




TECHSPOT NEWS PRODUCTS DOWNLOADS FORUMS Search

- Home Reviews Guides Product Finder Forums Downloads Drivers Extras Sign in About

- CPU & Motherboards Graphics Desktops Laptops Tablets Smartphones Storage Cases Input Software & Games

HOME > FEATURES

History of the Modern Graphics Processor, Part 3



TechSpot on: Follow

Most Popular

- Trending** | Featured
- Hacker demonstrates how to hijack an airplane using an Android app**
37 Comments
- The History of the Modern Graphics Processor, Part 3: The Nvidia vs. ATI era begins**
21 Comments
- Iranian scientist says he has built a machine that can predict the future**
15 Comments
- PC industry posts steepest quarterly sales decline in history**
21 Comments
- Next Xbox to feature AMD processor, not compatible with 360 titles**
35 Comments
- [More Trending Topics](#)

By Graham Singer on April 10, 2013

Find and Compare the Best Tech Products

Read user comments (22)

The Fall of 3Dfx and The Rise of Two Giants

With the turn of the century the graphics industry bore witness to further consolidation.

The pro market saw iXMICRO leave graphics entirely, while NEC and Hewlett-Packard both produced their last products, the TE5 and VISUALIZE FX10 series respectively. Evans & Sutherland also parted ways with the sale of its RealVision line to focus on the planetaria and fulldome projection systems.

In the consumer graphics market, ATI announced the acquisition of ArtX Inc. in February 2000, for around \$400 million in stock. ArtX was developing the GPU codenamed Project Dolphin (eventually named "Flipper") for the Nintendo GameCube, which added significantly to ATI's bottom line.

Also in February, 3dfx announced a 20% workforce cut, then promptly moved to acquire Gigapixel for \$186

Featured on Keyboards

- Metadot Das Keyboard Professional**
79 6 reviews
- Logitech Illuminated Keyboard**
87 24 reviews
- Razer BlackWidow Ultimate**
24 reviews

million and gained the company's tile-based rendering IP.

Meanwhile, S3 and Nvidia settled their outstanding patent suits and signed a seven-year cross-license agreement.



ATI GameCube GPU

VIA assumed control of S3 around April-May which itself was just finishing a restructuring process from the acquisition of Number Nine. As part of S3's restructuring, the company merged with Diamond Multimedia in a stock swap valued at \$165 million. Diamond's high-end professional graphics division, FireGL, was spun off as SONICblue and later sold to ATI in March 2001 for \$10 million.

3DLabs acquired Intergraph's Intense3D in April, while the final acts of 3dfx played out towards the end of the year, despite 2000 kicking off with the promise of a better future as the long-awaited Voodoo 5 5500 neared its debut in July. The latter ended up trading blows with the GeForce 256 DDR and won the high-resolution battle.

Where 3dfx was once a byword for raw performance, its strengths around this time laid in its full screen antialiasing image quality.

But where 3dfx was once a byword for raw performance, its strengths around this time laid in its full screen antialiasing image quality. The Voodoo 5 introduced T-buffer technology as an alternative to transformation and lighting, by basically taking a few rendered frames and aggregating them into one image. This produced a slightly blurred picture that, when run in frame sequence, smoothed out the motion of the animation.

3dfx's technology became the forerunner of many image quality enhancements seen today, like soft

shadows and reflections, motion blur, as well as depth of field blurring.

3dfx's swan song, the Voodoo 4 4500, arrived October 19 after several delays – unlike the 4200 and 4800 that were never released. The card was originally scheduled for spring as a competitor to Nvidia's TNT2, but ended up going against the company's iconic GeForce 256 DDR instead, as well as the much better performing GeForce 2 GTS and ATI Radeon DDR.

On November 14, 3dfx announced they were belatedly ceasing production and sale of their own-branded graphics cards, something that had been rumoured for some time but largely discounted. Adding fuel to the fire, news got out that upcoming Pentium 4 motherboards would not support the 3.3V AGP signalling required Voodoo 5 series.

84

90

82

Legion Hardware Reviews

- HIS Radeon HD 7790 iCooler Turbo 1GB
- QNAP TS-469L
- Gigabyte GeForce GTX Titan
- HIS Radeon HD 7850 iPower IceQ Turbo 4GB Crossfire
- Thecus N5550

Downloads and Drivers

Downloads	Drivers
↓ Speccy 1.21.491	
↓ 3DMark for Android 1.0.1-949	
↓ UVK Ultra Virus Killer 5.3.0.0	
↓ VLC Media Player 2.0.6	
↓ LG PC Sync Suite	
▶ More Downloads	


Latest Discussion

- Speaker got dropped, now it gets no power
[7 replies](#) on Audio and Video
- Gaming PC monitor, IPS or 120hz?
[2 replies](#) on Audio and Video
- Used GeForce 9800GT 1GB overheats -- should I buy it?
[14 replies](#) on Audio and Video
- Graphic adapter for Dell XPS M1530?
[4 replies](#) on Audio and Video
- Relation between fps and refresh rate
[14 replies](#) on Audio and Video

[▶ More Recent Discussion](#)

Subscribe to TechSpot

Get **free exclusive content**, learn about new features and breaking tech news.




Voodoo5 5500 AGP box art


The death knell sounded a month later for 3dfx when Nvidia purchased its IP portfolio for \$70 million plus one million shares of common stock. A few internet wits later noted that the 3dfx design team which had moved to Nvidia eventually got both their revenge and lived up to their potential, by delivering the underperforming NV30 graphics chip powering the FX 5700 and FX 5800 cards behind schedule.

The Nvidia vs. ATI Era Begins

Prior to the Voodoo 5's arrival, ATI had announced the Radeon DDR as "the most powerful graphics processor ever designed for desktop PCs." Previews of the card had already gone public on April 25, and only twenty-four hours later Nvidia countered with the announcement of the GeForce 2 GTS (GigaTexel Shader). The latter included Nvidia's version of ATI's Pixel Tapestry Architecture, named Nvidia Shading Rasterizer, allowing for effects such as specular shading, volumetric explosion, refraction, waves, vertex blending, shadow volumes, bump mapping and elevation mapping to be applied on a per-pixel basis via hardware.

The feature was believed to have made it to the previous NV10 (GeForce 256) chip but it remained disabled due to a hardware fault. The GTS also followed ATI's Charisma Engine in allowing for all transform, clipping and lighting calculations to be supported by the GPU. That said, ATI went a step further with vertex skinning for a more fluid movement of polygons, and keyframe interpolation, where developers designed a starting and finishing mesh for an animation and the Charisma core calculated the intervening meshes.





ATI Radeon DDR

The ATI Radeon DDR eventually launched for retail in August 2000. Backed by a superior T&L implementation and support for several of the upcoming DirectX 8 features, the Radeon DDR alongside the GeForce 2 GTS ushered in the use of DVI outputs by integrating support for the interface into the chip itself. The DVI output was more often found on OEM cards, however, as the retail variety usually sported VIVO plugs.

One downside to the Radeon DDR is that boards shipped with their core and memory downclocked from the promised 200MHz and 183MHz, respectively. In addition, drivers were once again less than optimal at launch. There were issues with 16-bit color and compatibility problems with VIA chipsets, but this did not stop the card from dominating the competition at resolutions higher than 1024x768x32. A price of \$399 for the 64MB version stacked up well versus \$349-399 for the 64MB GeForce 2 GTS, which it beat by a margin of 10-20% in benchmarks, and helped ATI maintain its number one position in graphics market share over Nvidia.

Nvidia wasn't doing all that bad for themselves either. The company reported net income of \$98.5 million for the fiscal year on record revenue of \$735.3 million, driven in large part by its market segmentation strategy, releasing a watered-down MX version of the card in June and a higher clocked Ultra model in August. The latter dethroned the Radeon in terms of performance but it also cost \$499. A Pro model arrived in December.

Besides releasing a GeForce 2 card at every price point, from the budget MX to the professional Quadro 2 range, Nvidia also released its first mobile chip in the form of the GeForce2 Go.

As 3dfx was undergoing its death throes in November, Imagination Tech (ex-Videologic) and ST Micro attempted to address the high volume budget market with the PowerVR series 3 KYRO. Typically ranging in price from \$80 to \$110 depending on the memory framebuffer, the card represented good value for the money in gaming at resolutions of 1024x768 or lower. It would have become more popular, had the GeForce2 MX

By the time 2001 dawned, the PC graphics market consisted of a discrete card duopoly, with both

arrived later, or not so aggressively priced at ~\$110.

The KYRO II arrived in April 2001 with a bump in clock speeds compared to the original and manufactured on a smaller 180nm process by ST Micro. But once again the card faced stiff

competition from the GeForce 2 MX. Nvidia rebadged the card as the MX200 and lopped 40% off its price, while adding a higher clocked MX400 card at the same price as the Kyro II.

When PowerVR failed to secure game development impetus for tile based rendering, and ST Micro closed down its graphics business in early 2002, Imagination Technologies moved from desktop graphics to mobile and leveraged that expertise into system on chip graphics. They licenced the Series 5/5XT/6 for use with ARM-based processors in the ultra portable and smartphone markets.

By the time 2001 dawned, the PC graphics market consisted of a discrete card duopoly, with both of them in addition to Intel supplying the vast majority of integrated graphics chipsets.

Meanwhile, Matrox and S3/VIA clung to the margins of traditional markets.

Building on the strides made with the GeForce 2 series, Nvidia unveiled the GeForce 3 on February 27, 2001 priced between \$339 and \$449. The card became the new king of the hill, but it really only came into its own at the (then) extreme resolution of 1600x1200, preferably with full screen antialiasing applied.



Nvidia's stock GeForce 3 card

Initial drivers were buggy, especially in some OpenGL titles. What the new GeForce did bring to the table was DirectX 8, multisampling AA, quincunx AA (basically 2xMSAA + post process blur), 8x anisotropic filtering as well as the unrivalled ability to handle 8xAF + trilinear filtering, and a programmable vertex shader which allowed for closer control of polygon mesh motion and a more fluid animation sequence.

There was also LMA (Lightspeed Memory Architecture) support -- basically Nvidia's version of HyperZ -- for culling pixels that would end up hidden behind others on screen (Z occlusion culling) as well as compressing and decompressing data to optimize use of bandwidth (Z compression).

Lastly, Nvidia implemented load-balancing algorithms as part of what they called the Crossbar Memory Controller, which consisted of four independent memory sub-controllers as opposed to the industry standard single controller, allowing incoming memory requests to be routed more effectively.

of them in addition to Intel supplying the vast majority of integrated graphics chipsets.



Nvidia NV2A inside Microsoft's Xbox

Nvidia's product line later added the NV2A, a derivative of the GeForce 3 with GeForce4 attributes that was used in Microsoft's Xbox game console.

At this point, Nvidia controlled 31% of the graphics market to Intel's 26% and ATI's 17%.

As Nvidia complemented the GF3 line-up with underclocked Ti 200 and overclocked Ti 500 models, ATI hurried to ramp up deliveries of the Radeon 8500. The card was built around the R200 GPU using TSMC's 150nm process (the same used by GeForce 3's NV20). The chip had been

announced in August and was eagerly awaited since John Carmack of id software talked it up saying it would run the new Doom 3 "twice as well" as the GeForce 3.

ATI's official R8500 announcement was no less enthusiastic. But reality kicked in once the card launched in October and was found to perform at the level of the underclocked GF3 Ti 200 in games. Unfinished drivers and a lack of workable Smoothvision antialiasing weighted heavily against the R8500 in its initial round of reviews. By the time the holiday season arrived, a second round of reviews showed that the drivers had matured to a degree and raised the R8500's performance in-between the Ti 200 and the standard GF3.

Spec comparison snapshot

	Core clock (MHz)	Pixel pipelines	Fill rate (Mpixels/s)	Texture units per pixel pipeline	Fill rate (Mtexels/s)	Memory clock (MHz)	Memory bus width (bits)	Memory bandwidth (GB/s)
GeForce3 Ti 200	175	4	700	2	1400	400	128	6.4
GeForce3	200	4	800	2	1600	460	128	7.4
GeForce3 Ti 500	240	4	960	2	1920	500	128	8.0
Radeon 64MB DDR	183	2	366	3	1100	366	128	5.9
Radeon 8500	275	4	1100	2	2200	550	128	8.8

Very competitive pricing and a better all around feature set (2D image quality, video playback, performance under antialiasing) made the card a worthy competitor to the GF3 and Ti 500 nonetheless.

ATI's sales for the year dropped to \$1.04 billion as the company recorded a net loss of \$54.2 million. The company began granting licenses to board partners to build and market graphics boards, while refocusing their resources on design and chip making.

ATI also debuted the Set-Top-Wonder Xilleon, a development platform based on the Xilleon 220 SoC which provided a full processor, graphics, I/O, video and audio for set-top boxes integrated into digital TV designs.

To complement Xilleon, ATI acquired NxtWave

<https://web.archive.org/web/20130412060547/http://www.techspot.com/article/657-history-of-the-gpu-part-3/>[9/12/2015 3:16:17 PM]



ATI Xilleon board

Communications for \$20 million in June 2002. The company specialized in digital signal processing and applications for set-top boxes and terrestrial

digital solutions.

Keeping up with their product launch cycle, Nvidia released the GeForce 4 in February 2002. Three MX parts, three mobile parts based on the MX models, and two performance Titanium models (Ti 4400 and Ti 4600) made up the initial line up -- built on TSMC's 150nm process. The GeForce 4 was effectively ready for release two months earlier but the launch was delayed to avoid eating into GeForce 3 sales over the holiday season.

The MX series cards were intended for the budget segment but they were still largely uninspiring as they were based on the old GeForce 2 architecture. MPEG2 decode added but the cards reverted to DirectX 7.0/7.1 support as the earlier GF2 MX line. Pricing at \$99-179 reflected the reduced feature set.

The Titanium models on the other hand were excellent performers and in some instances managed a 50+% increase in performance over the GeForce3 Ti 500. The Ti 4600 became the performance champ overnight, easily disposing of the Radeon 8500, while the Ti 4200 at \$199 represented the best value for money card.

But then came the [Radeon 9700 Pro](#) and promptly consigned every other card to also-ran status.



ATI Radeon 9700 Pro (FIC A97P)

Developed by a team that had originally formed the core of ArtX, the ATI R300 GPU delivered spectacularly and arrived very promptly. It was the first to bring DirectX 9.0 support, and by extension, the first architecture to support shader model 2.0, vertex shader 2.0, and pixel shader

2.0. Other notable achievements: it was the second GPU series to support AGP 8x -- SiS's Xabre 80/200/400 line was first -- and implementing the first flip-chip GPU package.

ATI complemented the line-up in October by adding a non-Pro 9700 at \$299 for those unable to part with \$399 for the top model. Meanwhile, the cut down 9500 Pro (\$199) and 9500 (\$179) reached down through mainstream market segments, and the FireGL Z1/X1 filled in the \$550-950 bracket for professional graphics. The All-In-Wonder 9700 Pro (\$449) was also added in December.

ATI's sales are likely to have taken a hit when it was found that many cards could be modded to their more expensive counterparts. Examples of this included the ability to turn a 9500 card into a 9700 using its reference board (with the full complement of memory traces), or a 9800 Pro to its XT counterpart. For the latter, a driver patch was made available to check if it would accept the mod, which consisted of soldering in a resistor or using a pencil to tweak the GPU and memory voltage control chip. Hard mods also included upgrading various 9800 models into a FireGL X2, while a patched/Omega driver had the ability to turn a \$250 9800 SE 256MB into a \$499 9800 Pro 256MB.

In addition to discrete graphics, ATI also introduced desktop integrated graphics and chipsets. These included the A3/ IGP 320 meant to be paired with AMD CPUs, RS200/IGP 330 & 340 for Intel chips, as well as the mobile series U1/IGP 320M for AMD platforms and RS200M for Pentium 4-M. All of them were complemented with ATI southbridges, specifically the IXP200/250.

SiS unveiled the Xabre line between the launch of the GeForce4 and the R300. The cards were consistently slower than Nvidia and ATI's offerings at the same price points, and were handicapped by the lack of vertex shader pipelines. This translated into a heavy reliance upon drivers and game developers to get the most out of software emulation, thus keeping SiS in the margins of desktop discrete 3D graphics.

The Xabre line also implemented "Turbo Texturing", where framerates were increased by drastically reducing texture quality, and lacked anisotropic filtering . All this did little to endear reviewers to the cards.

The Xabre line was the last under the SiS banner, as the company spun off its graphics division (renamed XGI) and merged with Trident Graphics a couple of months later in June.

The first of Nvidia's FX series arrived on January 27, 2003 with the infamous "Dustbuster" FX 5800 and the slightly faster (read: less slow) FX 5800 Ultra. When compared to the reigning champ, the ATI Radeon 9700 Pro (and non-Pro), the FX was much louder, it delivered inferior anisotropic filtering (AF) quality and antialiasing (AA) performance, and was overall much slower. ATI was so far ahead that a second-tier Radeon 9700 card launched five months earlier comfortably outperformed the Ultra, and it was \$100 cheaper (\$299 vs \$399).

About flip-chip GPU packages: Previous generations of graphics chips and other ICs used wire-bonding mounting. With this method, the chip sits on the board with the logic blocks sitting under the metal layers whose pads would be connected by thin wires arranged around the edges of the chip down to solder balls or pins on the underside. Flip-chip does away with the wire component through contact points (usually soldered in a ball grid array) directly on the "top" of the chip, which is then inverted, or "flipped" so that the solder points directly contact the substrate or circuit board. The chip then undergoes localised heating (reflow) to melt the solder that then forms the connection with the underlying contact points of the board.

The 3dfx design team which had moved to Nvidia got both their revenge and lived up to their potential, by delivering the underperforming NV30 graphics chip

behind schedule.

The NV30 chip was supposed to debut in August, around the same time as the Radeon 9700, but ramping problems and high defect rates on TSMC's Low-K 130nm process held Nvidia back. Some circles also argued that the company was strapped for engineering resources, with more than a few tied up with the NV2A Xbox console chip, the [SoundStorm APU](#), as well as the motherboard chipsets.

Looking to move things forward Nvidia undertook a project to have several FX series chips fabricated on IBM's more conventional Fluorosilicate glass (FSG) low-K 130nm process.

ATI refreshed its line of cards in March, starting with the 9800 Pro, featuring a R350 GPU that was basically an R300 with some enhancements to the Hyper-Z caching and compression instruction.

The RV350 and RV280 followed in April. The first of these, found inside the Radeon 9600, was built using the same TSMC 130nm low-K process that Nvidia had adopted. Meanwhile, the RV280 powering the Radeon 9200 was little more than a rebadged RV250 of the Radeon 9000 with AGP 8x support.

The same month saw ATI and Nintendo [sign a technology agreement](#) that would eventually lead to the Hollywood GPU for the Nintendo Wii console. ATI added a second console coup in August, when Microsoft [awarded the Xbox 360 GPU contract to them](#).


A scant three and a half months after the inglorious debut of the FX 5800, Nvidia took another shot with the NV35 (FX 5900 and FX 5900 Ultra). The new Detonator FX driver greatly improved AA and AF, almost matching ATI's solution in terms of quality. However the 5900 achieved what the 5800 could not. It knocked ATI's Radeon 9800 Pro from its spot as the fastest card around, although at \$499 apiece, few would actually take advantage of this.



Xbox 360 GPU (ATI C1 / Xenos)

As expected, ATI regained bragging rights in September with the release of the 9800 XT. Superior driver support – mainly with some DX9 games – also made the XT a better overall card than Nvidia's counterpart, ensuring that ATI ended the year with the performance crown. The 9700 Pro remained the standout mainstream board, while the FX 5700 Ultra at \$199 won the sub-\$200 price segment.

ATI bounced back with a \$35.2 million profit in 2003 after posting a \$47.5 million loss in 2002. A good chunk of this came from higher selling prices for the dominant 9800 and 9600 cards. Meanwhile, Nvidia retained 75% of the DirectX 9 value segment market, thanks to the popularity of the FX 5200.





Source DirectX 9.0 Effects Trailer, shown during ATI's presentation of the Radeon 9800 XT and 9600 XT

The newly formed XGI launched the Xabre successor in a staggered release between September and November. Renamed Volari, the card line-up ranged from the \$49 V3 to the dual GPU Duo V8 Ultra. The V3 was virtually a rebrand of Trident's Blade XP4 and a DX 8.1 part, while the rest of the series (V5 and V8) was developed from the previous SiS Xabre and featured DX9.0 support.

For the most part, all of the models underdelivered, with the exception of the entry-level V3 which offered performance equal to the GeForce FX 5200 Ultra and and Radeon 9200. The Duo V8 Ultra was priced ~20% higher than the Radeon 9800 Pro 128MB, yet delivered performance on par or lower than the 9600XT.

Another company making a comeback into desktop graphics was S3. Unfortunately, the buying public now generally saw desktop graphics as a two horse race – and S3 wasn't one of the two.

XGI's Volari line lingered on with the 8300 in late 2005, which was more or less on par with the Radeon X300SE/GeForce 6200 at \$49, as well as the Z9/Z11 and XP10. The company was reabsorbed back into SiS in October 2010.

Another company making a comeback into desktop graphics was S3. After the graphics division was sold to VIA for \$208 million plus the company's \$60 million debt, the restructured venture concentrated primarily on chipset projects.

DeltaChrome desktop cards were announced in January, but in time-honoured S3 fashion, the first S4 and S8 models didn't start appearing in the retail channel until December. The new cards featured most of the new must-haves of 2003;

DirectX 9 support, 16x AF, HD 1080p support, and portrait-mode display support.

Unfortunately, the buying public now generally saw desktop graphics as a two horse race – and S3 wasn't one of the two. While S3 was looking to keep competitive, ATI and Nvidia were driving each other to achieve ever-increasing levels of performance and image quality.

The DeltaChrome was succeeded by the GammaChrome in 2005.

Nvidia and ATI continued in 2005 their staggered launches. The former launched its first GDDR3 card in March as the FX 5700 Ultra, followed by the GeForce 6 series with the high-end 6800 range. The initial line up comprised the 6800 (\$299), GT (\$399), the Ultra (\$499), and an overclocked variant known as the Ultra Extreme (\$549) to counter ATI's X800 XT Platinum Edition. The latter was sold by a select band of add-in board partners.

The 6800 Ultra 512MB was added on March 14 2005 and sold for the unbelievable price of \$899 -- BFG added an overclocked version for \$999. The midrange was well catered for with the 6600 series in September.

Nvidia's feature set for the 6000 series included DirectX 9.0c support, shader model 3.0 (although the cards were never able to fully exploit this), Nvidia's PureVideo decode and playback engine,

and SLI support -- the multi-GPU performance multiplier IP that was acquired from 3dfx.



Reintroducing an old feature: SLI

Where the 3dfx implementation resulted in each processing unit being responsible for alternate line scans, Nvidia handled things in a few different ways. The company implemented split frame rendering (SFR), in which each GPU rendered the top or bottom half of the frame, alternate frame rendering (AFR) so GPUs rendered frames in turn, and in some cases the driver just disabled SLI depending on whether the game supported the feature. This last feature was a hit-or-miss early in driver development.

While the technology was announced in June, it required a motherboard with an nForce4 chipset to enable multi-GPU setups, and these didn't start reaching the retail channel in numbers until late November. Adding fuel to the fire, initial driver releases were sporadic (at best) until into the following year.

Reviews at the time generally mirrored current performance, showing that two lower tier cards (like the 6600 GT SLI which could be had for \$398) generally equalled one enthusiast card at lower resolutions and image quality. At highest resolutions and with antialiasing applied, however, single card setups still gained the upper hand. SLI and ATI's CrossFire performance was as erratic then as it sometimes is now, running the full gamut from perfect scaling to not working at all.

Nvidia's board partners immediately saw marketing opportunities with the re-invented tech, with Gigabyte offering a dual 6600 GT SLI card (the 3D1), followed by a dual 6600 (3D1-XL), and

While Nvidia's SLI was announced in June 2004, the required nForce4 motherboards didn't hit the retail channel in numbers until November, and initial driver releases were sporadic until

into the following year.

the 6800 GT (3D1-68GT). These cards not only required an nF4 chipset but also a Gigabyte branded motherboard as well.

Of the high-end single GPU cards, the 6800 Ultra and X800 XT/XT PE were fairly evenly matched, both in price and performance. But they weren't without their issues. The latter arrived in May and suffered supply constraints throughout its entire production life, while Nvidia's flagship 6800 Ultra was extremely late arriving in August and suffered supply constraints too depending on distribution area, since the card was only made available by a percentage of board partners.

The 6800 GT generally bested the X800 Pro at \$399, while the 6600 GT cleaned up in the \$199 bracket.

Intense competition with Nvidia that year didn't have an adverse effect on ATI's bottom line, as profit peaked at \$204.8 million for the year from nearly \$2 billion in revenue.

One quirk associated with the well-received 6600 GT was that it initially launched as a PCI Express card, at a time when PCI-E was an Intel-only feature for motherboards designed for Pentium 4 processors. These chips generally lagged in gaming performance behind AMD's offerings, which of course used the AGP data bus.

Nvidia's 7000 series started rolling off the assembly lines well before the 6000 series had completed its model line-up. The 7800 GTX arrived a full five months before the reduced bill of materials (BoM) 6800 GS saw the light of day. The first iteration of the 7800 series was based around the G70 GPU on TSMC's 110nm process, but quickly gave way to the G71-based 7900 series, made on TSMC's 90nm process.

While the naming convention changed from "NV" to "G", the latter were architecturally related to the NV40 series of the GeForce 6000. And while only fractionally larger than the NV40-45 at 334mm², the G70 packed in an extra eighty million transistors (for a total of 302 million), adding a third more vertex pipelines and 50% more pixel pipelines. In most cases, the G70 was superseded within nine months, and in the case of the GS and GTX 512MB, the figure was 3 and 4 months respectively.

At the entry level, the 7100 GS continued the use of TurboCache (the ability for the board to use some system memory), which was introduced with the previous generation GeForce 6200 TC.



Nvidia GeForce 7800 GTX

At the other end of the spectrum, the 7800 GTX 256MB hit retail on June 22 with an MSRP of \$599, though its actual street price was higher in many instances. ATI wrested the single-GPU crown back with the X1800 XT, but Nvidia countered with a 512MB version of the 7800 GTX thirty-five days later and promptly regained the title.

Two months later, ATI launched the X1900 XTX, which traded blows with Nvidia's flagship. This particular graphics horsepower race resulted in both cards being priced at \$650. One spinoff of the cards moving to a 512MB frame buffer was that gaming at 2560x1600 with 32-bit color and a high level of image quality enabled was now possible via dual link DVI.

ATI announced their multi-card Crossfire technology in May 2005 and made it available in September with the launch of the Xpress 200 Crossfire Edition chipset, and X850 XT Crossfire Master board. Due to a single-link TMDS, resolution and refresh rates were initially limited to 1600x1200 @60Hz, but a dual-link TMDS for 2560x1600 would soon replace it.



ATI's original CrossFire design required using an external Y cable

Unlike Nvidia's solution of two identical cards communicating via a bridge connector, ATI implemented a master card with TMDS receiver, which accepted input from a slave card via external dongle and a Xilinx compositing chip.

Like Nvidia's SLI, CrossFire offered alternative frame rendering (AFR) and split frame rendering (SFR), but also a rendering technique called SuperTiling. The latter offered a performance increase in certain applications, but it did not work with OpenGL or support accelerated geometry processing. Also like SLI, Crossfire faced its share of driver-related troubles.

ATI intended to have their R520 based cards – their first to incorporate Shader Model 3.0 – ready by the June-July timeframe, but the late discovery of a bug in the cell library forced a 4 month delay.

Initial launches comprised the X1800 XL/XT using the R520 core, the X1300 budget cards using the RV515 with essentially one quarter of the graphics pipelines of the R520, and the X1600 Pro/XT based on the RV530, which was similar to the RV515 but with a higher shader and vertex pipeline-to-TMU and ROP ratio.

Due to the initial delay with the R520, the GPU and its derivations were being replaced a scant three and a half months later by the R580-based X1900 series which used TSMC's new 80nm process. Continuing with the roll out, half the graphics pipeline resources went into the RV570 (X1650 GT/XT and X1950 GT/Pro), while a shrunk RV530 became the RV535 powering the X1650 Pro as well as the X1300 XT.

ATI's revenue rose to a record \$2.2 billion for the year, the highest in the company's history, aided by shipments of Xenos GPUs for the Xbox 360. Net profit, however, slumped to \$16.9 million.

ATI's revenue rose to a record \$2.2 billion in 2005, the highest in the company's history, aided by shipments of Xenos GPUs for the Xbox 360. Net profit, however, slumped to \$16.9 million.

By this stage, any graphics card launch not based on an Nvidia or ATI GPU was received with a certain amount of curiosity, if not enthusiasm. Such was the scene when S3's overhauled graphics line-up debuted in November.

The Chrome S25 and S27 promised good gaming performance based on their high clocks, but delivered a mostly sub-par product. Initial pricing at \$99 (S25) and \$115 (S27) put the cards in competition against Nvidia's 6600/6600GT and ATI's X1300Pro/X1600Pro, but neither S3 card stood up to the competition in any meaningful way, aside from power consumption. That slight advantage evaporated as ATI/AMD and Nvidia

addressed the HTPC and entry-level market segment, effectively killing S3's subsequent Chrome 400 and 500 series.

An added issue for S3 was that the cost of building the cards resulted in razor thin profits. The company needed high volume sales in a market dominated by two vendors. HTC were to acquire S3 in July 2012 for \$300 million, a move originally seen as leverage in HTC's and S3's separate legal disputes with Apple.

Nvidia and ATI continued to hog the press coverage in 2006.

ATI acquired Macrosynergy, a Shanghai based design and engineering centre with personnel working in California and previously part of the XGI group. Then in May the company bought BitBoys in a \$44 million deal.

Meanwhile, Nvidia's first foray into dual-GPU single board products came in March, following in the footsteps of ATI, 3dfx, and XGI. The 7900 GX2 would sandwich two custom boards essentially carrying a couple of downclocked 7900 GTXs. But Asustek didn't wait around for Nvidia's dual-GPU solution, however, and released its own take as the [Extreme N7800GT Dual](#) (\$900, 2000 units built), which paired two 7800 GT GPUs instead.

This card started Asus interest in limited edition dual-GPU boards, and possibly hardened Nvidia's attitude towards board partners', as Asustek products took the spotlight from their reference models at launch.

In the higher volume mainstream market, the 7600 GT and GS both provided solid performance and remarkable longevity, while ATI's X1950 XTX and Crossfire ruled the top end enthusiast benchmarks for single GPU cards. The X1900 XT and GeForce 7900 GT were fairly evenly matched in the upper mainstream bracket.



ATI's David Orton and AMD's Hector Ruiz officially announce the historic merger

After twenty-one years as an independent company, ATI was bought out by AMD on October 25 2006 for a total price of \$5.4 billion – split between \$1.7 billion from AMD, \$2.5 billion borrowed from lending institutions, 57 million AMD shares and 11 million options/restricted stock units valued at \$1.2 billion. At the time of the buy out, around 60-70% of ATI's chipset/IGP revenues were accrued from a partnership with Intel based motherboards.

Two weeks after the ATI buy-out, Nvidia ushered in

the age of unified shader architectures for PC graphics.

With a large part of Intel's IGP chipset market moving to Nvidia, market share dropped dramatically. The logic behind the buy was a seemingly quick path to GPU technology, rather than use the \$5.4 billion to develop AMD's own IP and add licenced technology where needed. At the time, AMD was aiming at the quick introduction of [Torrenza](#) and the associated Fusion projects.

Two weeks after the ATI buy-out, Nvidia ushered in the age of unified shader architectures for PC graphics. ATI's Xenos GPU for the Xbox 360 had already introduced the unified architecture to consoles.

This article is the third installment on a series of four. Next week we'll wrap things up, following the development of Radeon products under AMD's wing, the continued rivalry between GeForce and Radeon CPUs, the transition toward stream processing, and what the present a near future holds for graphics processors.

Part 1: (1976 - 1995) [The Early Days of 3D Consumer Graphics](#)

Part 2: (1995 - 1999) [3Dfx Voodoo: The Game-changer](#)

Part 3: (2000 - 2006) [The Nvidia vs. ATI Era Begins](#)

Part 4: (2006 - Present) [The Modern GPU: Stream processing units a.k.a. GPGPU](#)

- [22 comments](#) - Read user comments, write your own
- Share this:
- Contact the author: [Graham Singer](#)
- Check out [other recent features](#) published at TechSpot

[Back to TechSpot](#) ▶

Main Sections

[Technology News](#)
[Reviews](#)
[Guides](#)
[Product Finder](#)
[Downloads](#)
[Drivers](#)

Community

[TechSpot Forums](#)
[Recent Topics](#)
[Recent Comments](#)
[Community Activity](#)

Useful Resources

[Trending Topics](#)
[Featured on TechSpot](#)
[PC Buying Guide](#)
[Hot Downloads](#)
[RSS Feeds](#)

About TechSpot

[About Us](#)
[Advertising](#)
[News Archive](#)
[The TechSpot Blog](#)

TechSpot | Technology News and Analysis
© 2013 TechSpot, Inc. All Rights Reserved
TechSpot is a registered trademark
[Terms of Use](#) | [Privacy Policy](#)