 and Quadro 2 widens, but for the rest, the results stay the more or less the same.

High-End 3D CAD

We used the ProE-01 viewset of SPEC ViewPerf, which was created from traces of the graphics workload generated by the Pro/ENGINEER 2001 application from PTC. The most important model in this test, the PTC World Car, is composed of no less than 3.9 to 5.9 million vertices! This model is measured in shaded, hidden-line removal, and wireframe modes.

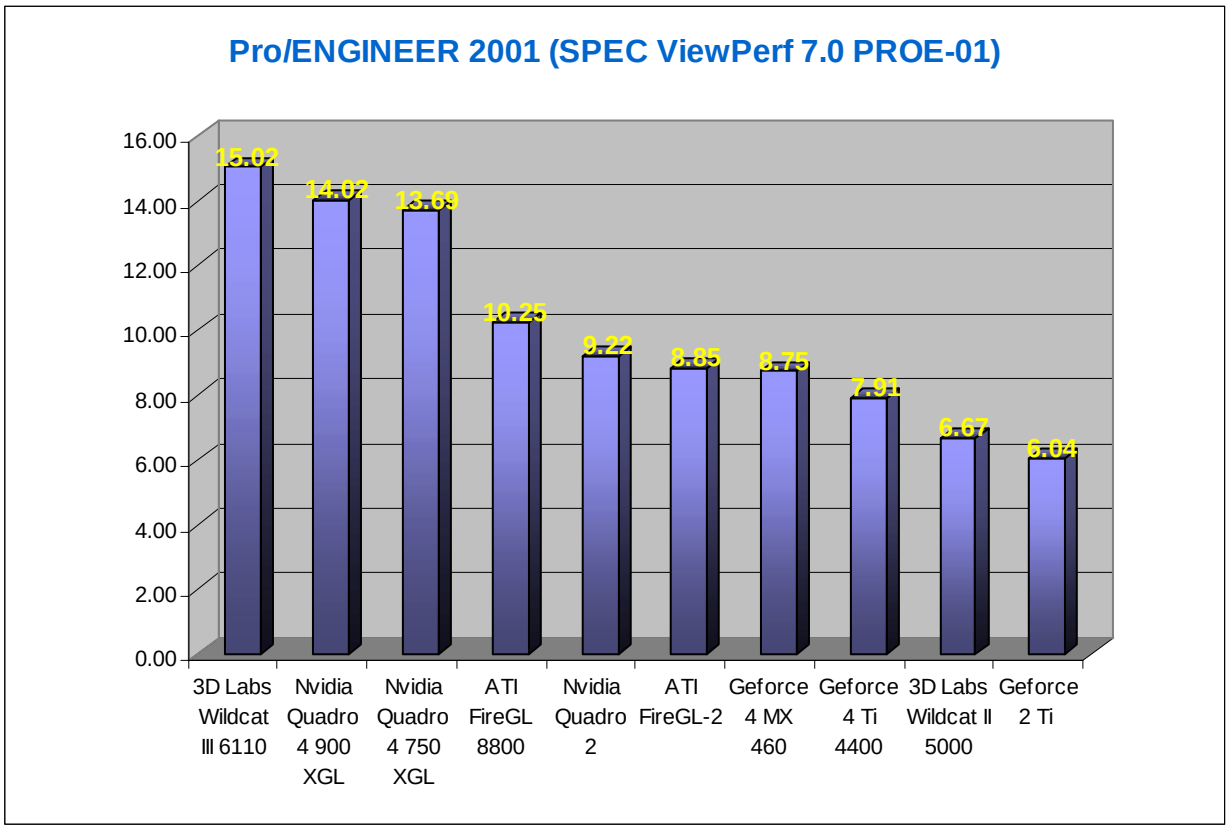

The Wildcat III comes out on top, but the days when Nvidia's GPUs were being swept away are over. Indeed, the Quadro 4 900 XGL is getting close to the Wildcat III 6110 in this test. The performance difference between modes was interesting as well.

Pro-E Viewset	3DLabs Wildcat III 6110	Quadro 4 900 XGL	ATI Fire GL 8800
PTC WorldCar shaded	5.37	7.18	4.16
PTC WorldCar shaded to show model detail	6.31	8.67	4.69
PTC WorldCar using hidden line removal mode	7.1	10.3	5.77
PTC WorldCar using wireframe mode	25.7	14.2	18.7
PTC WorldCar using anti-aliased lines	21.3	9.04	18.7

The Quadro 4 is superior in the shaded modes, but the Wildcat III wipes the floor with Nvidia flagship in line anti-aliasing modes. In our humble opinion, the wireframe and anti-aliased modes are more important than shaded modes. After all, most CAD professionals still spend most of their time there. Therefore, we may say that the Wildcat III is better than the weighted averages of SPEC ViewPerf indicate. Does the Wildcat III still rule in the high-end CAD market? We decided to find out with the UGS viewset.

The ugs-01 viewset was created from traces of the graphics workload generated by Unigraphics V17. The engine model used was taken from the SPECapc for Unigraphics benchmark. Three rendering modes were measured: wireframe, shaded, and shaded with transparency. The wireframe workloads are measured both in normal and anti-alised mode. All tests were repeated twice, once rotating in the center of the screen and then moving about the frame to measure clipping performance.

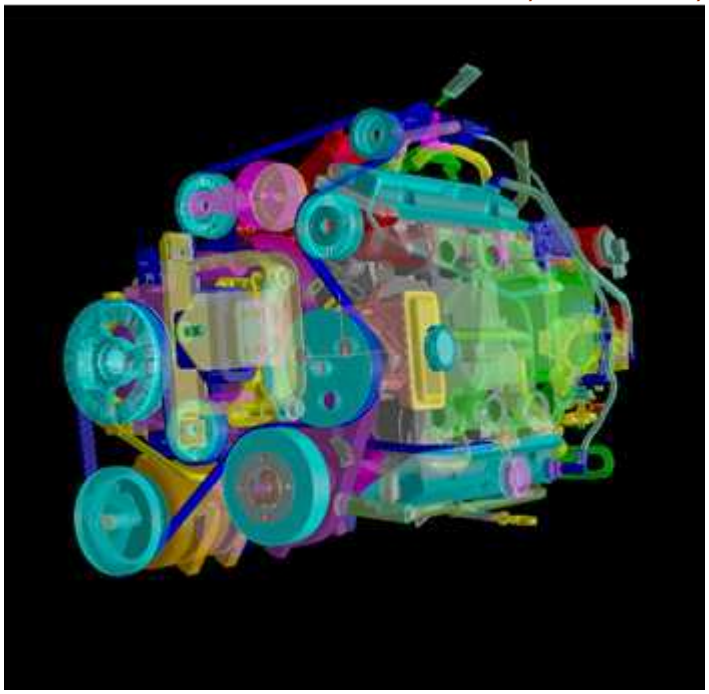

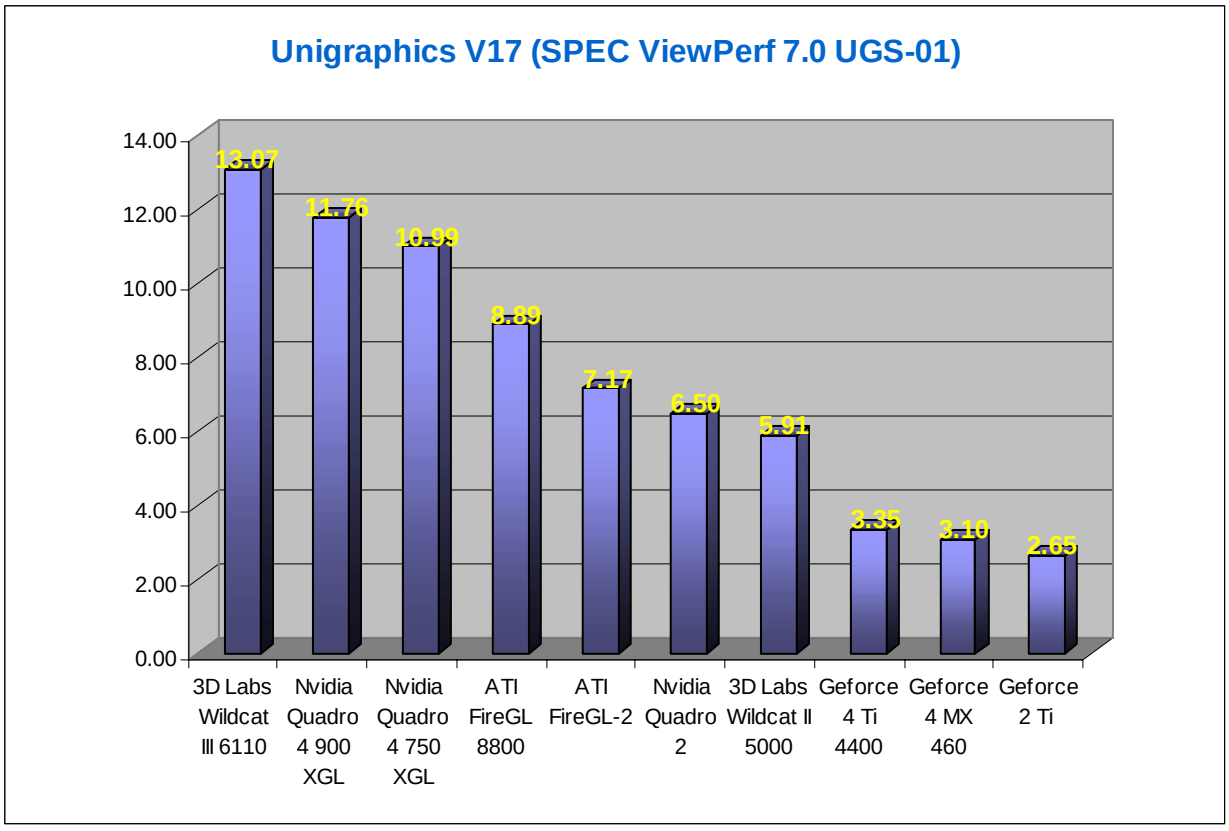

Again, the Wildcat III takes the lead. The difference between professional and gamers' cards here is substantial in this high-end CAD package. To be complete, we finish this CAD chapter with a more specialized 3D CAD package: DesignReview (DRV-08).

DesignReview is a 3D computer model review package specifically tailored for plant design models consisting of piping, equipment and structural elements such as I-beams, HVAC ducting, and electrical raceways. It allows flexible viewing and manipulation of the model for helping the design team visually track progress, identify interferences, locate components, and facilitate project approvals by presenting clear presentations that technical and non-technical audiences can understand.

On the construction site, DesignReview can display construction status and sequencing through vivid graphics that complement blueprints. After construction is complete, DesignReview continues as a valuable tool for planning retrofits and maintenance. DesignReview is a multi-threaded application that is available for both UNIX and Windows NT.

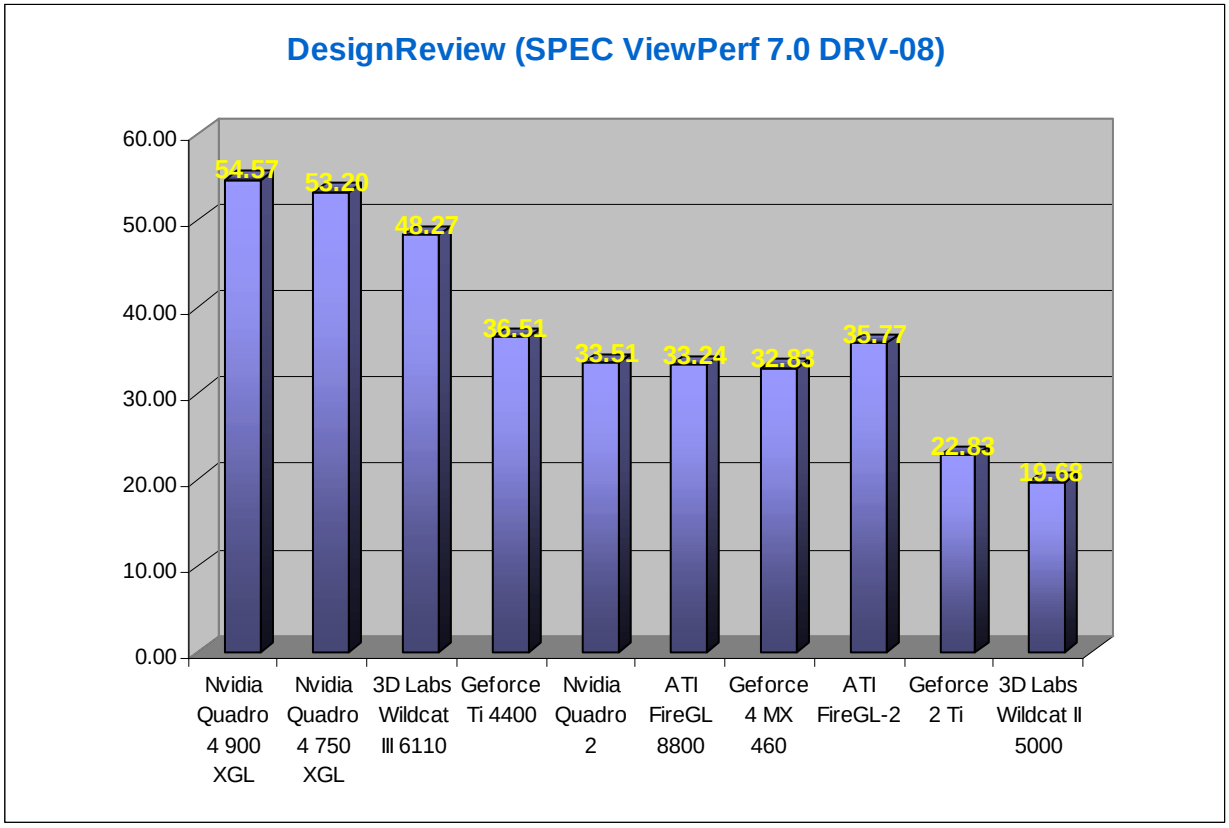

The emphasis of this test is on smooth shading, and the result is a landslide victory for the Quadro 4's super shading engine.

Scientific Visualization

From the SPEC ViewPerf description page:

> IBM's Visualization Data Explorer (DX) is a general-purpose software package for scientific data visualization and analysis. It employs a data-flow driven client-server execution model and is currently available on Unix workstations from Silicon Graphics, IBM, Sun and Hewlett-Packard. The OpenGL port of Data Explorer was completed with the release of DX 2.1.
>
> The tests visualize a set of particle traces through a vector flow field. The width of each tube represents the magnitude of the velocity vector at that location. Data such as this might result from simulations of fluid flow through a constriction. The object represented contains about 3,000 triangle meshes containing approximately 100 vertices each. This is a medium-sized data set for DX.

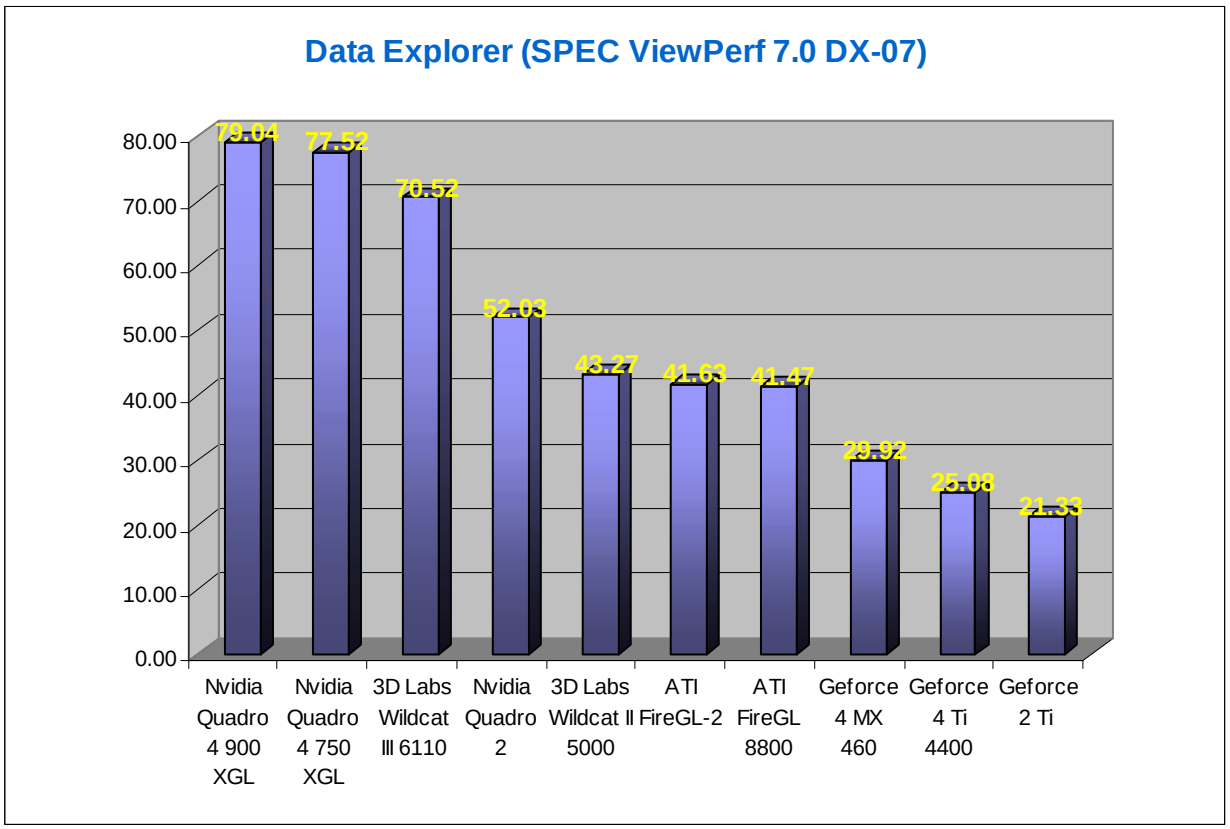

We are not experts in scientific visualization, but the Quadro 4 wins hands down. The low performance of ATI's solutions is rather remarkable, as the Wildcat II 5000 beats them.

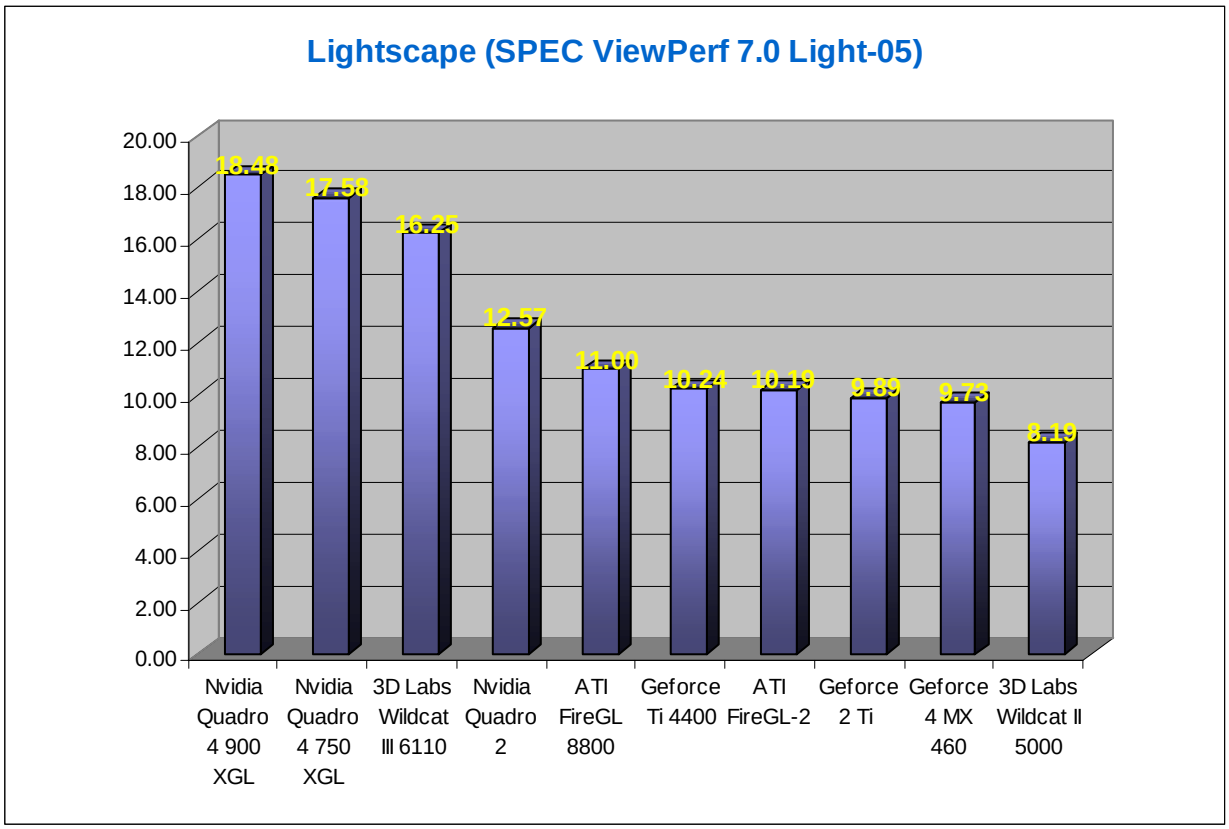

In the case of ViewPerf's Light-05 test, both the Quadro 4 900 XGL and 750 XGL lead the pack, with the Wildcat III 6110 bringing up third place.

Digital Content Creation: Maya 4

Typical Digital Content Creation (DCC) packages such as 3D Studio Max and Maya tend to stress OpenGL accelerators somewhat differently than CAD applications. Few CAD users will use a textured mode, while in DCC software, good acceleration of textured previews is a must. Most CAD packages use gouraud shading, while DCC software uses more complex forms of shading like smooth Blinn shading (such as our Maya 4 benchmark) or phong shading.

Maya, a creation of Alias|Wavefront is a must for any workstation review. The package has been used to help produce major blockbuster films like Lord of the Rings, Harry Potter, and Final Fantasy: The Spirits Within, to name a few. Many best selling games on the PS2 console platform, like Gran Turismo 3: A-Spec and Tekken Tag Tournament, have been created with the help of Maya as well. The days when Maya was a high-end package only available to the richest game and film studios are over. The complete version of Maya is now available for around $2000 and a crippled, but still powerful version, "**Maya Personal Learning Edition**," is available for free.

A list of qualified hardware is available **here**.

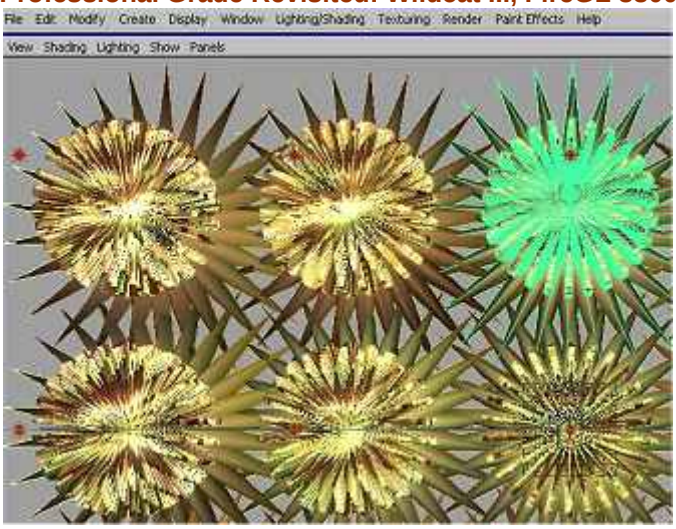

To test Maya 4, we used the 3D test from the **Maya Test Center**. We also wanted to evaluate the influence of the monitor's resolution, so we tested at 1024x768, which is considered as the absolute minimum workable resolution, and 1600x1200, the typical resolution for DCC work.

Maya 4 1024x768	Wireframe	Shaded	Shaded + light	Textured	Textured + light
Wildcat III 6110	**92.2**	**30.1**	9.8	**23.7**	11.7
GeForce 4 Ti 4400	65.9	27.1	6.4	17.3	6.4
Nvidia Quadro 4 900 XGL	62.2	25.9	7	16.7	7
Nvidia Quadro 4 750 XGL	62	25.9	6.4	16.5	6.4
ATI FireGL-2	44.3	21.1	**18.9**	12.6	**12.5**
GeForce MX460	43.6	15	10.3	13.7	10.3
GeForce 2 Ti	49.2	20.5	3.1	16.5	3.1
Elsa Gloria III (Quadro 2)	42.8	18.1	8.4	13.4	8.4
ATI Fire GL 8800	37.5	18.8	9.4	15	9.4
Wildcat II 5000	42	10	5.2	9	3.9

The Maya benchmark proves what the predictions of the Indy3D benchmark: the Wildcat III's six T&L engines offer serious polygon power. ATI seems to master lighting very well. While the Quadro's performance implodes because it is not able to cope with the l5 lights used in this scene, ATI's Fire GL 8800 is capable of offering frames rates which are no less than 34% higher than the chips of arch rival Nvidia.

The FireGL supports 16 hardware lights and comes with a super lighting engine but is not able to handle textures very well.

One look at the performance data at 1600x1200 and we'll see that the tables have turned.

Maya 4 1600x1200	Wireframe	Shaded	Shaded + light	Textured	Textured + light
Nvidia Quadro 4 900 XGL	**61.4**	**24.4**	6.6	**16.7**	6.6
Wildcat III 6110	56.5	17.3	9.8	16.6	11.7
GeForce 4 Ti 4400	51.9	22.3	6	16.5	6
Nvidia Quadro 4 750 XGL	51	21.4	6	16.5	6

GeForce 2 Ti	31.8	13.6	3.1	16.5	3.1
GeForce MX460	23.9	13.6	9	12.2	9
ATI FireGL 8800	25	12.5	8.3	12.5	8.3
Elsa Gloria III (Quadro 2)	22.2	11.2	7.3	10.8	7.3
ATI FireGL-2	17.3	8.2	8	5.1	5.1
Wildcat II 5000	18.2	3.7	3.6	2.8	3.1

The support for only 8 hardware-accelerated lights remains a weakness of typical gaming 3D accelerators. Nevertheless, Nvidia performs very well at higher resolutions thanks to its powerful rendering engine and 10 GB/s memory interface. While performance on the Quadro 4 cards hardly decreases at higher resolutions, the Wildcat's shading and wireframe performance takes a dive.

The Wildcat II 5000 performs poorly, and so will its brother, the 5110. The Wildcat II 5110 can outperform the Wildcat II 5000 by about 40 to 75 percent (deduced from **here**), a performance level which is still inadequate to compete with the newest generation of DCC OpenGL cards. On a positive note, it is clear that the Wildcat III is far superior to the Wildcat II when it comes to texturing and shading. Based on our previous assumptions that the Wildcat II 5110 performs about 40% faster in texturing than the Wildcat II 5000, we can safely say that the Wildcat III 6110 handles textures about 4 times faster than the Wildcat II 5110. If you work frequently with more than 8 lights, the Wildcat III is the most balanced solution.

The low-priced ATI FireGL 8800 could use some triangle power, but it offers decent performance in all circumstances. The FireGL-2, on ther other hand, comes with a strong lighting engine, but performs very poorly at higher resolutions.

It is also interesting to note that despite the fact that it runs at lower clockspeeds (250 MHz versus 300 MHz), the Quadro 2 is able to beat the Quadro 4 900 XGL in the "shaded + light" test. The most likely explanation for this is that the vertex shader of the Quadro 4 has to emulate a "static" lighting engine when calculating lights in Maya.

3D Studio Max 4.26

No need to introduce 3DSMax, as it is by far the most popular 3D modeling and animation package on Earth. Our first test comes once again from the SPEC ViewPerf 7.0 benchmark suite.

Each model was measured with two different lighting models to reflect a range of potential 3ds max users. The high-complexity model uses five to seven positional lights as defined by the SPECapc benchmark and reflects how a high-end user would work with 3ds max. The medium-complexity lighting models uses two positional lights, a more common lighting environment.

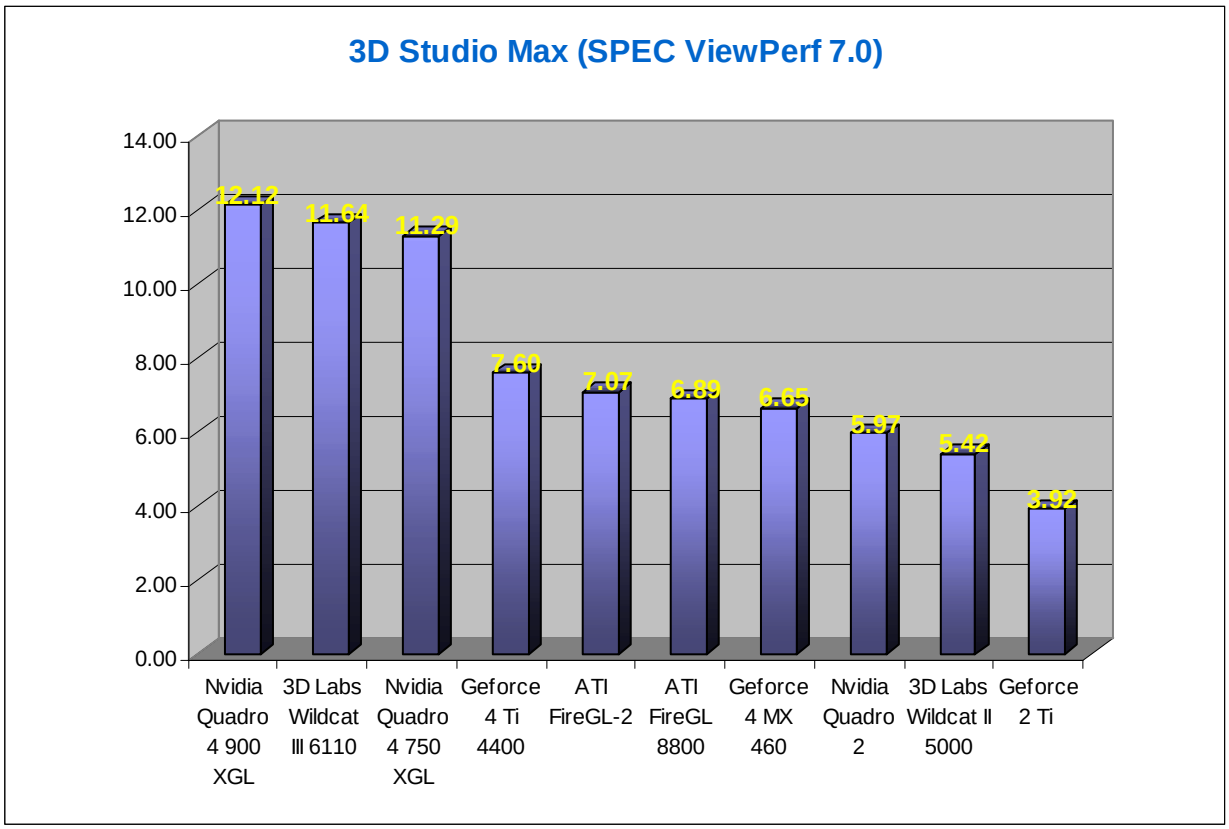

Building further on the Elsa Gloria DCC ("Quadro 3"), which was already targeted towards 3DSMax users, the Quadro 4 900 XGL is the fastest 3DSMax accelerator today. It is important to consider that the SPEC ViewPerf test uses the OpenGL plug-in driver from Discreet for benchmarking. As the Wildcat III 6110, the FireGL 8800, and Quadro 4 all ship with specially optimized OpenGL drivers, the SPEC test only gives an idea of the "raw 3DSMax performance."

Contrary to the screenshot below, all optimized drivers were set to Quality and anti-aliased lines were always enabled. Below you can see the optimized driver settings for the FireGL 8800:

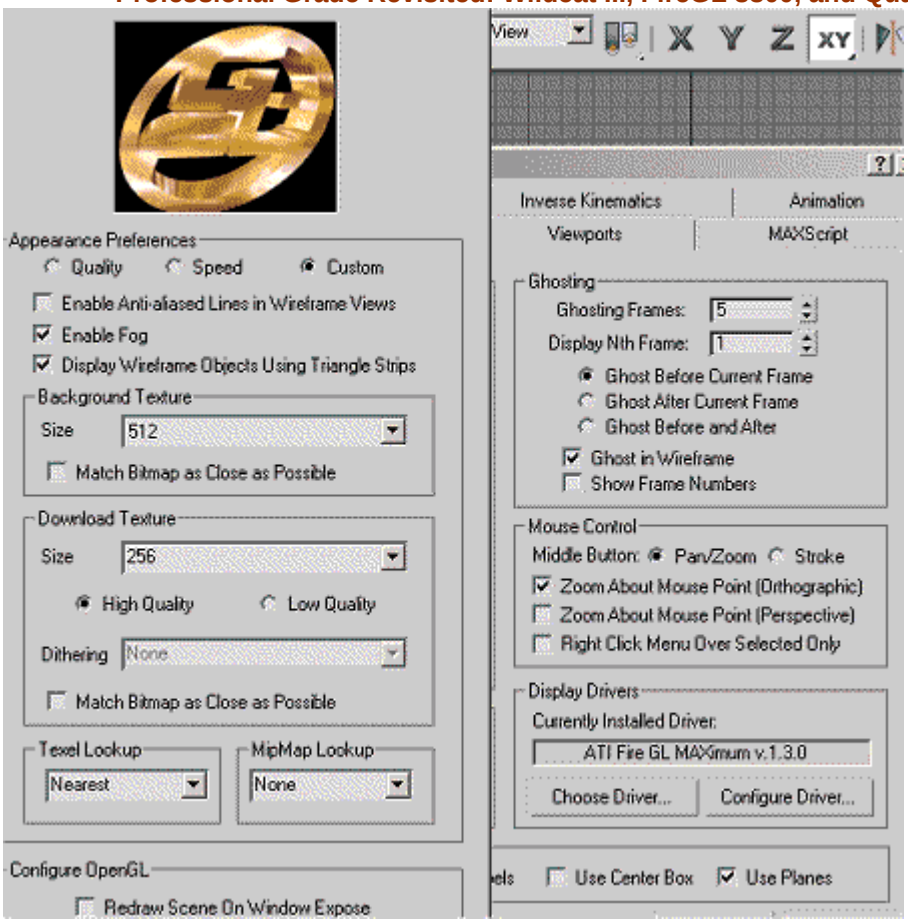

And the optimized driver settings for the Quadro 4:

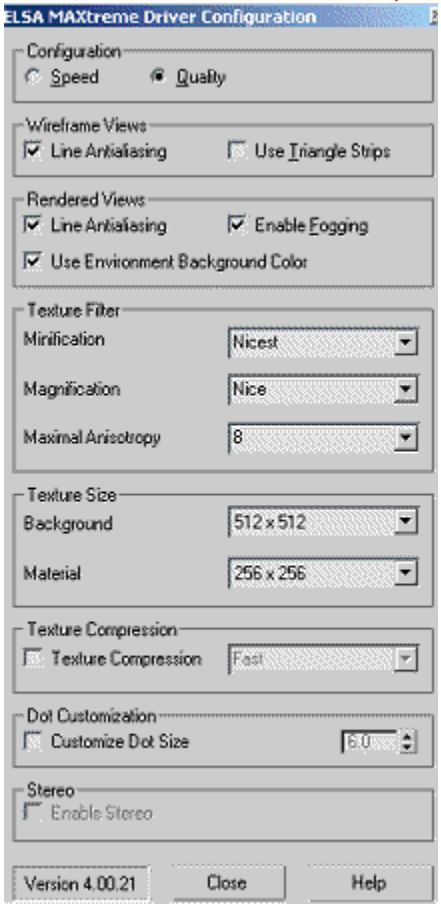

We started with SPECapc 3DSMax 4.26 benchmark, but we came to the conclusion that the benchmark could not produce reliable numbers on our platform. The reason is that the SPECapc test uses up to 1.1 GB of RAM and we were not able to outfit our system with more RAM at the time of the testing. So, if you see SPECapc 3DSMax 4.26 benchmarks on a system with less than 1 GB of RAM, please consider that there may be some issues with swapping that may affect those results. At certain points, our system started swapping like mad and the test became more a harddisk benchmark than anything else. The error margin on the scores between different tests was 40% and, in some cases, even higher.

We are trying to upgrade our testbed, but for now we decided to run some "manual" real world benchmarks. We measured how long it takes - in the different modes - to preview the 100 moving frames of the architecture scene in 3D Studio Max 4.26. The test was run in two different modes, the first using smooth shading and highlights, with the second rendering in wireframe mode with line anti-aliasing.

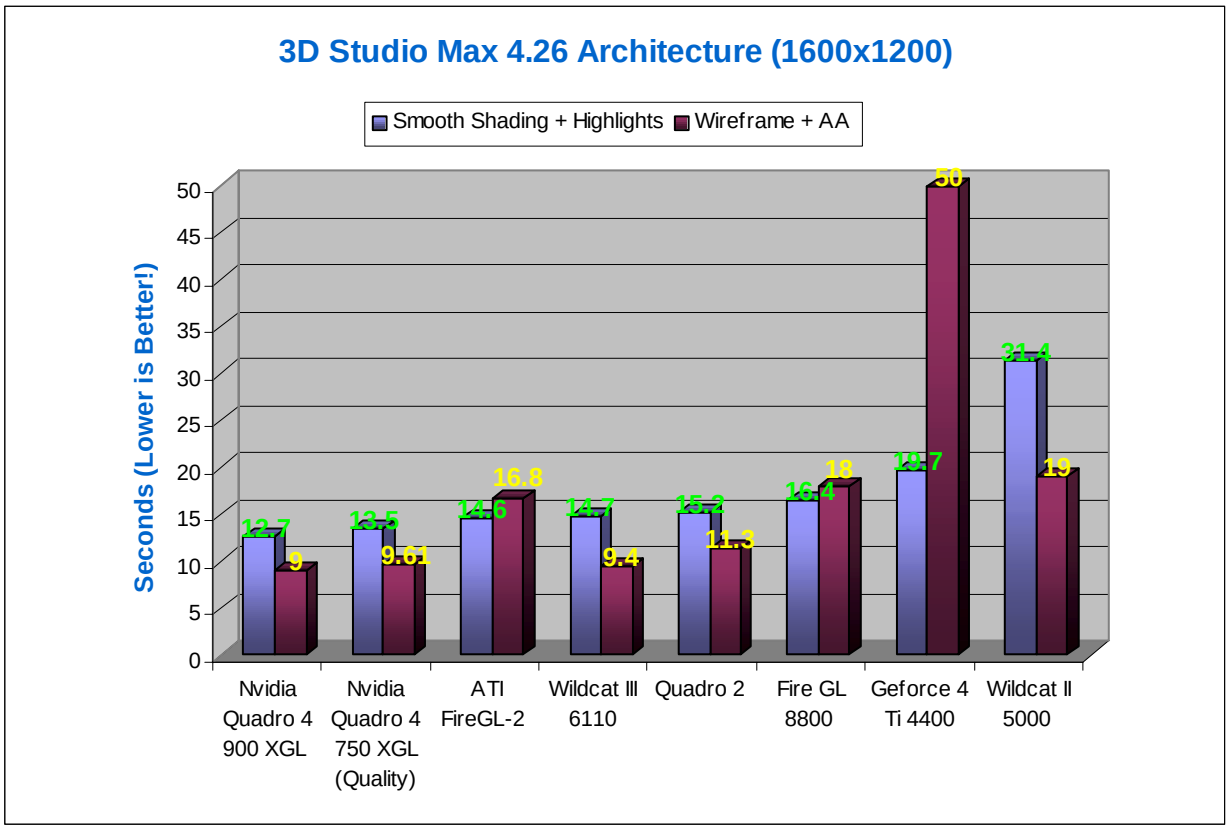

Nvidia's Maxtreme driver does some real magic and widens the gap between the Quadro 4 boards and the competition.

We also tried the "lit wireframe" mode, which produced some strange results.

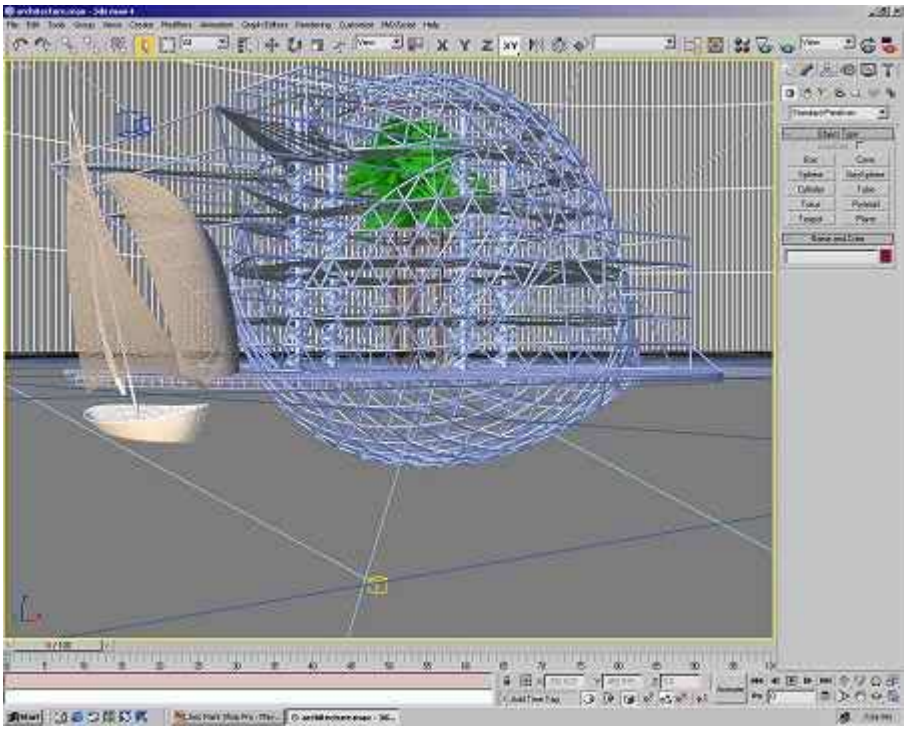

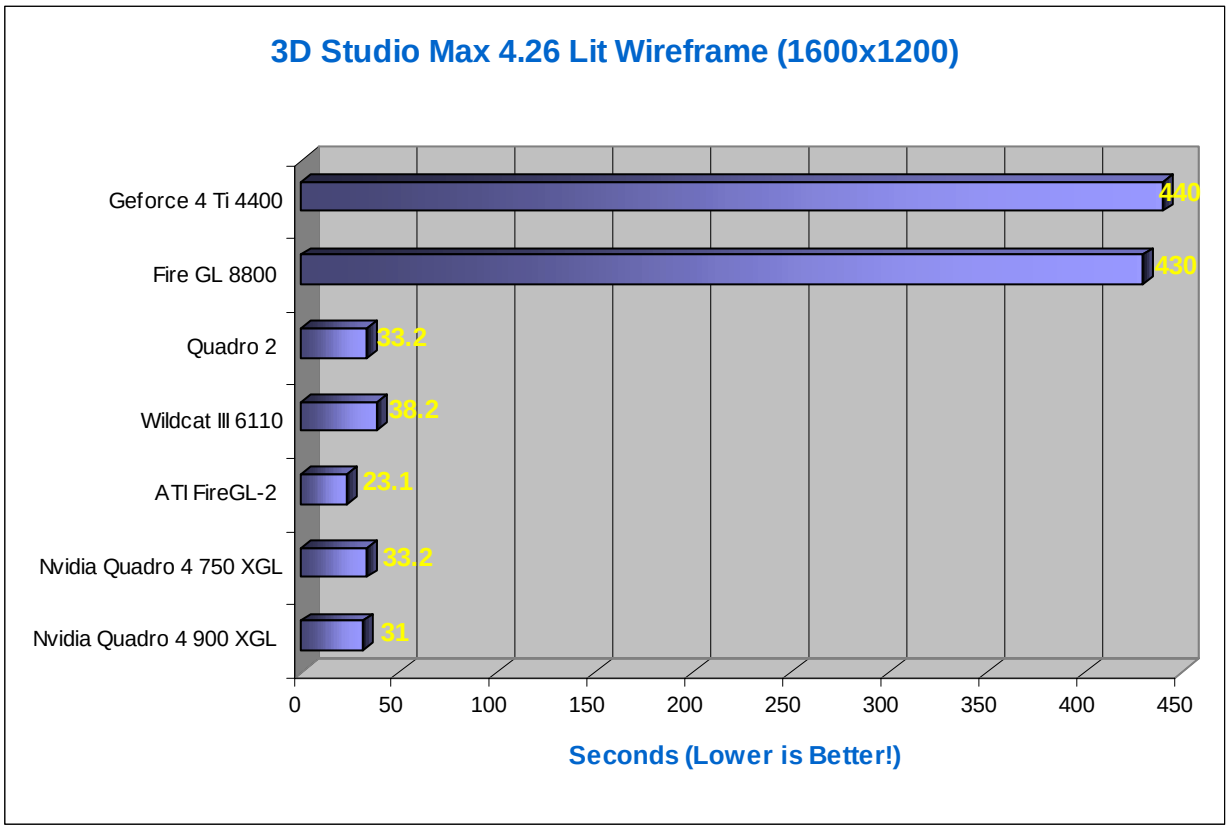

We are not entirely sure how many artists use this mode, but some cards do not seem to have hardware support for it. The FireGL 8800 and the gaming cards completely failed to accelerate this mode. It seems that the mode is running on the CPU only in this case.

Conclusion

Before we start with the conclusion, I'd like to make note of a few points. While benchmarking and working with these cards has taken weeks, we are well aware you might still have some questions. How do the cards perform on slower Pentium 4 CPUs, how do they perform on a dual-processor Athlon platform? How do their dual display modes compare? We'll find out in an upcoming article, as these questions deserve good answers.

So what can we conclude so far?

The Wildcat III 6110 is, on average, the fastest card for professionals, especially for high-end CAD applications. Contrary to its older brother, the Wildcat II 5110, we could not observe any real weakness: the third Wildcat performs excellently whether you work in wireframe, shaded, or textured moded with or without lights. It is the best performer in almost all CAD tests - especially with line AA enabled, and it offers superb image quality. The Maya benchmarks make it clear that no matter what mode you are working in, the Wildcat III's performance is among the best. The Blinn shading performance should have been slightly better, as the Quadro 4 is faster at high resolutions, but the Wildcat III makes up with superb texturing and lighting performance.

3DLabs can not afford to rest on its laurels, however. The most expensive card gets beaten in 3DSMax by the Quadro 4 and the latter costs between $600 (XGL 900) and $1000 (XGL 750) less. The performance difference in CAD has been reduced significantly, compared to the "Quadro 2-Wildcat 5110" days. For those willing to pay the high price, the Wildcat III can still offer unique features to high-end users: like an enormous onboard texture buffer and excellent image quality (probably the best).

The Quadro 4 XGL 750 is probably the most dangerous opponent 3DLabs has ever had to face. While the Quadro 2 produced pretty ugly pictures (for an OpenGL card), the Quadro 4 offers excellent image quality and a revolutionary price/performance ratio for this market. It is by far the best card for 3DSMax, and a very strong one for CAD as well.

The 3DLabs Wildcat II 5000 is still a decent wireframe CAD solution, but the cheaper Quadro 4 750 XGL sweeps the floor with 3DLabs' budget product in all other disciplines. 3DLabs needs a new and better budget card and needs it fast.

The Quadro 4 750 GXL is also a FireGL 4 killer. While ATI still positions this card as their high-end CAD solution, it is clear that the FireGL 2 and FireGL 4 are no match for the Quadro 4 in both DCC and CAD applications. At high resolutions, the Quadro 4 consistently walks all over the FireGL 2 and the slightly higher-clocked FireGL 4 will likely have the same problems.

ATI's FireGL 8800 fixes some of the weaknesses of the FireGL 2 and 4, but fails to convince us in wireframe mode. The FireGL 8800 is no CAD card, but is thanks to its very low price ($550, less than $300 for the FireGL 8700) it is an attractive low budget solution. Nevertheless, we feel ATI should try to get better wireframe performance out of its drivers. Otherwise, it will be hard to resist Nvidia's Quadro 4, which is - in many cases - the best card in the low-end and mid-range markets.

Summary

3DLabs Wildcat III 6110

Pros

- Fastest CAD solution available with exceptional line anti-aliasing performance
- Only high-end accelerator with 128 MB of texture memory and 32 hardware accelerated lights
- Performance is very good in all cases (wireframe, shading, texturing, lighting)
- Superb image quality

Cons

- Very expensive compared to the Quadro XGL 750
- 2D speed is not the best of the industry

Quadro 4 750 XGL / 900 XGL

Pros

- Excellent image quality
- Extremely powerful MaxTreme drivers, best 3DSMax performance
- Excellent 2D Speed
- 750 XGL offers by far the best price/performance ratio in the industry.

Cons

- Why pay so much more for the 900 XGL?
- Mediocre/poor performance with more than 8 lights
- Mediocre line anti-aliasing performance in some high-end CAD applications

ATI Fire GL 8800

Pros

- Good DCC performance, especially with many lights
- Good line anti-aliasing performance
- Good 2D Speed
- Low price

Cons

- Mediocre image quality, especially for CAD
- Disappointing wireframe performance
- Mediocre Blinn shading performance at high resolutions

Special thanks to Ping-Che Chen for his assistance on the **General Message Board**.