

# DirectX<sup>®</sup> 10.1

Enabling breakthrough graphics on the ATI Radeon<sup>™</sup> HD 3800 series

---



# Table of Contents

<b>Introduction</b> .....	<b>1</b>
<b>The Evolution of DirectX</b> .....	<b>1</b>
<b>Next-Generation Image Quality</b> .....	<b>4</b>
Lighting and Shadows.....	4
Global Illumination.....	5
Ambient Occlusion .....	7
<b>Anti-aliasing Improvements</b> .....	<b>8</b>
Custom Anti-Aliasing.....	9
<b>Tighter Specification</b> .....	<b>10</b>
New texture format requirements.....	10
New multi-sample anti-aliasing requirements .....	11
Higher precision requirements .....	11
<b>Conclusion</b> .....	<b>11</b>

## Introduction

Microsoft® DirectX® 10.1 is the latest application programming interface from Microsoft that unlocks the state of the art in GPU technology, represented by the ATI Radeon™ HD 3800 series. Key features of DirectX 10.1 include an updated shader model, improved anti-aliasing support, more flexible data access, and tighter specifications for better application compatibility. These features will enable exciting new techniques, such as real-time global illumination, that will define the future direction of interactive 3D graphics.

DirectX 10 was one of the most significant updates to the API since its inception. DirectX 10.1 represents an evolutionary update that addresses some of the limitations identified after the specification was finalized. DirectX 10.1 support will be coming to the Windows Vista™ operating system with the release of a service pack planned for early 2008.

For many years, ATI was at the forefront of DirectX technology development, working actively with Microsoft to identify and implement new graphics features. The merger of ATI with AMD in 2006 not only continues this tradition, but also enables new possibilities for platform-level synergies between the GPU, CPU, and system chipset.

The new ATI Radeon HD 3800 series of GPUs are the first to be designed for DirectX 10.1, as well as other cutting edge technologies, including PCI Express 2.0, Unified Video Decoder (UVD), hardware accelerated tessellation, and power efficient 55nm transistor design. The products are perfectly positioned to deliver a superior gaming experience at compelling prices.

This paper describes the new features of DirectX 10.1, and provides a number of examples showing how they can be put to use. To help illustrate these techniques, AMD has created an accompanying interactive game called PingPong. This game makes extensive use of the DirectX 10.1 features on ATI Radeon HD 3800 series products to highlight the benefits in a fun and informative way.

## The Evolution of DirectX

DirectX 10.1 maintains the overall structure and programming model of DirectX 10, while providing numerous enhancements. The vertex, geometry, and pixel shader instruction sets have been updated to Shader Model 4.1.

The new features of DirectX 10.1 can be divided into three general categories: new shading and texturing capabilities, anti-aliasing improvements, and tighter specifications. The following table highlights some of the key features in each of these categories, as well as some of the benefits they provide.

	FEATURE	FUNCTION	BENEFITS
Shader and Texture Improvements	<i>Cube Map Arrays</i>	Allow reading and writing of multiple cube maps in a single rendering pass	Efficient Global Illumination in real time for complex, dynamic, interactive scenes Enable many ray trace quality effects including indirect lighting, color bleeding, soft shadows, refraction, and high quality glossy reflections
	<i>Separate Blend Modes per-MRT</i>	Allows pixel shaders to output to multiple buffers (MRTs), each with their own blend mode	Efficient deferred shading for improved performance in complex 3D scenes
	<i>Increased Vertex Shader Inputs &amp; Outputs</i>	Doubled from 16 to 32 128-bit values per shader	Improved performance for complex shaders
	<i>Gather4</i>	Allows a 2x2 block of unfiltered texture values to be fetched in place of a single bilinear filtered texture lookup	Higher quality shadows with improved filtering Fast procedural noise computation to add more visual variety to 3D scenes Higher order filtering for high quality fluid simulation on the GPU Improved performance for Stream Computing applications
	<i>LOD instruction</i>	New shader instruction that returns the level of detail for a filtered texture lookup	Custom texture filtering techniques for optimized performance and quality Fast parallax occlusion mapping for improved 3D surface detail

	FEATURE	FUNCTION	BENEFITS
Improved Anti-Aliasing	<i>Multi-sample buffer reads and writes</i>	Allow individual color and depth samples in a multi-sample buffer to be accessed directly by a shader	Custom edge detect filters for high quality anti-aliasing with optimized performance and reduced memory footprint
	<i>Pixel Coverage Masks</i>	Enable programmable anti-aliasing in a pixel shader	Faster adaptive anti-aliasing Improved anti-aliasing quality with HDR rendering Improved anti-aliasing compatibility and performance with deferred shading High quality volumetric rendering for atmospheric effects High quality depth of field post-processing effects
	<i>Programmable AA Sample Patterns</i>	Allows programmers to define their own sample patterns for each pixel	Temporal anti-aliasing Improved image quality for multi-GPU anti-aliasing
Tighter Specification	<i>FP32 filtering required</i>	Filtering of 128-bit floating point texture formats now a requirement instead of an optional feature	Encourages use of these high precision data formats by ensuring hardware compatibility
	<i>Int16 blending required</i>	Blending of 64-bit integer pixel formats now a requirement instead of an optional feature	
	<i>Minimum 4x MSAA support required</i>	Multi-sample anti-aliasing with at least 4 samples per pixel must be supported for all 32-bit and 64-bit pixel formats	Ensures anti-aliasing can behave identically on all DirectX 10.1 GPUs
	<i>Standardized AA sample patterns</i>	Pre-defined sample locations for 2x/4x/8x/16x AA modes that hardware must support	Encourages support for anti-aliasing by improving consistency
	<i>Increased precision for floating point operations</i>	0.5 ULP precision required for all floating point math (add/subtract/multiply/divide) and blending operations	Eliminates rounding errors Matches IEEE standard requirements for these operations

## Next-Generation Image Quality

The new features of DirectX 10.1 will enable many new rendering techniques that can enhance the realism and quality of 3D games that use them. Rather than covering all of the possibilities here, we will focus on some of the most promising ones, which deal with light and shadows.

### *Lighting and Shadows*

One discontinuity between real time and non-real time rendering today is the latter's use of global illumination for physically based, realistic lighting. The process of determining the brightness and color of light reflecting off any given point is critical to giving rendered scenes a sense of depth, and helping viewers gauge the location and movement of objects in 3D space.

For example, indirect or bounced lighting is used in the ATI Radeon PingPong game to help the player visually identify the positions of objects. When one object gets close to another brightly colored object, bounced/reflected light causes color bleeding, and this provides an important visual cue for the closeness of the two objects. So when a white ball is approaching a red wall, the part of the ball facing the wall will take on a reddish hue from the light bouncing off the wall. The human perceptual system processes this information and understands that the ball is near the wall. This gives the player a much richer experience, and makes it easier to make strategic gameplay decisions.

There are a number of different lighting methods used in 3D rendering today. These include rasterization techniques, such as light/shadow mapping, which are prevalent in real-time applications and interactive games; ray tracing techniques, which are commonly used for offline rendering applications such as CG films; and other techniques such as radiosity, which attempt to combine the benefits of rasterization and ray tracing. On today's microprocessors, the performance of the latest ray tracing techniques still trails far behind that of rasterization.

Light/shadow mapping works by rendering a version of the image from the point of view of each light source to one or more textures. For each pixel, lookups into these textures are performed to determine the amount and color of light contributed by each, and also whether or not each light source is visible from that pixel (if not, it is considered shadowed). Surface shader programs then evaluate and blend all the contributions, taking into account material properties, to determine the final pixel color.

This lighting method is straightforward for capturing standard diffuse lighting, but handling reflections generally requires separate cube maps to be rendered. It is relatively fast and handles dynamic scenes well, but reflection and shadow quality can suffer from limited resolution of light/shadow maps. It also doesn't handle large numbers of point or area light sources well, and it doesn't capture indirect lighting.

Ray tracing works by casting rays out from the viewpoint toward every visible pixel on the screen. It then determines the first point at which each ray intersects an object, and casts rays into each of those points from all directions. A shader is executed at each point to determine the amount and color of light reflecting off of it. This process is good at handling reflections, refraction, and shadows, and it can be repeated as necessary for multiple bounces to get indirect lighting.

However, ray tracing requires the processing of billions of rays per second to achieve real time frame rates at typical display resolutions (in the 1 to 4 megapixel range). It also needs complex data structures to store the locations of all objects in the scene, which makes it inefficient for dynamic scenes (since they require these data structures to be updated each frame). Finally, ray tracing is inefficient when it comes to handling shading, since complex surface shaders must be executed repeatedly for each pixel to handle many incoming rays.

Since ray tracing is prohibitively expensive on today's hardware, most games today handle indirect lighting with a "constant ambient term" (i.e. uniform illumination with no particular source that contributes to the entire scene). This is a very rough approximation at best, and it doesn't handle color bleeding or shadows.


Radiosity is another lighting and shadowing technique used in some recent games and offline rendering, that combines some of the benefits of both light/shadow mapping and ray tracing. It works by using a method similar to ray tracing to pre-compute and store lighting values for a scene. The stored values are then looked up at render time to allow fast indirect lighting. However, the extensive pre-computation required means that radiosity doesn't work well for dynamic scenes with moving or changing light sources.

## Global Illumination

Global illumination is a rendering technique that combines the benefits of light/shadow mapping with indirect lighting and support for practically unlimited dynamic light sources, as well as realistic reflections and soft shadows. With DirectX 10.1, developers can use cube map arrays and geometry shaders to implement global illumination efficiently in real time, even with thousands of physically modeled objects in a complex, interactive scene.

This technique works by dividing the scene into a 3D array of cubes. For each cube face, a simple version of the scene is rendered from the point of view of someone at the center of the cube looking outward (referred to as a "light probe"). The resolution and detail of these images generally doesn't need to be as high as that of the final image, but is easily scalable according to the level of accuracy and performance required. When the six faces of each cube are complete, they are stored in a cube map texture.

The next step is to convert each cube map into a compressed spherical representation using spherical harmonics. With this representation, it becomes possible to determine the amount and color of light falling on any point in the cube from any and all directions with a few simple math operations. For points between the light probes, lighting values can be interpolated between values taken from adjacent cube maps.



*Illustration of light probes used to capture global illumination.*

The cube map textures can also be used to create high quality reflections on shiny or glossy objects in the scene. In the PingPong game, looking closely you can see that each of the thousands of balls has a reflection of the environment on its surface.

This method of calculating lighting is highly scalable. The level of quality can be controlled by changing the size and number of cubes used, as well as the level of detail in the cube face images. With DirectX 10.1 cube map arrays, large numbers of cube maps can be rendered to and sampled from simultaneously in parallel, making this technique particularly effective on single GPU or multi-GPU systems.

If required for extremely complex scenes, multiple light bounces can be simulated by repeating this process and accumulating the results before rendering the final image. Furthermore, the quality of the lighting can be decoupled from the number of objects or polygons in the final scene; in the case of the PingPong game, the number of balls that can be rendered is limited only by the speed of the physics simulation.




*Scene from the ATI Radeon PingPong game. On the left, global illumination is being used to provide realistic lighting and shadows with thousands of physically simulated objects that the player can interact with. The right side of the image shows the scene without global illumination. Using ray tracing techniques to render highly dynamic scenes like this would be practically impossible to achieve in real-time on existing consumer hardware.*




## Ambient Occlusion

Ambient occlusion is a shading technique that determines how much light can reach a given point from all directions, and how much is occluded by intervening objects. It is most often calculated by casting rays in every direction from the surface. Rays which reach the “background” (typically defined as walls, floors, sky, etc.) increase the brightness of the surface, whereas a ray which hits any other object contributes no illumination. As a result, points surrounded by a large amount of geometry are darker than points with little surrounding geometry. This is an important factor because it captures local visibility changes beyond what is provided by the global illumination alone.



*Scene from the ATI Radeon PingPong game with ambient occlusion enabled on the left, and disabled on the right. Note how the soft shadows resulting from the ambient occlusion make it much easier to identify the positions of the balls relative to the walls, floor, and each other.*

In the PingPong game, a dynamic ambient occlusion technique adds significantly to the effectiveness of the global illumination. It gives the balls soft-edged shadows as they approach each other and as they approach the walls of the scene. These contact shadows are an important visual cue for resolving the positions of objects in 3D space. This type of shadowing would be impractical to achieve using traditional shadow mapping or shadow volume techniques, given the large number of moving objects present in this game.



*Soft-edged shadows rendered using ambient occlusion.*

The new Gather4 feature of DirectX 10.1 can help further improve the quality of ambient occlusion shadows (and traditional shadow maps as well). These techniques use special textures to store shadow data, with a single data value per location (unlike standard color textures that store red, green, blue, and alpha values for each location). Gather4 takes advantage of this by sampling 4 values from these textures in the same time it takes to sample a single value from a standard texture. These extra samples can be used to implement more sophisticated filtering for smoother shadows.

In summary, DirectX 10.1 enables a breakthrough in lighting quality with real-time global illumination effects, including indirect lighting, color bleeding, soft shadows, reflection and refraction.


## Anti-aliasing Improvements

Jagged, shimmering edges caused by aliasing can be very distracting in rendered images. These artifacts are a consequence of insufficient rendering and/or display resolution. They can be reduced or eliminated using anti-aliasing (AA) filtering techniques, which work by taking multiple samples per pixel and blending them together. The trick to doing this effectively is to choose sample points and weights so as to maximize edge smoothing without blurring out detail, while maintaining good performance.

The most commonly used anti-aliasing technique today is multi-sample anti-aliasing (MSAA), but this only works on polygon edges; it doesn't address texture aliasing or shader aliasing. The ATI Radeon™ HD 2000 series introduced custom filter anti-aliasing (CFAA), which enabled sophisticated programmable filters. The latest Catalyst drivers support 4 different filter types that can be applied to almost any game using DirectX® 9 or earlier. The most advanced of these filters is the edge detect filter, which detects and anti-aliases all edges (including those in textures and shaders).

## Custom Anti-Aliasing

DirectX 10.1 allows custom anti-aliasing filters to be implemented with pixel shaders. Custom filters can offer improved quality in certain cases where standard MSAA can have issues, such as with HDR lighting and deferred shading techniques. All DirectX 10.1 compatible hardware must support a minimum of 4x MSAA. Also, the specification now includes some pre-defined AA sample patterns, in contrast to earlier versions of DirectX where they were left entirely up to each particular GPU to define.



*Edge pixels, detected by the edge detect AA filter and highlighted in red in this image from the ATI Radeon PingPong game, are assigned more samples than other pixels in the image. The inset at the lower right shows how this anti-aliasing removes jagged edges and shimmering artifacts from the edges of the balls; the effects are more pronounced when the image is in motion.*

A new feature of DirectX 10.1 allows all AA buffers to be accessed directly by shaders. Previously, it was only possible to access multi-sampled color buffers; it was impossible to access information from a depth buffer for each sample individually. This allows developers to implement more advanced custom AA techniques using a combination of shaders and dedicated hardware, much like ATI Radeon HD GPUs do today with CFAA.

ATI Radeon GPUs also introduced support for adaptive AA, which smooths out jagged edges in partially transparent textures (such as those used to render foliage and chain-link fences). DirectX 10.1 expands on this capability by introducing sample coverage masking, which provides control over the specific sample locations where pixel shaders are executed.

This allows developers to extend the adaptive AA technique to address more types of aliasing artifacts. Custom sample patterns can also be specified to complement the basic set that must be supported by all compatible hardware. Many of these capabilities were already present in previous generations of ATI Radeon GPUs, but the DirectX 10.1 allows them to be exposed directly to developers for the first time.

In summary, DirectX 10.1 finally gives developers the tools they need to eliminate all types of aliasing artifacts from interactive real-time games, and can deliver a major step forward in image quality.

## Tighter Specification

GPU compatibility issues have historically been a significant roadblock that has slowed the adoption of new 3D features by developers. These arise when certain elements of the programming interface they are using are interpreted slightly differently by different GPUs, causing unexpected performance drops, image quality problems, error messages, or even crashes. Rather than spend time working around these problems, many developers preferred to avoid them entirely by targeting lowest common denominator GPU feature sets.

DirectX 10 made major strides toward eliminating these issues by more tightly defining the required GPU behavior for each function and instruction, and greatly reducing the number of optional features that might be present on one DirectX 10 GPU but not present on another. Many of the improvements in DirectX 10.1 were introduced to take the API even further down this path.

### *New texture format requirements*

One obstacle that has prevented developers from using the higher precision texture and output formats in recent versions of DirectX has been the limits on what operations each GPU can perform on them. DirectX 10.1 improves this by requiring all compatible GPUs to support texture filtering of 32-bit floating point formats, and blending operations on 16-bit integer formats.

## *New multi-sample anti-aliasing requirements*

Multi-sample anti-aliasing is a well established technique for improving image quality. However, over the years, many enhancements have been made to the basic technique. Since these modifications can yield significantly different output on different GPU models, developers often tended to avoid supporting them directly in their games. DirectX 10.1 mandates a minimum anti-aliasing quality requirement as well as a set of pre-defined sample patterns. This ensures that a consistent, high quality level for anti-aliasing is supported on all compatible GPUs, while still supporting new techniques on individual GPUs.

## *Higher precision requirements*

All data formats have a limited amount of precision they can support, which depends on the number of bits available. However, operations done on these formats do not necessarily produce output that takes full advantage of all the available precision; in some cases, approximations are used that can cause rounding errors in the least significant bits of the output. This practice can cause unpredictable behavior in some cases (for example, when errors build up due to iterating an operation many times).

DirectX 10.1 addresses these issues by mandating that basic math operations on floating point values up to 32-bits take advantage of the full precision available, thus ensuring identical results down to the last bit on all compatible GPUs (and even CPUs as well).

## **Conclusion**

DirectX 10.1 offers incremental improvements to the graphics programming interface that address limitations of DirectX 10, and unlock new graphical techniques that will take the quality of 3D graphics to the next level in 2008 and beyond.

One of the key advantages of these improvements is support for real-time global illumination, delivering lighting and shadow quality in real-time that can match that of ray tracing techniques used in CG films. Other advantages include improved anti-aliasing techniques to clean up distracting shimmering artifacts, and tighter specifications for improved compatibility. The ATI Radeon HD 3800 series GPUs are the first products to bring these features and benefits to the PC.

## Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2007 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ATI, the ATI logo, Avivo, Catalyst, Radeon, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Microsoft, Windows, Vista, and DirectX, are registered trademarks, of Microsoft Corporation in the United States and/or other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners.