Review

NVIDIA® GeForce™ 6800 GT PCI Express

by **John Reynolds**

Introduction

NVIDIA's mantra last year for their GeForce FX chips was the depth of their GPUs' pipelines, eschewing a wide architecture for a narrow (think four pixel pipelines) yet deep one. With the announcement of the NV40 architecture this spring, NVIDIA revealed a dramatic change in its approach with the GeForce 6800 series, new GPUs that boast up to four times the number of pixel pipelines as their predecessors. In addition to this wider architecture, NV40 also incorporates hardware support for such cutting-edge features as DirectX 9.0c's Shader Model 3.0 and floating point blending. NVIDIA traveled a somewhat rocky road last year, so has this design change paid off for them? In a word: yes. The high-end chip of the new GeForce 6 series, the 6800 Ultra, is roughly twice as fast as the company's best GPU from last year, bringing NVIDIA back to a very strong position with the market from both a performance and technology perspective.

The NVIDIA GeForce™ 6800 GT chip

The GeForce 6 series is comprised of three models: the 6800 Ultra, the 6800 GT, and the vanilla 6800. The flagship product, the GeForce 6800 Ultra, has a core clock speed of 400 MHz and 550 MHz GDDR3 memory, and a suggested price of $500. The Ultra is a dual-slot board due to the size of its cooling solution, and NVIDIA recommends users to run a power supply rated at 350W or higher (particularly if overclocking) with the Ultra boards. The 6800 GT is a slower, single-slot version of its Ultra sibling, with 350/500 MHz chip and memory speeds and a more popular $400 pricing. Both the Ultra and GT models boast 16 pixel pipelines and six vertex shader units. Last, the plain 6800 model is a 12-pipe, $300 mainstream solution clocked at 325/350 MHz. Worth noting is that these are reference specifications NVIDIA provides to their AIB (add-in board) vendors, and are therefore subject to change as the vendors attempt to differentiate their products with special or unique "samples".

NV40 Architecture Overview

The GeForce 6800 Ultra and GT, as mentioned above, are a 16x1 architecture, and such a wide, parallel design does not make for a small chip. Produced on a 130nm fabrication process, these new GPUs have an unprecedented

transistor count of over 220 million. The reference board NVIDIA sent is a PCI Express solution that has a 256-bit memory interface, with dual DVI and one TV-out connectors driven by two 400 MHz DACs. Moreover, the NV40 architecture upon which the GeForce 6 series of chips is based has the richest feature-set for GPUs currently available on the market.

The NVIDIA GeForce™ 6800 GT PCI Express

And while SimHQ is a site that focuses on simulations-based PC gaming, we would be amiss if not providing a brief look at this architecture:

CineFX 3.0 Shading Architecture

- Full DirectX 9.0 support
- Shader Model 3.0 support
- Infinite Shader lengths
- MRT (multiple render targets) support
- 16 textures per rendering pass
- 32-bit and 16-bit floating point format support

High-Precision Dynamic-Range

- Full floating point support through entire pipeline
- 16-bit floating point frame buffer blending

Intellisample 3.0

- Rotated-grid, multisampling anti-aliasing
- 16x anisotropic filtering
- Lossless compression algorithms for color, texture, and z-data
- Fast Z clear

UltraShadow II

- Stencil shadow performance acceleration

Advanced Video and Display

- Dedicated on-chip video processor
- MPEG video encode and decode
- Digital Vibrance 3.0 control

Shader Model 3.0, an evolutionary step from the 2.0 model, lifts a number of restrictions within its predecessor for ease of programming while adding new features such as dynamic branching and vertex texturing, among others. And while game support for the new model (which will be exposed in DirectX 9.0c, included with SP2 for Windows XP) is, for the moment, very sparse — currently limited to Far Cry's recent 1.2 patch, which is still not yet available — NVIDIA can certainly be expected to do their utmost to work with developers to take advantage of their new GPUs' hardware support for the 3.0 model. In addition to SM 3.0, the 6800s' ability to perform floating point blending allows for accumulation effects such as motion blur (reminiscent of 3dfx's T-buffer) and soft shadows. And the 32-bit floating point precision throughout the entire pipeline, from input to the rendered output, allows for HDR (high dynamic range) effects, which can be used to create more realistic scenes by working with a superior gradient of brightness ranges. All of the above are unique to NVIDIA's NV40 architecture, as no other GPU geared toward the mainstream or gaming markets currently available supports these features.

Naturally the above capabilities will have to be specifically incorporated into games by developers, but NVIDIA should be congratulated for being the first to bring these features to market. But the hardware support for such advanced features and a 16x1 architecture that also includes two shader units per pixel pipeline does not come for free in terms of the transistor count of the GeForce 6800s, and with over 220 million transistors even the flagship Ultra chips are clocked at speeds considerably lower than those of the previous generation. Let us take a look at how a PCIe version of the 6800 GT handles the SimHQ benchmark suite.

Test System Setup

- Intel Pentium 4 Extreme Edition 3.4 GHz
- Intel D925XCV motherboard (BIOS CV92510A.86A.0193)
- 1 GB (2x512 MB) Micron DDR2 533 MHz RAM (4-4-4-12)
- NVIDIA GeForce 6800 GT PCI Express graphics board (ForceWare 61.45)
- Maxtor MaXLine III (16 MB buffer) native SATA HD (2x) in a RAID 0 array (NTFS)
- Windows XP Professional (SP1)
- DirectX 9.0b

The benchmark suite that will be used to evaluate this test system is listed **here**. Again, unless specified otherwise all games are configured to their highest settings, and 32-bit color and trilinear texture filtering are the default baseline during testing. Also, Windows XP is configured to have Automatic Update, System Restore, and all unnecessary startup services disabled. Fraps 2.2.1 is used to record performance scores unless otherwise noted.

The GeForce 6800 GT being reviewed is a PCI Express board clocked at 350 MHz with 256 MB of 500 MHz (1 GHz effective) GDDR3 memory, which unlike the Ultra PCIe board is identical to the AGP 8x version. And being a PCIe design, the reference board has one 6-pin power connector rather than a 4-pin molex connector. The beta ForceWare 61.45s driver set — NVIDIA currently does not have WHQL drivers for the GeForce 6800 series — were used and manually configured for high quality settings, with both trilinear and anisotropic filtering optimizations disabled.

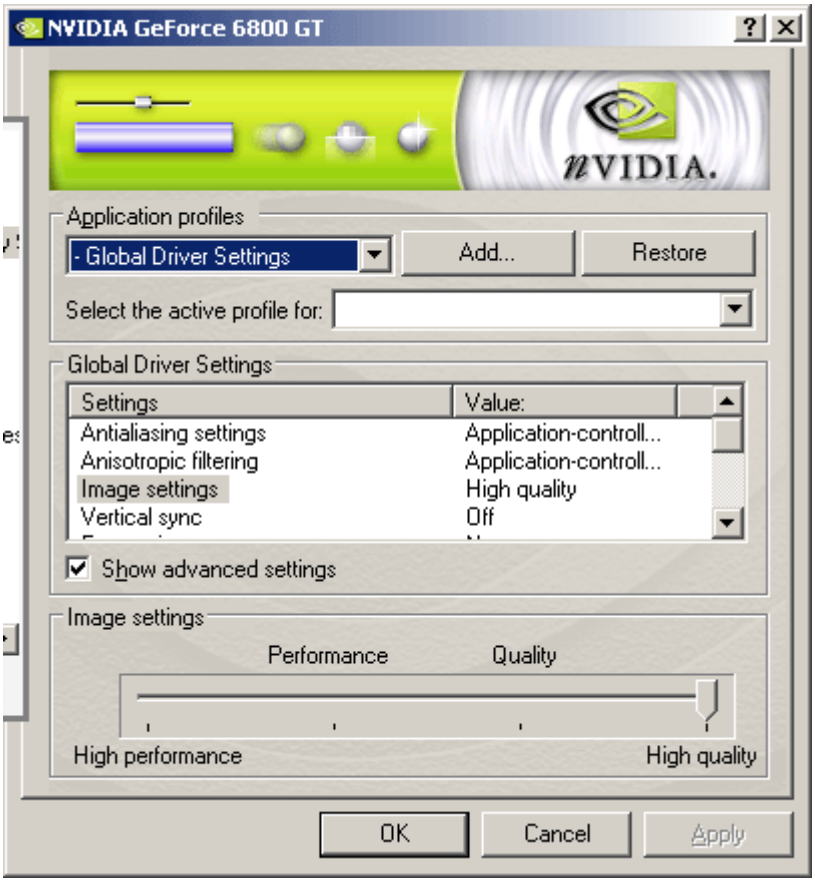

Also worth noting is that NVIDIA's current PCI Express solutions are not native PCIe boards, instead making use of a HSI (high-speed interconnect) bridge chip to communicate between the graphics bus and the chip itself. Capable of realizing the full upstream bandwidth of the PCIe x16 bus, the HSI chip should not be the cause of any performance issues for those with PCI Express motherboards as current gaming software is simply not limited by the graphics bus (PCIe or AGP).

Benchmarks

Note: The High Quality test setting designates benchmark scores for when both 4x AA and 8x AF are enabled.

Lock On: Modern Air Combat was tested using the MiG-29 Intercept demo. For readers interested in doing comparison testing, please check the "How SimHQ Tests" page for exact LOMAC settings used during testing — several features, such as water detail, were lowered to avoid poor frame rates throughout all test settings. The Intercept demo was run until the six minute mark.

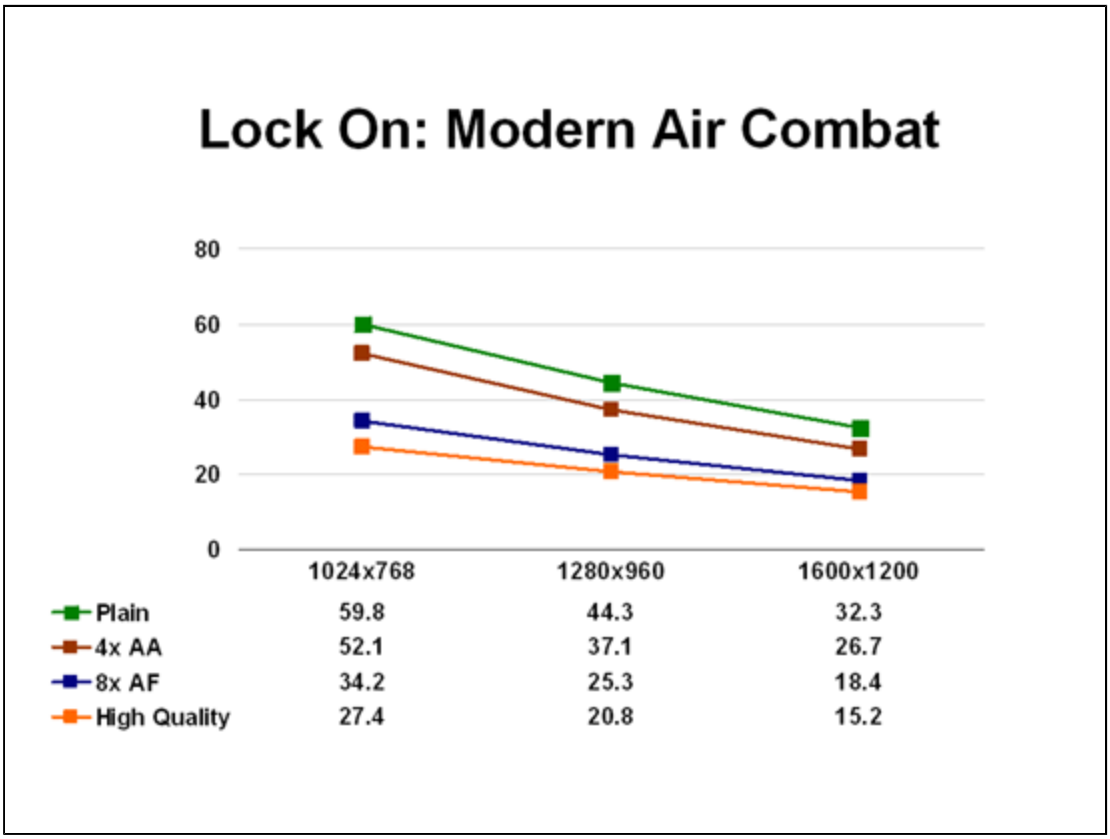

With the filtering optimizations disabled, it is interesting to see anisotropic filtering impose a much higher performance penalty than anti-aliasing, roughly 25% compared to AA's 15%. And once both features were enabled Lock On is playable only at the resolution of 1024x768. For curiosity's sake, additional testing was done with the game's graphics set to its in-game of Low, and at 1600x1200 with high quality settings the frame rate jumped up 50% to 23.3 (from 15.2), and to 19 fps at Medium settings.

Microsoft's **Flight Simulator 2004** is a strongly system-limited title that tends to scale better with processor rather than GPU upgrades. SimHQ's testing consists of an in-house dusk flight over the city of Hong Kong, with an external camera view set from behind the plane. Frame rate recording is stopped once the plane touches down. All game options were set to their highest option, save for Ground Scenery Casts Shadows, which was disabled.

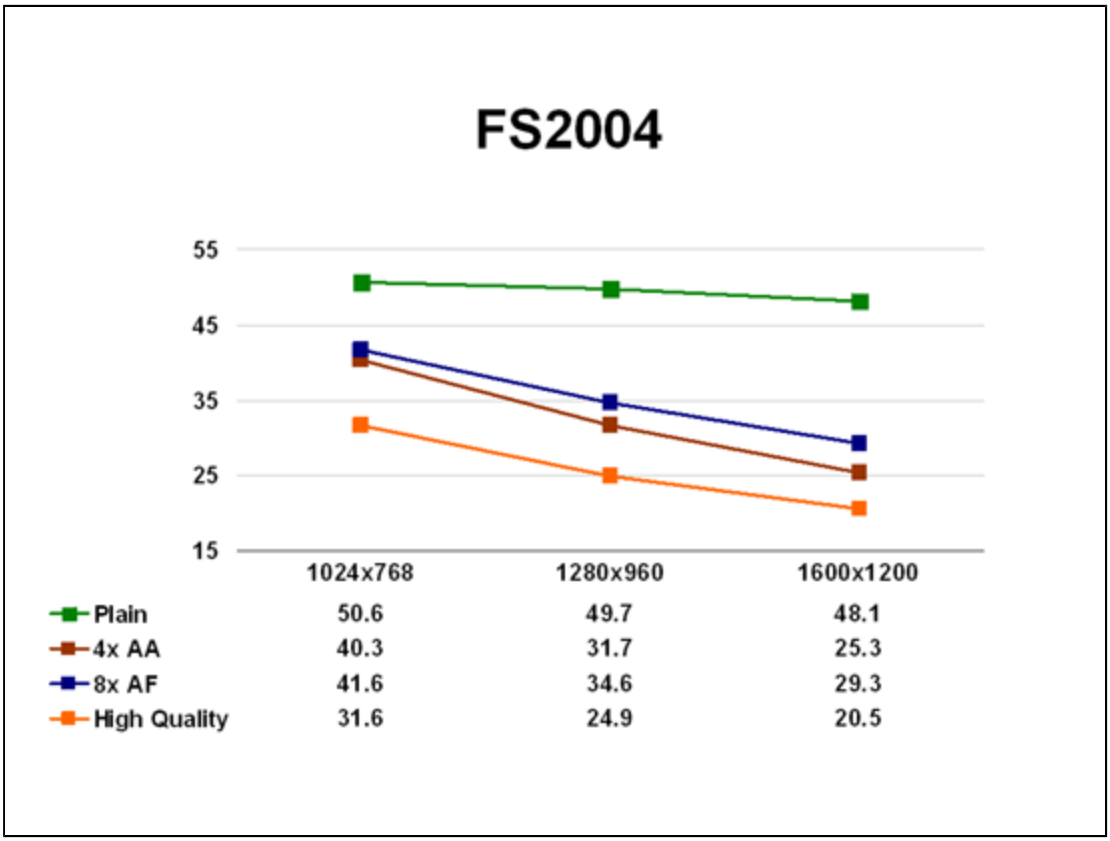

At 1024x768 both AA and AF cost the game roughly 20% of its performance, yet with the frame rate loss sharply rising at the higher resolutions. Anti-aliasing incurred a slightly higher loss for all three tested resolutions, and depending on personal preference for what constitutes “playable” only the two lower resolutions offer usable performance with both features enabled.

The **IL-2: Sturmovik Forgotten Battles - Aces Expansion Pack** represents SimHQ’s non-modern flight simulation test. Using OpenGL rather than D3D, the landscape option was set to perfect. Testing consisted of using Fraps to record the frame rate during the first two minutes of the Black Death track. The use of this track is a change from previous testing, and the switch was made since the Black Death track simply displays more action than that of the Bf109 Introduction.

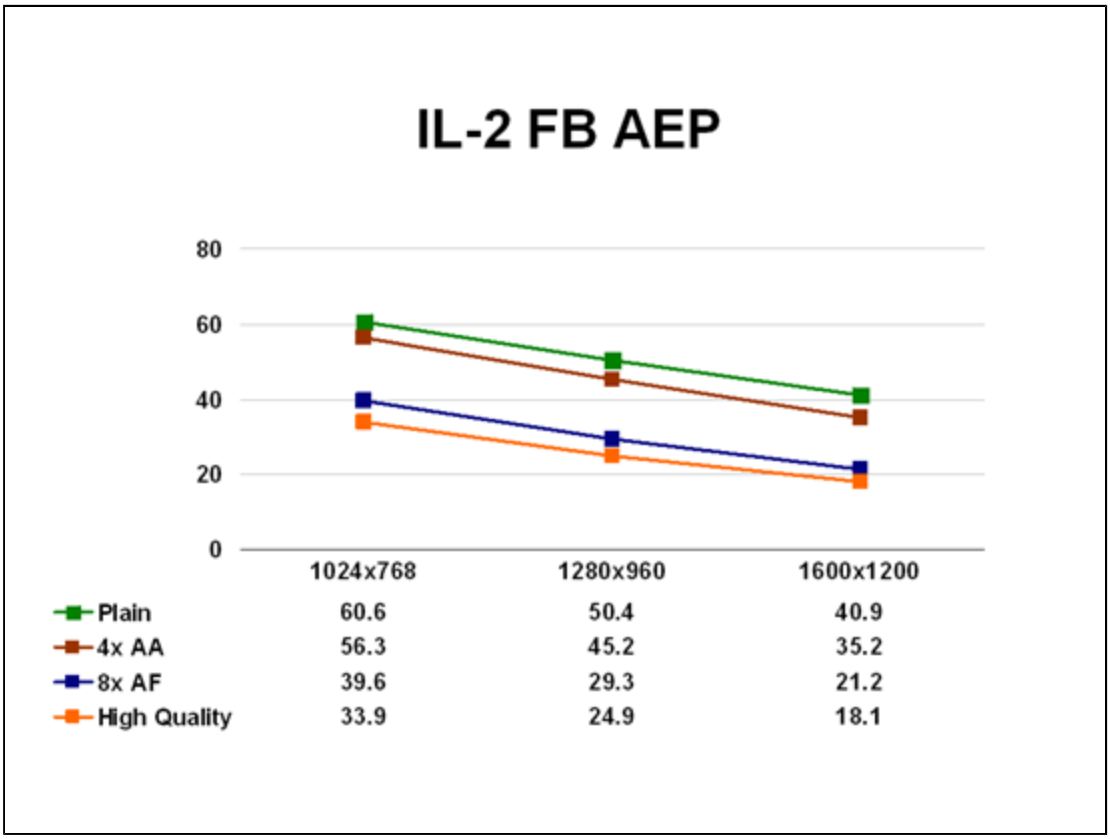

Testing with 4x anti-aliasing gave a consistent near 10% performance loss across the tested resolutions. Anisotropic filtering, on the other hand, imposed a surprisingly strong penalty, dropping the frame rates 35% to 50% as the resolutions scaled upward, making the resolution of 1600x1200 unplayable even with a P4 3.4EE CPU. Struck by this performance loss, we reran the 8x AF tests again, this time with both trilinear and anisotropic filtering optimizations enabled in the driver panel. The results were consistently 20% higher, with frame rates of 47.7, 35.5, and 25, respectively; however, these optimizations did result in a noticeable increase in texture aliasing during gameplay, though how much of a concern this may be is, of course, subjective to the end user's preferences.

The venerable **Falcon 4** has received a recent overhaul with the release of SuperPAK 4.1. Testing was conducted using SimHQ's demo, a low level, air-to-ground dusk mission that consists of two Falcons using Mk20s and Mavericks. The outbound flight route takes the Falcons over the city and into a hot combat zone.

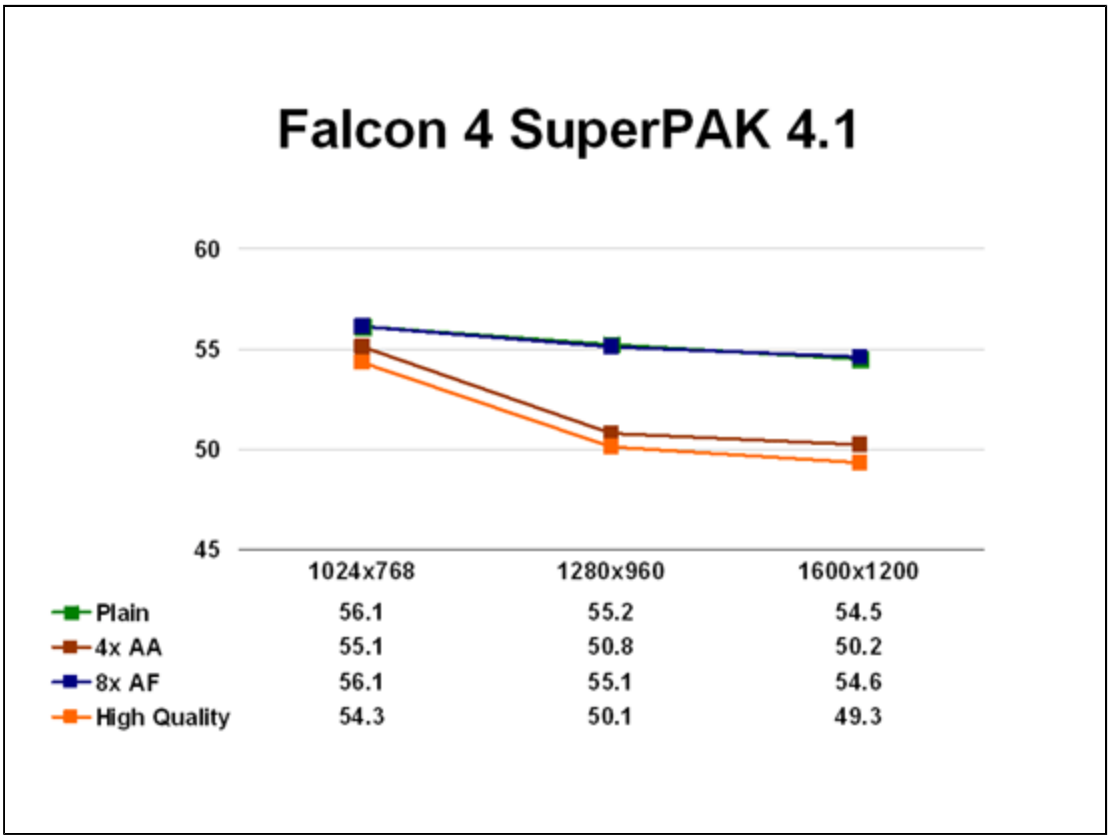

With a game based on such an aging, albeit overhauled, graphics engine, Falcon 4 completely fails to scale with faster graphics subsystems. Even the high quality settings at the highest resolution barely moved the simulation 10% away from is performance at a vanilla 1024x768. Perhaps it is time this title was put to pasture, so to speak, in terms of continued use for hardware testing? Feel free to voice your opinion **here** in SimHQ's Hardware Forum.

For the sake of consistency, **Far Cry** (v1.1) was tested with all graphical options placed at their highest settings and sound disabled. The Research map was used, with testing consisting of playing completely through the map in God mode using the same basic path each time. Anti-aliasing was enabled via the game's menu option, and anisotropic filtering via the ForceWare driver's profile for Far Cry.

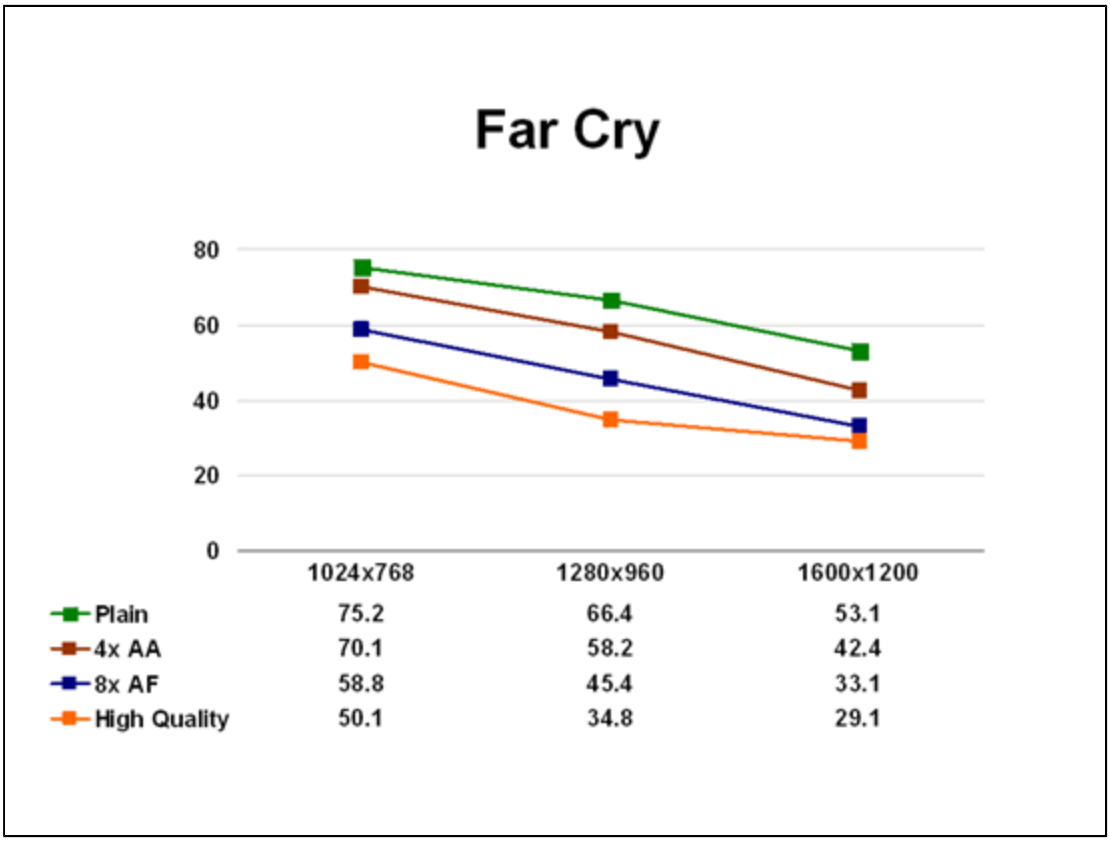

Like IL-2: Forgotten Battles, Far Cry's engine also imposed a much steeper performance penalty for anisotropic filtering than anti-aliasing. The latter's loss varied from 10% to 20% across the tested resolutions, while higher filtering cost the game 20-30% depending upon the resolution; again enabling the filtering optimizations resulted in a score of 70.3 for 8x AF at 1024x768, once more a roughly 20% performance improvement. Yet combined, AA and AF render Far Cry unplayable at the resolution of 1600x1200 when playing with all in-game settings placed at their highest option.

Call of Duty's Dawnville demo was used to test this OpenGL title based on id Software's Quake 3 engine. The included timedemo utility was used to record performance rather than Fraps, and the "com_maxfps" command was used to raise the default frame limit of 85 fps.

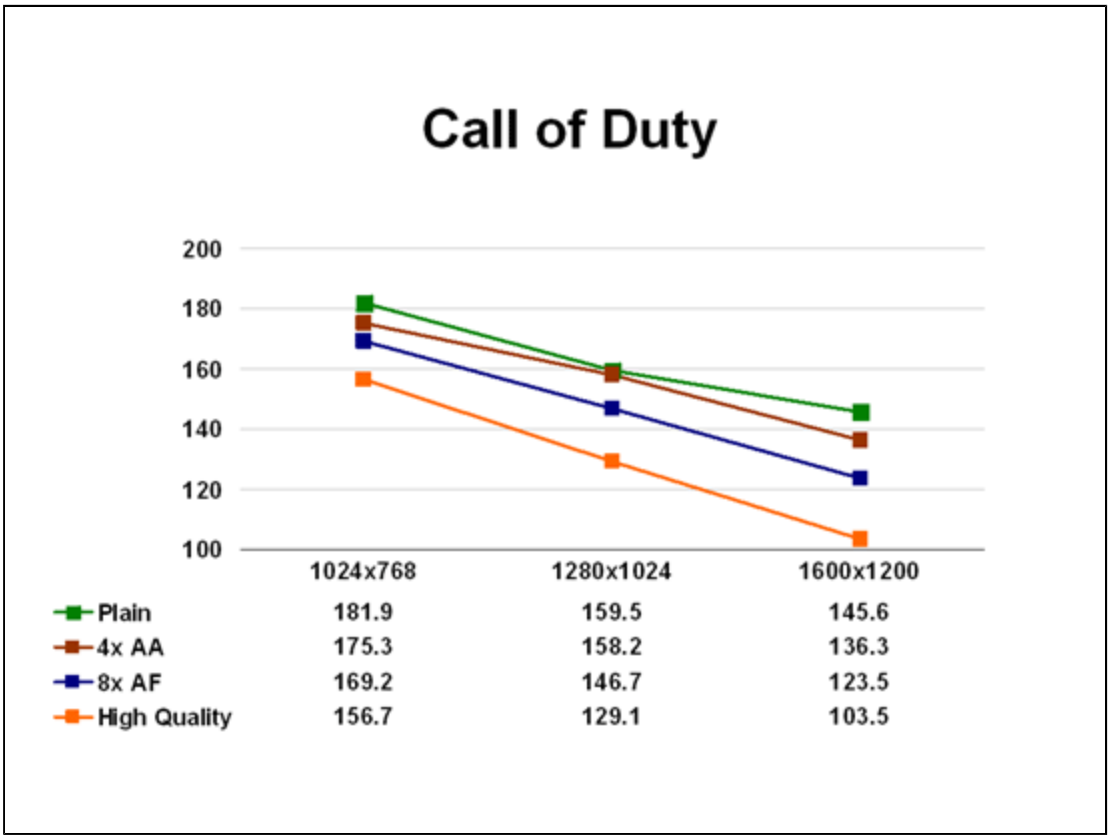

With scores as high as listed above, additional reviewer commentary hardly seems necessary. Once again, anisotropic filtering incurs a higher performance penalty than anti-aliasing, yet nothing that impacts the frame rate too terribly.

NASCAR Racing 2003 Season was tested using SimHQ's crowded Daytona replay. The camera was fixed to Earnhardt's cockpit for the view mode, and all graphical options were set to their highest.

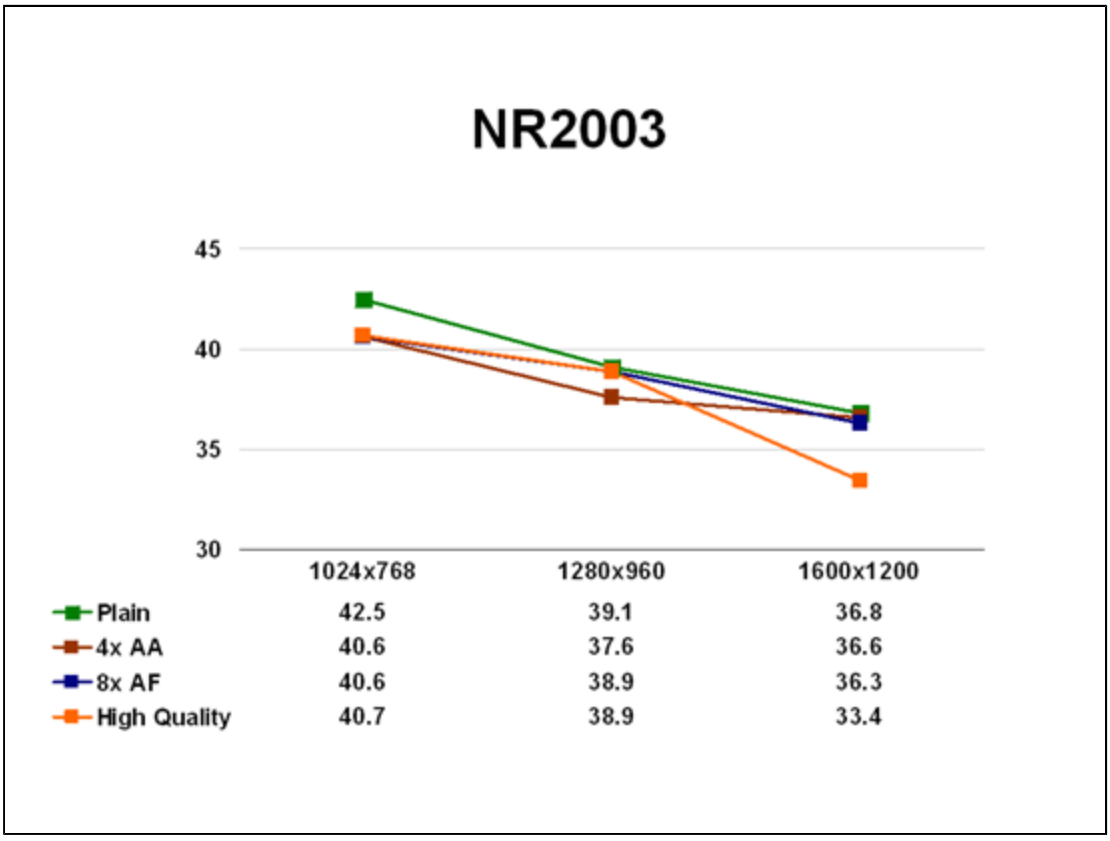

NASCAR is clearly CPU-limited as its failure to scale with the test settings indicates. Regardless, the game remains playable at even the highest resolution and settings tested.

Image Quality

For years NVIDIA GPUs have been limited to an ordered-grid sampling pattern for their 4x anti-aliasing, which produced an inferior quality when dealing with polygon edges running along near-vertical or horizontal angles. Yet the 6800s now employ a rotated-grid pattern, which can be seen below:

4x AA Sampling Pattern

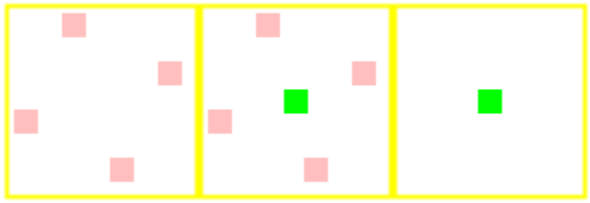

Unfortunately, due to a hardware ROP (render output) limitation, the 6800s can only perform a maximum of 4x multi-sampling AA. Yet the 6800 series also supports an 8xS anti-aliasing mode, which is a combination of multi- and super-sampling methods. The sampling pattern for the 8xS mode can be seen here:

8xS Sampling Pattern

The GPU companies moved away from super-sampling several years ago due to its performance loss from rendering additional color data for each sub-sample. Yet one advantage of super-sampling is that the additional color data created would result in greatly reduced texture aliasing (e.g. texture shimmering); and as seen above, the 8xS mode writes color data into two rather than one of its sub-samples and this incurs a higher fill rate cost as shown by the additional IL-2: Sturmovik Forgotten Battles - Aces Expansion Pack testing below:

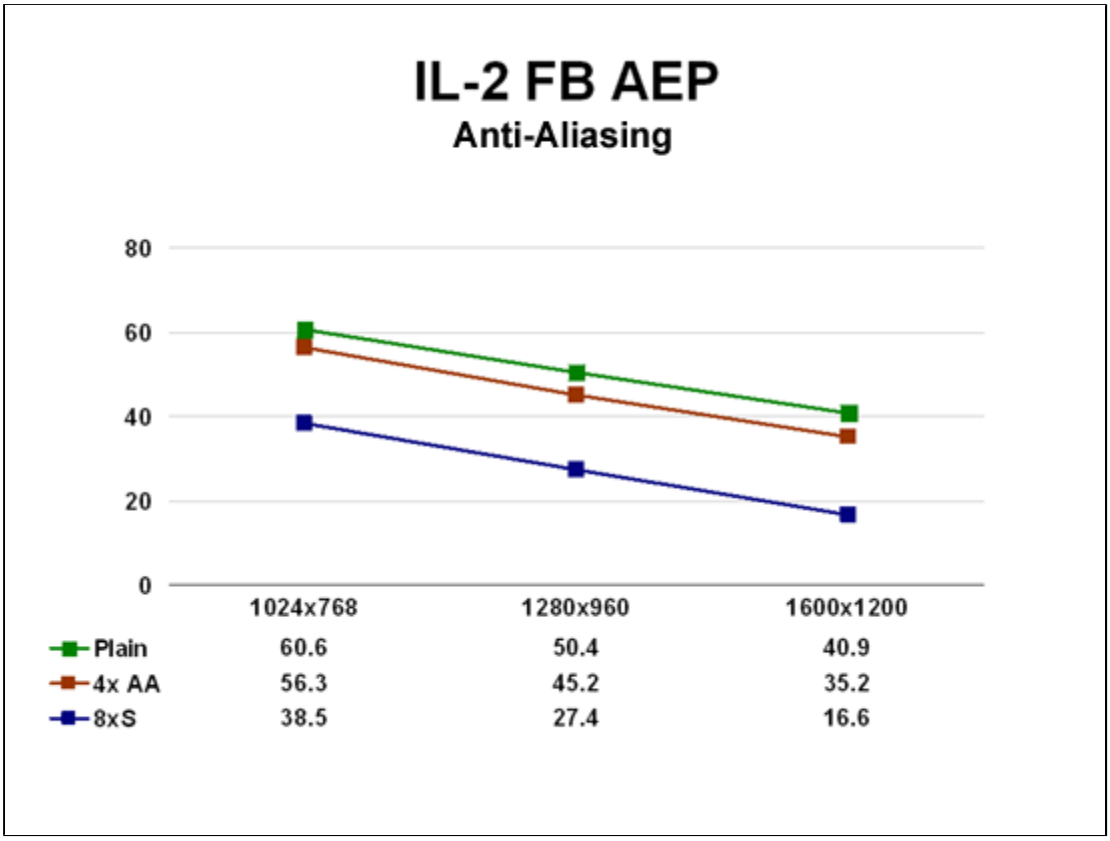

Whereas 4x multi-sampling dropped performance across the tested resolutions by roughly 10%, 8xS incurs a much higher frame rate loss, one that ranges from 40-50% depending upon the resolution. However, 8xS remains a viable mode at 1024x768 in newer titles, is particularly useful for many older games, and is generally worth enabling when possible due to the image quality improvements it provides. Note that all game screenshots throughout this review were taken at 1024x768 with 4x AA.

A theme throughout the benchmarking portion of this review has been the rather strong performance impact the 6800 GT's anisotropic filtering had on certain games with the card's filtering optimizations disabled. Using the Direct3D AF Tester utility to capture the images below, we can clearly see the differences the optimization settings have on texture filtering:

Direct3D Anisotropic Filtering (AF) Test Results

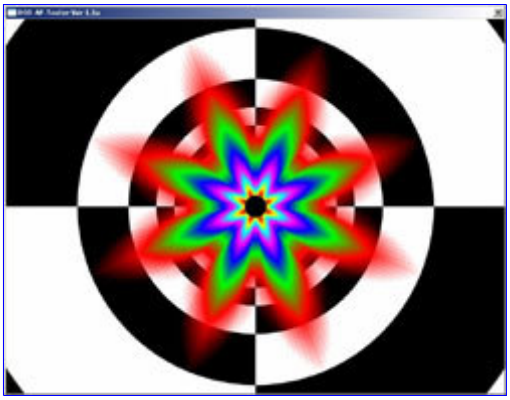

Optimizations Off

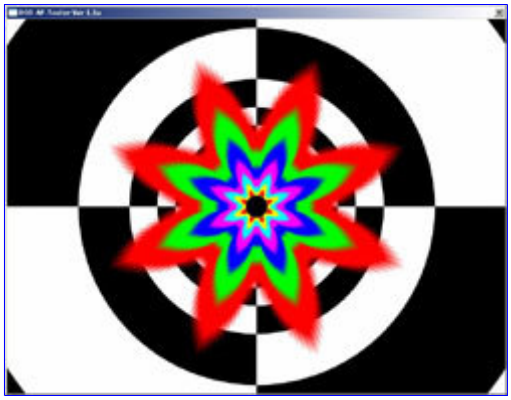

Optimizations On

Most noticeable is the decreased amount of blending that occurs between MIP maps (color bands), which is a result of the trilinear optimization. This effect is often not apparent during gameplay, though for those with sharper eyes the choice is available with the 61.45s to disable all optimizations. NVIDIA is to be commended for giving the end user the option of choosing either better image quality or somewhat lessened IQ for increased performance when it is needed.

Overclocking

With the 6800 GT's 350 MHz chip speed and 500 MHz GDDR3 memory, it trails the flagship Ultra settings by some 10%; which, considering the suggested retail pricing for both models makes the GT an attractive product in its price-performance ratio. Yet to see how a 6800 Ultra would fare, the CoolBits registry value was added to enable the clock frequencies option in the driver panel. SimHQ, however, cannot stress strongly enough that the overclocking results of one review board cannot possibly begin to serve as an adequate sampling of all similar products, and to please keep this in mind while reading the following section.

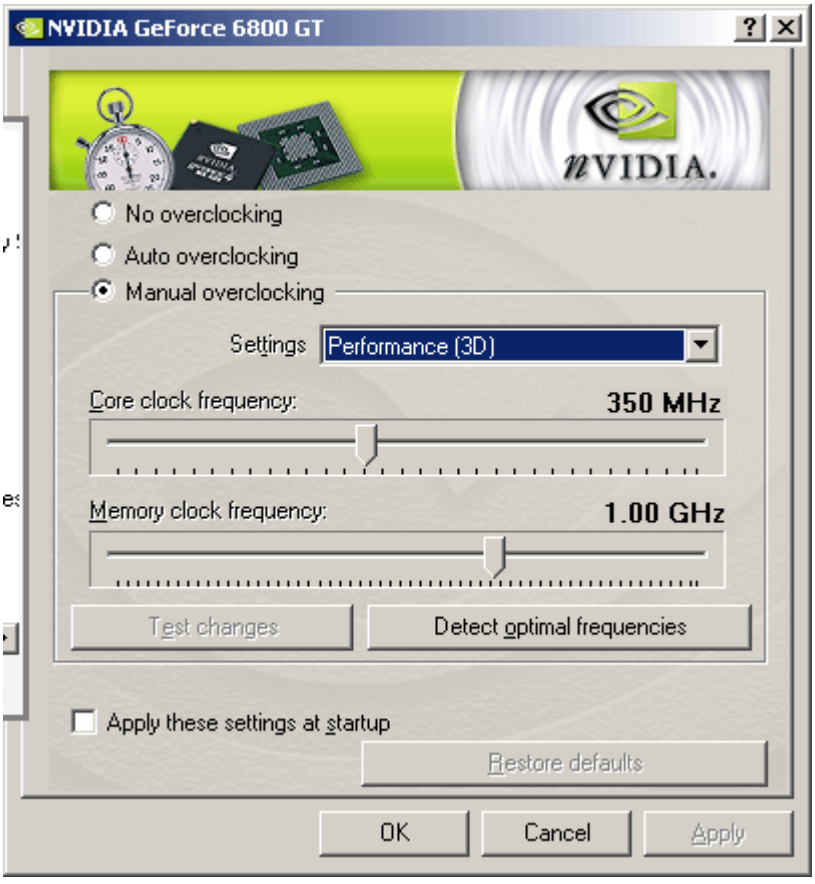

The reference GT overclocked to Ultra speeds (400/550 MHz) without issues, and because it scaled so well with resolution increases we decided to again use the IL-2: Sturmovik Forgotten Battles - Aces Expansion Pack.

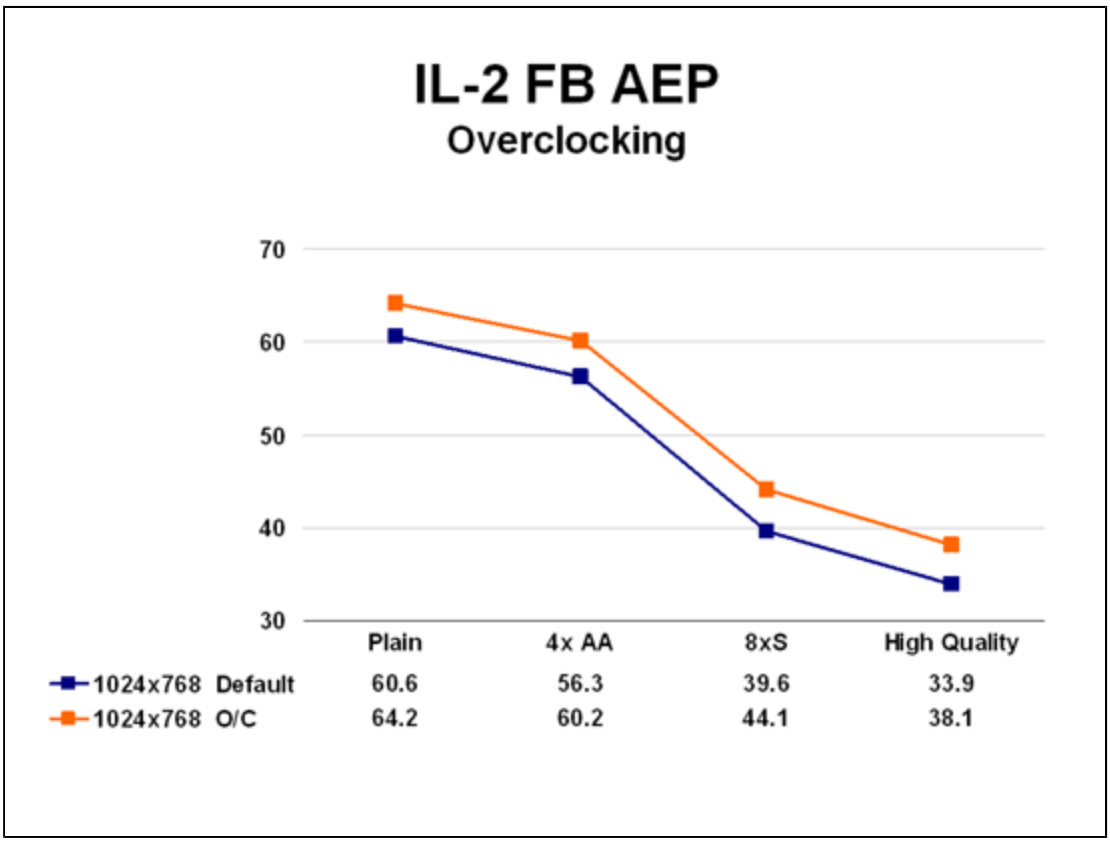

The overclocked results are almost but not quite consistently 10% faster than the default scores, closely matching the core and memory speed increases. It is worth noting that with the reference cooling solution the 6800 GT ran reliably overclocked at Ultra speeds during hours of gaming, though this is again certainly no guarantee that all GT cards will do so. Also note that the PCIe 6800 Ultra cores are clocked at 435 MHz compared to the AGP version's 400 MHz.

Gallery

Here are some screenshots taken with the test GeForce 6800 GT card.

Lock On: Modern Air Combat

Far Cry

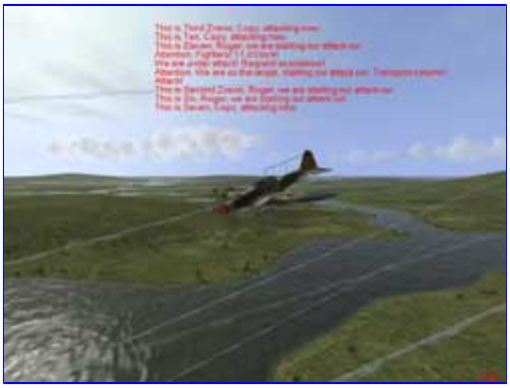

IL-2 Sturmovik: FB - AEP

NASCAR Racing 2003 Season

Conclusions

The PCIe GeForce 6800 GT offers excellent performance for its price. The card burned through SimHQ's benchmark suite setting new highs in frame rates, with no visual glitches or anomalies, and no compatibility or stability issues. The GT boards could very well become the Ti 4200s of their generation. NVIDIA has also improved the image quality that is produced by their chips, with better anti-aliasing quality at 4x samples and support now for 16x anisotropic filtering. And the ForceWare driver panels give end users a simple yet elegant interface by which they can quickly adjust various settings and options; moreover, NVIDIA's drivers have for months now supported individual game profiles, which enables users to easily apply different settings for each title they have installed.

The GeForce 6 series appears to herald NVIDIA's return as a company that aggressively engineers chips capable of fighting for both the performance and technology crowns. The new design also signifies the company's abrupt shift away from its former deep architecture of the FX chips with the 6800's widely parallel pipelines and the substantial performance gains it realizes, positioning the board models very competitively for the different price points they address. And hardware support for Shader Model 3.0 and advanced features such as the high dynamic range and accumulation effects possible due to the full floating point precision formats and buffer blending set NVIDIA's new GPUs apart from the competition. How much developer attention these features receive, however, will be determined by whether or not the PC gaming industry settles upon SM 2.0 or 3.0 as the primary inflection point for DirectX 9 support. Regardless, all else being equal — price, performance, compatibility, image quality, etc. — it is hard not to favor the part that offers more advanced, forward-looking features.