GRAPHICS GLOSSARY OF TERMS

AGP
The accelerated Graphics Port (AGP) provides a high bandwidth, low overhead connection between the graphics device and system memory. This connection, combined with DMA transfers, allows effective and efficient communication between the system and graphics. AGP provides 264MB/sec (AGP-1X), 528MB/sec (AGP-2X) to 1GB/sec (AGP-4X) of dedicated bandwith, up to 8x the performance of PCI. AGP improves graphics performance for 2D and 3D graphics and for video. In addition to high bandwidth, AGP provides hardware and software support for sharing memory between the system and graphics -- see "AGP Memory".

AGP Memory
AGP memory refers to memory that is on the computer (system memory). The graphics card is able to directly access this memory as if it were its own through the AGP bus. Data transfer through the AGP bus is faster than transferring data through the PCI bus, since the AGP bus is dedicated solely to the graphics card, whereas the PCI bus is shared among other peripherals.

AGP Texturing
A texture that is in AGP memory that can be used as the source texture for a polygon. Only the pixels that are actually read need to be transferred to video memory, rather than unnecessarily loading the whole image. With non-AGP textures, the entire texture must be loaded into video memory before it can be used as a texture. If there is not enough available video memory, then some other texture in video memory must be released. This results in texture thrashing.

AGP Writing
Images that are rendered into surfaces that are in AGP memory. The target surface does not have to be in video memory.

Additive Color
A method of generating many colors by combining a small number of primary colors. This is the color model used on CRT displays, with red, green, and blue as the three primary colors.

Aliasing
Visual artifacts produced by using pixels for display. The most familiar example is 'stairstepping' or jagged edges on near horizontal or near vertical lines.

Alpha
A fourth color component (Alpha is the A in RGBA)
The RGB components represent color (Red, Green, Blue); the A (Alpha) component represents opacity. Alpha values can range from completely transparent to totally opaque.

Alpha Blending
Used to render polygons that are partially transparent. The degree of transparency is the alpha value, which can range from completely transparent to completely opaque. Useful for rendering effects such as smoke, explosions, water or glass.

Ambient Light
A lighting model that produces a constant illumination on all surfaces, regardless of orientation. For example, a red ball would look like a red circle if there were no highlights or shading to indicate depth. This is the easiest type of light to simulate. Commonly used in computer graphics to ensure that all objects in a scene are visible.

Animation
Generating repeated renderings of a scene, with smoothly changing viewpoint and/or object positions, quickly enough that the illusion of motion is achieved.

Antialiasing
A rendering technique that assigns pixel colors based on the fraction of the pixel's area that's covered by the primitive being rendered. Antialiased rendering reduces or eliminates jagged lines or edges that result from aliased rendering.

Application Program Interface (API)
Application Program Interface (API) or middleware, is a layer of software that sits between the application program/user interface and the hardware. APIs define what capability is available, enabling an application to perform graphics or other system operations.

Architecture Review Board (ARB)
The OpenGL Architecture Review Board (ARB), an independent consortium formed in 1992, governs the OpenGL specification. Composted of members from many of the industry's leading graphics vendors, the ARB defines conformance tests and approves OpenGL enhancements. In 1999 the board includes representatives from 3Dlabs, Compaq, Evans & Sutherland, Hewlett Packard, IBM, Intel, Intergraph, NVIDIA, Microsoft, and SGI.

Aspect Ratio
The ratio of width to height of a rectangular area

Bi-Linear Interpolation
A texture display mode. Instead of retrieving a single texel (see texel) from a texture map, the four texels nearest the texture coordinate are retrieved and averaged together to produce a single value. This is done to minimize texture aliasing.

B-Spline Curve
A mathematical representation of a curve which allows curves to be stored as parameters, which reduces the amount of information necessary to describe a complex curve and increase performance.

Back Face Cull
The removal of surfaces in the image that are facing away from the viewer. Performance is increased by avoiding unnecessary drawing of these facets.

Bitmap
A rectangular array of pixels. Each pixel contains a color. Pixels in a bitmap are addressed by their X and Y location in the array.

Bitmapped Display
A raster scan display that stores pixel value in a frame buffer.

Bitplane
A rectangular array of bits mapped one to one with pixels. The framebuffer is a stack of bitplanes.

Blending
Combining two color components to produce a single color, usually as a linear interpolation between the two components.

Buffer
A group of bitplanes that store a single component (such as depth or green) or a single index (such as the color index or the stencil index). Sometimes the red, green, blue, and alpha buffers together are referred to as the color buffer, rather than the color buffers.

Clipping
The elimination of the portion of a geometric primitive that's outside the half-space defined by a clipping plane. Points are simply rejected if outside. The portion of a line or polygon that's outside the half-space is eliminated, and additional vertices are generated as necessary to complete the primitive with the clipping half-space. Geometric primitives and the current raster position (when specified) are always clipped against the six half-spaces defined by the left, right, bottom, top, near, and far planes of the view volume. Applications can specify optional application-specific clipping planes to be applied in eye coordinates.

Color Index
A single value that represents a color by name, rather than by value. OpenGL color indices are treated as continuous values (for example, floating point numbers) while operations such as interpolation and dithering are performed on them. Color indices stored in the framebuffer are always integer values. Floating-point indices are converted integers by rounding to the nearest integer value.

Color Index Mode
An OpenGL context is in color index mode if its color buffers store color indices, rather than red, green, blue and alpha color components.

Color Interpolation
Gradual transition from one color to another across a polygon or line.

Color Map
A table of index-to-RGB mappings that's accessed by the display hardware. Each color index is read from the color buffer, converted to an RGB triple by lookup in the color map, and sent to the monitor.

Component
A single, continuous value (for example, floating point) that represents an intensity or quantity. Usually, a component value of zero represents the minimum value or intensity, and a component value of one represents the maximum value or intensity, though other normalizations are sometimes used. Because component values are interpreted in a normalized range, they are specified independent of actual resolution. For example, the RGB triple (1,1,1) is white regardless of whether the color buffers store 4, 8, or 12 bits each.

Context
A complete set of OpenGL state variables. Note that framebuffer contents are not part of OpenGL state, but that the configuration of the framebuffer is.

Convex
A polygon is convex if no straight line in the plane of the polygon intersects the polygon more than twice.

Convex Hull
The smallest convex region enclosing a specified group of points. In two dimensions, the convex hull is found conceptually by stretching a rubber band around the points so that all of the points lie within the band.

Contouring
The process of using different colored lines are shaded bands on a surface to represent regions of different value or ranges of values.

Coordinate System
In n-dimensional space, a set of n linearly independent vectors anchored to a point (called the origin). A group of coordinates specifies a point in space (or a vector from the origin) by indicating how far to travel along each vector to reach the point (or tip of the vector).

Culling
The process of eliminating a front or back face of a polygon so that I isn't drawn.

Current Matrix
A matrix that transforms coordinates in one coordinate system to coordinates of another system. There are three current matrices in OpenGL: the modelview matrix transforms object coordinates (coordinates specified by the programmer) to eye coordinates; the perspective matrix transforms eye coordinates to clip coordinates; the texture matrix transforms specified or generated texture coordinates as described by the matrix. Each current matrix is the top element on a stack of matrices. Each of the three stacks can be manipulated with OpenGL matrix-manipulation commands.

Depth
Generally refers to the z window coordinate.

Depth-cueing
A rendering technique that assigns color based on distance from the viewpoint.

Direct3D
Direct 3D is Microsoft's standardized 3D programming interface.

Dithering
A technique for increasing the perceived range of colors in an image at the cost of spatial resolution. Adjacent pixels are assigned differing color values; when viewed from a distance, these colors seem to blend into a single intermediate color. The technique is similar to the half-toning used in black-and-white publication to achieve shades of gray.

Double-buffering
A technique for dividing the frame buffer into two areas, one "back" buffer and one "front" buffer. This allows the application to display completed frames in the front buffer while building a new incomplete frame in the back buffer. After the new frame is completed it is swapped with the front buffer. The user only sees a series of completed frames.

Face
One side of a polygon. Each polygon has two faces: a front face and a back face. Only one face or the other is ever visible in the window.

Flat Shading
Refers to coloring a primitive with a single, constant color across its extent, rather than smoothly interpolating colors across the primitive.

Fogging
Used to add the effect of the scenery fading off into a fog, or hide the far clipping plane. By hiding the far clipping plane, users won't see polygons suddenly popping into view in the distance.

Frame Buffer
The portion of video memory where pixel data is stored. Pixel data includes color information, and may include other information such as Z value (depth). The frame buffer resides in the 3D card's onboard video memory.

Geometric Primitive
A point, line, or polygon.

Gouraud Shading
A shading algorithm that takes the 3 colors defined at each point of a triangle and smoothly interpolates them throughout the surface of the triangle. It is often used to make round objects look more round and smooth.

Graphics Pipeline
The internal process of creating an image from user definition to final output. The pipeline determines how the data is processed in hardware and software, both color information and geometric coordinates.

Graphics Plane
Each set of bits in a frame buffer is referred to as a graphics plane and is commonly used to display pictures.

Hidden Surface Removal
To imitate real life, computer graphics must ensure that all surfaces in front hide the surfaces behind them. This process is called hidden-surface removal.

Lighting
The process of computing the color of a vertex based on current lights, material properties, and lighting-model modes.

Luminance
The perceived brightness of a surface. Often refers to a weighted average of red, green, and blue color values that gives the perceived brightness of the combination.

Matrix
A two-dimensional array of values. OpenGL matrices are all 4x4, though when they are stored in client memory they're treated as 1x16 single-dimension arrays.

MIP-Mapping
Multiple textures of increasing size are used to represent one texture. When a polygon needs to use the texture as a texture map, the texture whose size most closely matches the size of the polygon that is to be rendered is used. This reduces rendering artifacts such as "sparkling' and "moire" patterns as well as blockiness.

Immediate Mode
A rendering mode in which the graphics system processes each drawing command as soon as it receives it. The graphics system does not retain any information about the command after acting on it.

Mixed Mode
A combination of both immediate and retained mode rendering from within the same application. In this case there is no guarantee that the graphics database contains all the data that is necessary to render a compete picture.

Retained Mode
A rendering mode in which the graphics system stores each drawing command in a specialized graphics database, sometimes referred to as a structure or display list, upon receipt. Actual drawings are generated by a separate traversal of the display list. The structure is retained in memory and can later be modified or edited by special commands. Some people refer to this as structure mode.

OpenGL
The OpenGL is an open graphics API originally developed by SGI. Since 1992, the OpenGL Architecture Review Board (ARB) has guided the OpenGL specification. OpenGL is used with many video adapters for many 3D applications, from games to high-end CAD. It is a standard graphics API.

Overlay Planes
Frame buffer memory that stores pixel data, which is independent of the image buffer. This is often used for menus, window borders, and other user interface images.

Palettized textures
A form of texture compression. Used to reduce the size of a texture when the texture does not have many unique colors. 4-bit palletized textures can have 16 (2^4) different colors. 8-bit palletized textures can have 256 (2^8) different colors.

Perspective Correction
A technique that is used to give the proper illusion of depth when pictures are texture mapped onto polygons. Without perspective correction, warping will occur and the resulting image will not look realistic.

Phong Lighting Model
An empirically derived lighting model which accounts for the ambient, diffuse, and specular reflections of light from a point
on a surface.

Pixel (Picture Element)
The smallest element of a display surface that can be independently assigned a color.

Polygon
A near-planar surface bounded by edges specified by vertices. Each triangle of a triangle mesh is a polygon, as is each quadrilateral of a quadrilateral mesh.

Primitive
A point, a line, a polygon, a bitmap, or an image.

Radiosity
An illumination model in which the visibility of any portion of a surface of an object is assessed (pre-processed) relative to every other surface and light source. Radiosity techniques can produce soft shadows, unlike ray tracing.

Rasterize
Convert a projected point, line, or polygon, or the pixels of a bitmap or image, to fragments, each corresponding to a pixel in the framebuffer. Note that all primitives are rasterized, not just points, lines, and polygons.

Raster Scan Display
A CRT display device where each horizontal line is scanned on the display sequentially, setting pixel intensities to form a picture.

Ray Tracing
An illumination and shading model which takes into consideration not only the direct light sources but also all the reflected, refracted, and transmitted light in a scene. Ray tracing approximates the scene by tracing rays from the eye through each pixel and into the environment, which provides realistic images with sharp-edged shadows.

Reflection
In computer graphics, a surface characteristic used to determine color intensity in lighting models, and therefore helps determine how the surface looks.

Diffuse Reflection
The scattering of light in all directions. This type of reflection occurs when the object has a dull, matte surface. The surface appears to have the same brightness from all viewing angles.

Specular Reflection
The type of reflection that occurs when light hits a shiny surface. It produces highlights on the object. It is seen when the light rays from the light source are reflected directly into the viewer's eye.

Refresh Rate
The number of times per second the display screen is scanned. Also called vertical scan rate.

Rendering
The process of producing images or pictures. Rendering techniques such as shading, light modeling, or depth cueing are sometimes used to make an image look realistic. Rendering is the primary operation of OpenGL.

Resolution
For a CRT, the maximum number of displayable pixels in the horizontal and vertical directions. For a printer or plotter, the number of pixels per inch.

Rotation
A geometric transformation that causes points to be re-oriented about an axis.

Scaling
Changing the size of an object without changing its location or orientation.

Scan Converter
The hardware that converts the primitive commands and data to pixels. This is the part of the graphics pipeline that converts endpoints of lines and polygons to complete lines and filled polygons.

Scan Line
A row of pixels on the CRT display. Also refers to the row of pixel values in the frame buffer.

Shading
The process of interpolating color within the interior of a polygon, or between the vertices of a line, during rasterization.

Shading Model
The algorithm used to create the intensity and color of the visible portions of a primitive.

Shadow
An area that is totally or partially obscured from light by an object. In computer graphics, ray tracing and radiosity are two illumination models that can create realistic shadows.

Stipple
A one- or two-dimensional binary pattern that defeats the generation of fragments where its value is zero. Line stipples are one-dimensional and are applied relative to the start of a line. Polygon Stipples are two-dimensional and are applied with a fixed orientation to the window.

Solids Modeling
Computer graphics that require not only specifying the surface characteristics, but also the inside volume characteristics of an object.

Surface Modeling
A surface model is created when primitives are filled with color. A surface model has a "skin" on the outside, but is hollow on the inside.

Texel
A texture element. A texel is obtained from texture memory and represents the color or the texture to be applied to a corresponding fragment.

Texture
Texture, or texture map, is a 2D bitmap image. It is applied to 3D surfaces 3D surfaces through a 2D to 3D mapping process. Texture mapping is used to add visual detail to 3D geometry without increasing the amount of 3D geometry required.

Texture Mapping
Texture mapping is a fast, effective technique for laying bitmapped images over solid models, like wrapping a package with patterned paper.

Texture Memory
Memory used to store texture maps. This memory is generally allocated exclusively to storing textures. Graphics systems may include an entire memory subsystem dedicated just to texture.

Texture Thrashing
Constant swapping of texture memory content to replace old textures with new ones requested by the application, causing a performance slowdown. Several approaches are used to minimize texture thrashing, including use of OpenGL texture objects, providing more texture memory, using AGP textures, and modifying applications to use smaller texture maps.

Texture Transparency
Texture maps are commonly specified with RGBA, containing both color and transparency information. A powerful graphics technique is to combine simple geometry with complex textures. For example, a photograph of a tree can be texture mapped onto a single polygon. The texels corresponding to leaves or trunk are opaque and all others are transparent. The result is a high quality image with minimal system overhead.

Transformation
Translation, scaling, and rotation of a geometric object.

Translation
The changing of position of an object without changing its shape, size, or orientation.

Transparency
Transparency allows realistic rendering of surfaces like glass and window screens. Specified as "alpha" value.

Triangle Setup Engine
Special circuitry in the graphics chip that computes the mathematical aspects of what is needed to draw 3D images on the screen. This frees the CPU to perform other computations.

Triple Buffering
Triple buffering is a technique of dividing the frame buffer into three areas: two drawing buffers and one display buffer. This allows the application to display a frame of animation from one buffer while drawing the next frame into the other buffer and then start rendering in the third buffer while the other two buffers are still in use, therefore further increasing performance.

Tri-Linear Filtering
An extension of Bi-linear filtering for better image quality. For Tir-linear filtering, Bi-linear Filtering is done twice, against two MIPmap levels. Hence, a total of eight texels are averaged together to produce a single pixel. Tri-Linear Filtering eliminates "texture crawling" or "texture sparkles" on moving textured surfaces.

True Color System
A 24-plane graphics subsystem which produces the complete range of 16.7 million available colors.

Viewing Coordinates
A coordinate system that defines the positions of objects relative to the location of a viewer or synthetic camera model.

Viewing Transformation
The transformation that converts modeling coordinates into virtual device coordinates.

Volumetric Rendering
Common in scientific visualization and medical applications, volumetric rendering involves the display of information from "real world" applications through direct conversion to pixel values without creating any geometric data.

Wireframe
The drawing of a model by tracing features such as edges or contour lines without attempting to remove invisible or hidden parts, or to fill surfaces.

Z-Buffer (also known as Depth Buffer)
The memory available to store the depth information of a pixel to be used for hidden-surface removal. When a polygon is rendered with z buffering, each of its pixels' depth (z-value) is compared with the corresponding value stored in the z-buffer. If the value stored in the z-buffer is less than the depth of the new arriving pixel, it is decided that this pixel is visible and should therefore be rendered. The z-buffer is then updated with the pixel's depth. If however, the value is greater, the pixel is rejected and will not be rendered, as this means it is behind of what has already been drawn. Z-buffering is used to ensure that objects are rendered in the right order, that is, objects in the back should not appear in front of objects that are in the front.

Z-Buffering (also known as Depth Buffering)
A hidden-surface-removal technique in which objects are scan-converted to pixel data and -depth values, and then inserted into a pixel array and a depth array. Where the object's pixel depth is closer to the viewer than that of the currently stored pixel/depth pair, the new color data is written into the frame buffer, replacing the old color, plus the Z-buffer is updated for that location on the screen.