

ATI Technologies

WHITE PAPER



SMOOTHVISION™

Visit ATI.com

• Introduction

Anti-aliasing has quickly become a necessary tool in the creation of high-quality computer-generated images. Unfortunately, many anti-aliasing solutions up to this point have been limited to a small number of possible options, all of which usually cause a significant drop in performance, or sacrifice image quality at the expense of maintaining a reasonable level of performance. With the introduction of ATI's SMOOTHVISION™ anti-aliasing, high quality anti-aliased graphics can now be achieved while maintaining a high level of performance, enabling computer games to look more detailed and realistic.

With SMOOTHVISION™, users are able to choose from a variety of different anti-aliasing modes to best meet their image quality or performance needs. SMOOTHVISION™ supports a High-Performance anti-aliasing mode and a High-Quality anti-aliasing mode. Both of these modes give users the ability to choose between 2x, 3x, 4x, 5x, and 6x sampling, allowing for a total of 10 different anti-aliasing settings. This allows users to quite easily select both the desired level of visual quality and the corresponding performance level.

SMOOTHVISION™ makes use of a jittered sampling table, allowing for a variety of different sample positions. Using a variable sampling pattern eliminates problems associated with redundant sampling methods implemented in other graphics accelerators. Redundant sampling can cause image artifacts that are quite noticeable along edges, due to the repetitive nature of the sample colors. Using jittered sampling patterns enables highly realistic anti-aliased edges.

The High-Quality SMOOTHVISION™ setting provides users with unparalleled image quality and jagged edge removal, while maintaining very playable frame rates. Using the High-Quality 2x SMOOTHVISION™ setting provides superior image quality to current 4x multi-sampling techniques, while maintaining a significant performance advantage. The High-Performance SMOOTHVISION™ mode offers users quality anti-aliased images superior to the competing anti-aliasing solution, at a high level of performance.

In order to better appreciate the need for anti-aliasing, it is important to have a solid understanding as to how an image is displayed on a computer monitor and its relation to anti-aliasing theory.

• Displaying an Image

After the rendering stage of the graphics pipeline is completed, the final image is stored within a segment of video memory called the “frame buffer”. A different amount of frame buffer memory is required depending on the resolution of the generated image: as the resolution is increased, the amount of required memory also increases. Due to the increased memory demands associated with higher resolution images, frame rate will always decrease as the resolution is increased. The digital image stored in graphics memory must then be converted into an analog signal so that the image can be displayed on a computer monitor. Using a digital to analog (DAC) converter, every pixel making up the image is converted into a varying set of voltage signals. These different voltages are used in conjunction with an electron gun, which fires electrons at phosphors located on a monitor, and results in a lit pixel. A common term associated with monitors is refresh rate or VSYNC, (i.e. 60 Hz, 75 Hz), which indicates how many times a complete image is drawn per second.


Based upon the above description, it should be apparent that all images are displayed with a large group of colored pixels, where each pixel occupies a finite area of space. As the resolution is increased, the area of each pixel is decreased, and a greater number of pixels are needed to draw the scene, hence improving the detail of the image. A typical image consists of 1024 horizontal rows, with each row containing 768 pixels, giving a total of 786432 pixels in a 1024x768 resolution image.

When an image is displayed on a monitor, it is an attempt to describe a real life object with a finite amount of detail (or pixels). Aliasing will occur if an insufficient amount of data is used to describe an image. Aliasing would not be visible to the human eye, if images could be displayed at very high resolutions (i.e. 6000 x 4000), but the technology required to display images at such high resolutions is still years away. Anti-aliasing is a technique used to try and alleviate the problems related to aliasing without having to draw images at unattainable resolutions.

• Anti-Aliasing Theory

Anti-aliasing is a sampling technique used to create more realistic looking images, by removing the stair stepping effect seen on the edges of most objects within computer generated images. As previously mentioned, this blocky look is a result of the fact that every image displayed on a monitor is actually composed of a finite number of pixels. See Figure 1.

Figure 1: Anti-aliasing illustration



The size of a pixel is fixed for a corresponding resolution, so the only way the blocky look of an image can be removed is by altering the color or brightness of a pixel. Anti-aliasing is a process used to alter the color of pixels in an effort to maintain the intended image fidelity of the original modeled object. A more accurate pixel color is generated through averaging a set of pixel samples.

To gain a better understanding as to how anti-aliasing works and what is actually being corrected, one must understand how using a jittered sample pattern could provide an improved level of image quality. If a standard sampling pattern is used, some edges will be anti-aliased much better than others, creating an image with less than desirable quality. The reason for this is that the position and angle that edges occur will change on a per-pixel basis. Hence, when an ordered sampling pattern is used, samples taken from some pixels may not properly describe the edge transition from one edge to another, resulting in incorrectly colored pixels. Using a jittered sampling pattern provides higher quality anti-aliasing, as a result improved distribution of pixel samples.

An analogy, which helps clarify the importance of sampling, is the creation of an audio CD. An audio CD is used to convert and store a continuous analog audio signal, as a set of discrete digital samples. A total of 44100 samples are taken per second to create a high quality audio signal. As fewer samples are taken per second, the audio quality of the signal deteriorates significantly. This is not to say that maximizing the number of samples taken is always the best solution. In the audio case, the human ear can only detect frequencies up to about 20 KHz; hence taking audio samples above this frequency is pointless. This same idea can be applied to displayed images; the human eye is much more sensitive to graphical anomalies along the vertical and horizontal axis. So if sub-pixel samples are taken from a more random distribution, a higher quality image can be produced using a smaller set of samples.

A significant problem occurs when the same samples are used multiple times in the generation of an anti-aliased image. Using the same samples multiple times causes blurring, and will always lead to severe degradation in image quality. Again, one can easily imagine that the quality of an audio signal would also be greatly compromised if the same audio data samples were used at multiple points in time.

• Anti-Aliasing Methods

The two most commonly used anti-aliasing solutions are “Super sampling” and “Multi-sampling”. Both of these methods blend sub-pixel samples to correct aliasing, but generate pixel samples in different ways. The exact distinction between these two methods is somewhat unclear, and has been defined mostly by marketing material more than anything else.

When it comes to anti-aliasing the only real difference between these methods is whether or not the image textures are accessed multiple times or only once. The sampling method which only accesses a set of texture once has generally been called multi-sampling, which is in fact counter intuitive to what is actually happening. Super sampling on the other hand accesses texture data multiple times, as it actually renders a set of new pixels.

• Super Sampling

“Super sampling” is commonly known as the anti-aliasing technique that renders every image at a much higher resolution than the actual display image, and then scales and filters the high resolution pixel samples together to create the anti-aliased pixels at the original resolution. For example, implementing 4x super-sampling at 800x600 actually renders a scene in a memory buffer at a resolution of 1600x1200. Every pixel in the 800 x 600 anti-aliased scene would then be generated through the averaging of four sample pixels from the 1600 x 1200 resolution image.

• Multi-Sampling

The second type of anti-aliasing technique is known as “Multi-sampling”, which can be implemented in two different ways:


1. The first method renders the scene multiple times at the original resolution, but shifts each copy by a distance of less than half a pixel in a random direction. Each slightly shifted image is rendered into a separate buffer in graphics memory. The final image is then created through blending the slightly shifted images together. This method has the same memory bandwidth and performance requirements as super sampling, as each shifted image must access its own set of textures separately.

2. The second multi-sampling approach obtains samples in a very similar way to the first multi-sampling method, except that it uses the same texture data taken from the first sample for all subsequent samples taken within each pixel. Although using the same texture information repeatedly offers an improvement in performance, the quality of the final anti-aliased image will be less than that of an anti-aliased image created through super sampling. Figure 2 and Figure 3 clearly illustrate the far superior image quality as a result of SMOOTHVISION™ compared to the multi-sampling technique used by the competition. This type of multi-sampling approach is an edge-based process, which means that it only affects pixels that contain edges, where the color abruptly changes. Pixels not located along an abrupt edge will look exactly the same after edge-based multi-sampling has been implemented, resulting in much lower quality textures. Note that the image quality provided with SMOOTHVISION™ High-Quality 2x anti-aliasing looks better than the multi-sample 4x ordered grid approach.

Figure 2:
SMOOTHVISION™ High-Quality 2x setting in Max Payne. Note the amazing texture detail seen in the green and white sign that says “JOHNSON”. Note that the image quality of the SMOOTHVISION™ High-Quality 2x setting is better than that of a 4X ordered grid anti-aliasing technique.



FIGURE 3:
Ordered grid 4x multi-sampling in Max Payne. The impact of using less texture data becomes very apparent. The sign that appears to be clear above is now blurred.



The Impact of Using Different Sampling Methods

For any specific level of anti-aliasing, both multi-sampling and super sampling generate the same number of pixel samples. However, the quality and method by which these samples are obtained have a large impact on the quality of the final anti-aliased image.

This is quite apparent when comparing ordered grid (OG) sampling to rotated grid (RG) sampling. The difference in the sampling pattern for OG and RG sampling is shown below in Figure 4.




Figure 4: Ordered Grid sampling on the left and Rotated Grid sampling on the right

The difference between the two methods becomes especially apparent along nearly horizontal and nearly vertical edges, where aliasing artifacts are most visible. In the ordered grid case, only 0, 2, or 4 pixel samples may be taken per pixel, which leads to non-gradual pixel color transitions along anti-aliased edges, decreasing the quality of the final image. Rotated grid sampling allows for a larger number of possible pixel samples, giving a more natural color transition across anti-aliased edges. Using a jittered sampling pattern provides the ultimate in natural edge color transition, providing the highest in anti-aliasing image quality.

The images found on the next page clearly demonstrate the superiority of ATI's SMOOTHVISION™ anti-aliasing technology. Figure 5 compares two different 2x anti-aliasing techniques. The image on the left uses multi-sample anti-aliasing combined with an algorithm that blends outside pixel samples. Using these techniques clearly leads to degraded texture quality and over all decreased image quality. The Image on the right uses the 2x SMOOTHVISION™ High-Performance setting. The SMOOTHVISION™ 2x High-Performance image has a greater level of visual quality as well as a higher level of performance when compared to the image on the left.

Figure 6 shows the image quality of a 4x multi-sampling technique compared with ATI's SMOOTHVISION 2x High-Quality setting. Again the superior image quality is quite apparent in the screen shot that uses SMOOTHVISION™ High-Quality 2x anti-aliasing. The 4x multi-sampled image has textures that look blurred and washed out compared to the image that used SMOOTHVISION™ anti-aliasing. Also note that a much higher performance level is possible with SMOOTHVISION™ 2x High-Quality anti-aliasing than the 4x multi-sampling approach.




Figure 5: Quake3 screen shot: The image on the left uses 2x multi-sample anti-aliasing and blends 3 exterior samples per pixel as well to create the final pixel color. Note the blurring as a result of the excessive blending. The image on the right uses the SMOOTHVISION™ 2x High-Performance mode. Note the superior texture quality and improved anti-aliased edges in the image on the right.




Figure 6: Quake3 screen shot: The image on the left uses 4x multi-sampling, and the image on the right uses the SMOOTHVISION™ 2x High-Quality mode. Note the superior texture quality and improved anti-aliased edges in the image on the right. Also note that the performance of SMOOTHVISION™ 2x High-Quality is much faster than the multi-sampled 4x method.

• How SMOOTHVISION™ Works

SMOOTHVISION™ is a super-sampling technique that gives users a huge number of different sampling options, by varying the number of samples taken per pixel and changing the sample processing technique. In both the High-Quality settings and High-Performance settings a total of 5 different anti-aliasing levels are possible, giving a total of 10 different settings. ATI's implementation of anti-aliasing provides a substantial improvement over other forms of anti-aliasing found in other GPUs. Instead of shifting every pixel in the same direction, groups of pixels can be individually shifted in different directions. This technique can be used to generate a more random distribution of pixel samples, leading to higher quality anti-aliased images.

SMOOTHVISION™ uses a large group of 16-sample blocks that cover a variable number of pixels, which is determined by the level of anti-aliasing being implemented. Each pixel within every group of 16 samples has their actual sample positions selected from a pre-programmed 8-sample jitter table. See Figure 7.

As the level of anti-aliasing is increased, each 16-sample block will cover a smaller number of pixels. For example in the 2x anti-aliasing case, a total of 8 pixels are covered by the 16 sample group, and in the 4x case, a total of 4 pixels are sampled 4 times each. See Figure 8 for a 4x SMOOTHVISION™ example.

A non-ordered pixel sampling approach provides the optimal solution, since the eye has a much harder time seeing artifacts in the resulting anti-aliased image, being less sensitive to non-ordered sampling patterns.

High-Quality SMOOTHVISION™ anti-aliasing uses a different texture color per sample, providing superior image quality compared to other competing anti-aliasing techniques, as a greater amount of color data is used to create the final anti-aliased image. Combined with superior sampling pattern of SMOOTHVISION™, the High-Quality SMOOTHVISION™ provides users with an unparalleled anti-aliasing solution.

High-Performance SMOOTHVISION™ anti-aliasing uses a jittered sample pattern for high-quality jagged edge removal, combined with an advanced sample blending and filtering algorithm. This advanced algorithm allows SMOOTHVISION™ to take fewer samples per pixel by blending in samples from other outside pixels. This results in a High-Performance anti-aliasing setting, while maintaining a level of image quality very similar to competing anti-aliasing solutions.




Figure 7: 8 possible predetermined sample locations for a single pixel




Figure 8: Possible pixel sample locations for the SMOOTHVISION 4x setting

- **Conclusion**

SMOOTHVISION™ anti-aliasing offers users a variety of different image quality and corresponding performance settings, which allow users to configure their own gaming experience.

SMOOTHVISION™:

- Uses a jittered sampling pattern to provide the absolute highest quality of anti-aliased images possible.
- Gives users a flexible interface that allows them to pick from a variety of different anti-aliasing settings. The flexibility offered by SMOOTHVISION™ includes the ability to specify a large number of sample levels (2x, 3x, 4x, 5x, and 6x) as well as a High-Quality and High-Performance mode, for a total of 10 different possible settings.