



(19) **United States**
(12) **Patent Application Publication**
Hunter

(10) **Pub. No.: US 2008/0186329 A1**
(43) **Pub. Date: Aug. 7, 2008**

(54) **FULL-SCENE ANTI-ALIASING METHOD AND SYSTEM**

division of application No. 10/837,234, filed on Apr. 30, 2004, now Pat. No. 7,119,812, which is a division of application No. 10/233,106, filed on Aug. 28, 2002, now Pat. No. 6,922,199.

(75) Inventor: **Gregory M. Hunter**, San Jose, CA (US)

Correspondence Address:
Kimton N. Eng, Esq
DORSEY & WHITNEY LLP
Suite 3400, 1420 Fifth Avenue
Seattle, WA 98101

Publication Classification

(51) **Int. Cl.** *G09G 5/00* (2006.01)
(52) **U.S. Cl.** **345/611**

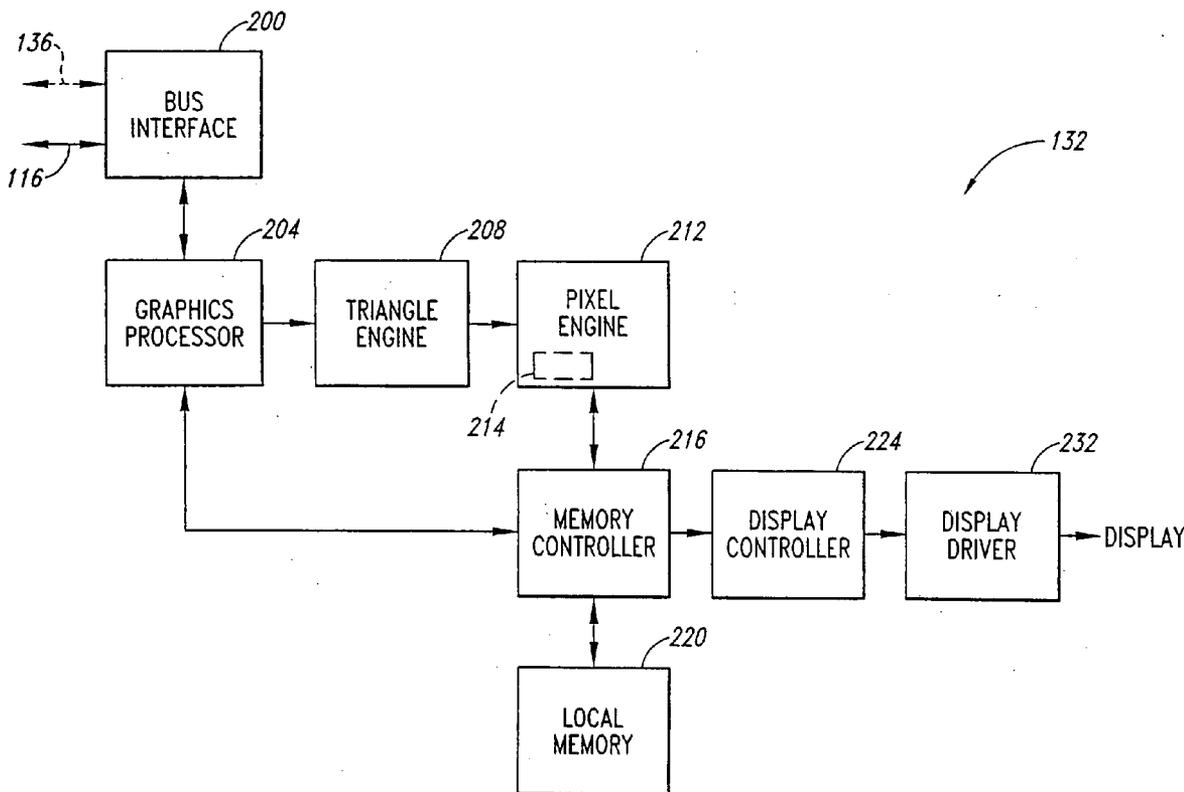
(73) Assignee: **Micron Technology, Inc.**, Boise, ID (US)
(21) Appl. No.: **12/079,840**
(22) Filed: **Mar. 27, 2008**

(57) **ABSTRACT**

A method and system for performing full-scene anti-aliasing for an image through a technique of rotating and unrotating rasterization of a scene and rendering a resulting image. A scene is rasterized at a first angle relative to a first coordinate system to generate a plurality of pixels, which are then applied to a polygon surface that is rendered at a second angle equal to the inverse of the first angle. Thus, the resulting image is re-oriented with respect to the first coordinate system.

Related U.S. Application Data

(60) Continuation of application No. 11/477,326, filed on Jun. 28, 2006, now Pat. No. 7,369,138, which is a



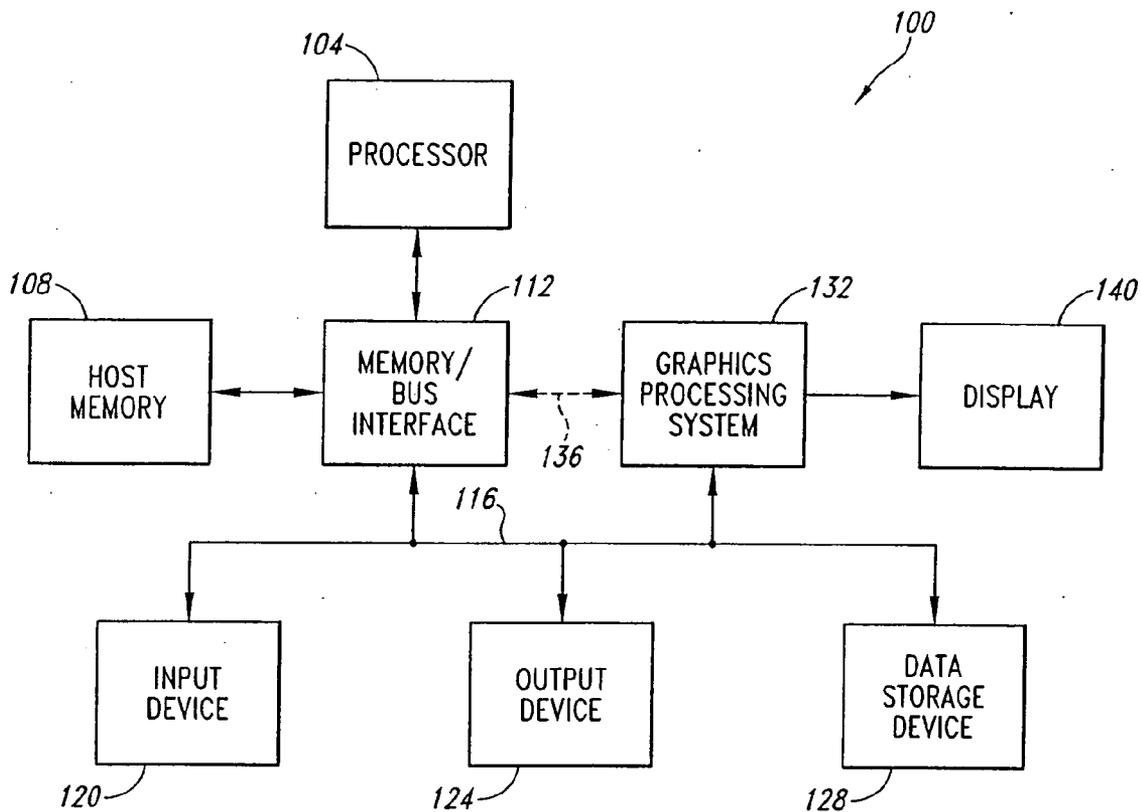


Fig. 1

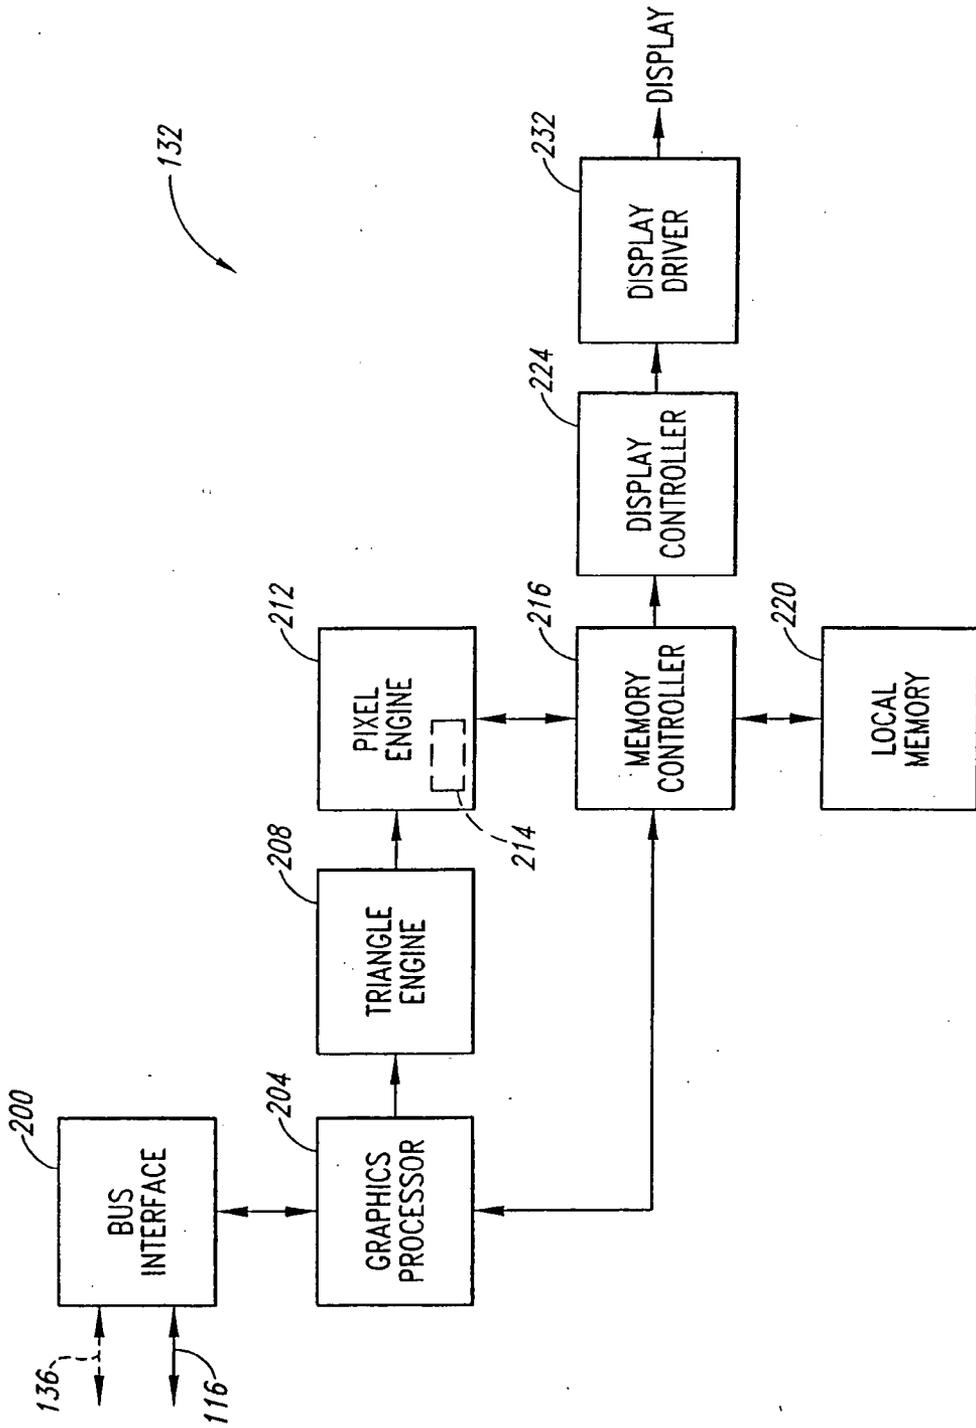


Fig. 2

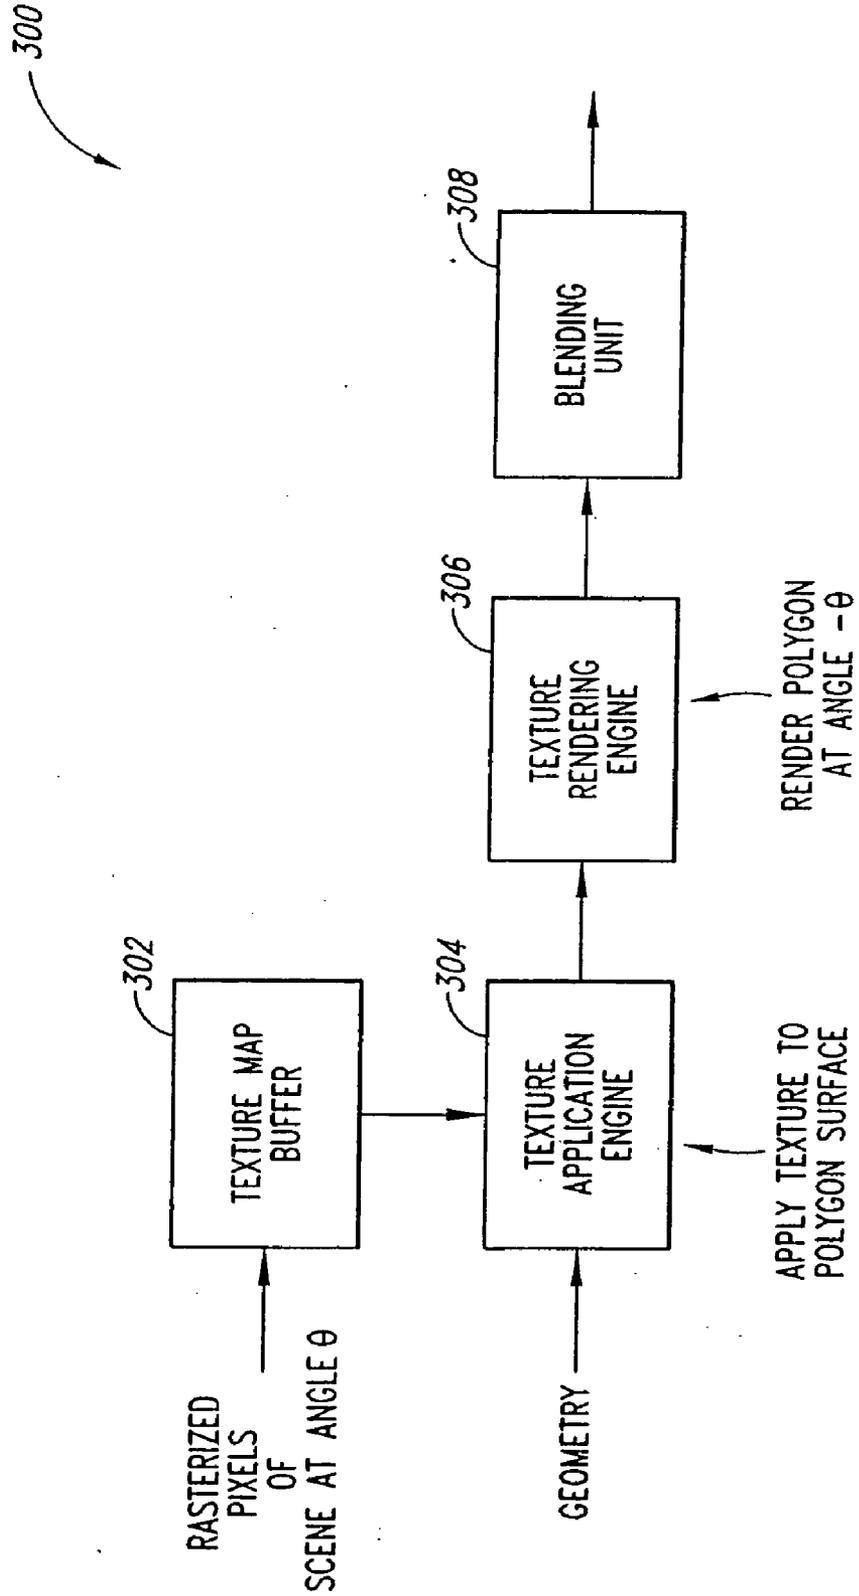


Fig. 3

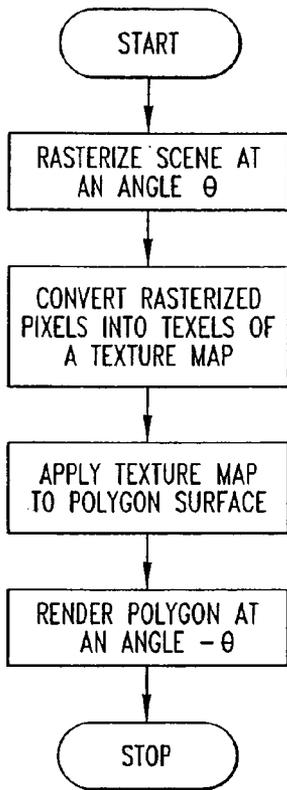


Fig. 4

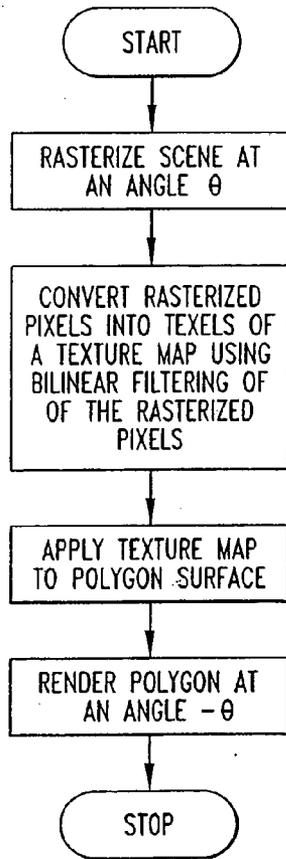


Fig. 5

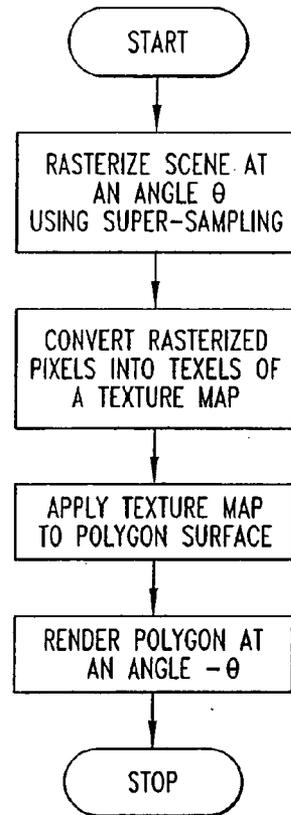


Fig. 6

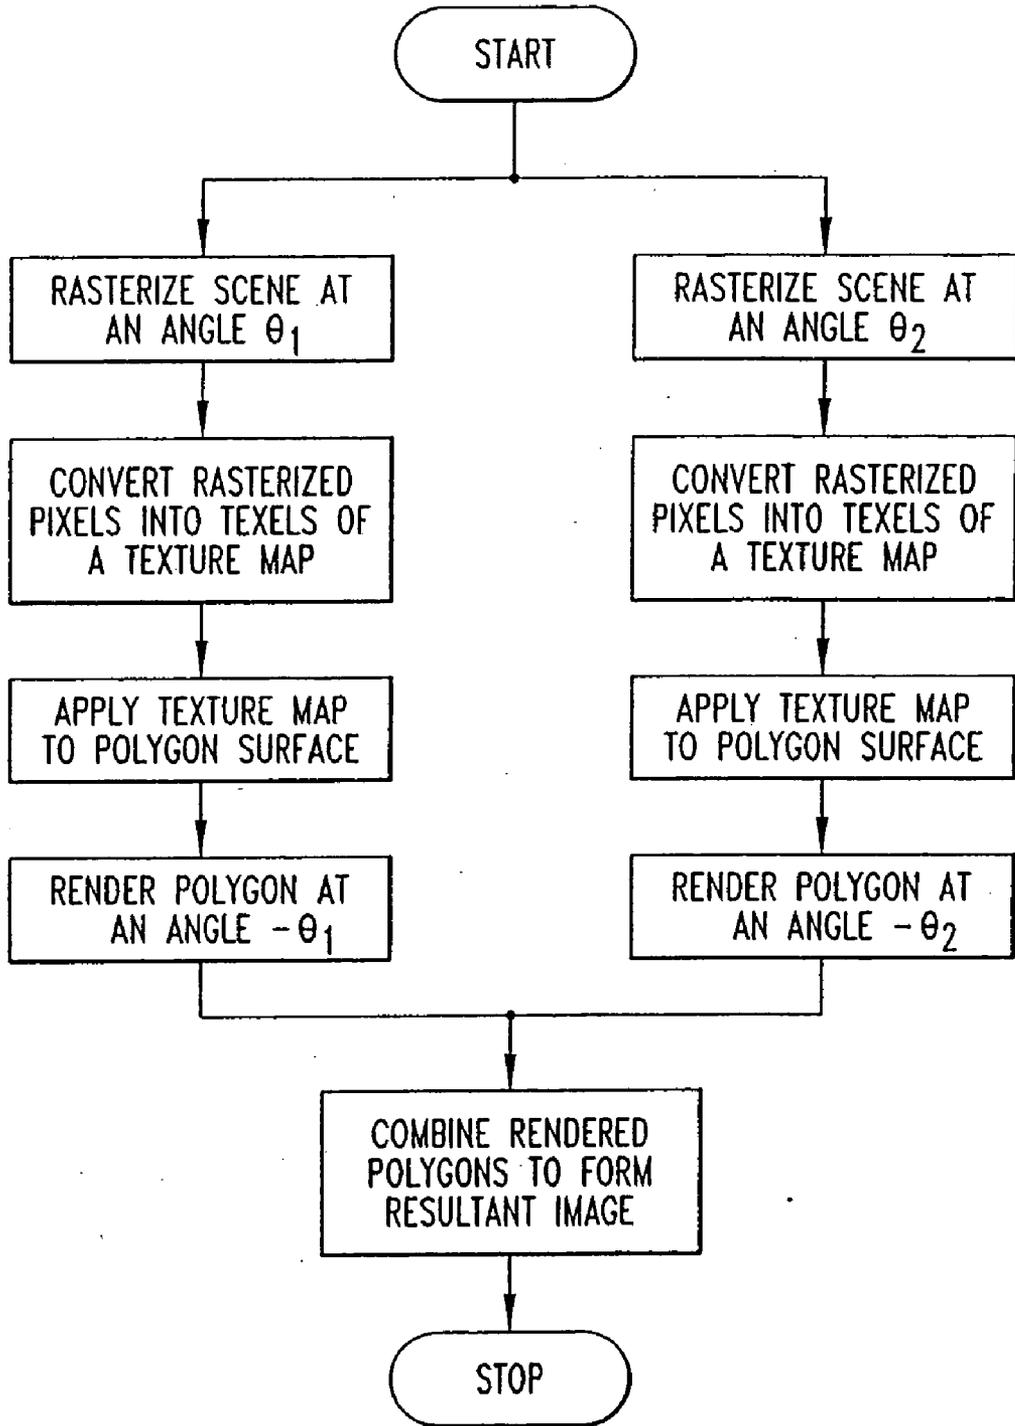


Fig. 7

FULL-SCENE ANTI-ALIASING METHOD AND SYSTEM

TECHNICAL FIELD

[0001] The present invention is related generally to the field of sampled imagery, and more particularly, a system and method for generating anti-aliased sampled images and sampling images by computation or acquisition while reducing aliasing.

BACKGROUND OF THE INVENTION

[0002] In computer graphics, input geometry or the like is typically used to represent a scene from which a sampled image is created. A sampled image is a grid, or “raster,” of discrete samples referred to as pixels. The pixels are generally arranged along two perpendicular axes corresponding to the horizontal and vertical axes of the image. Rendering a scene is the process of creating such an image of the scene. This process is also called rasterization. The value of a pixel is determined by calculating a sample value from the scene. The rasterized image, represented by the grid of pixels, is typically displayed visually, such as being printed or displayed upon a computer display monitor. The number of pixels and number of colors or values used in rendering a graphics image limit, to some extent, the quality, accuracy, and various characteristics of the image displayed, (for example, the image sampling accuracy, smoothness, detail, and resolution). A great amount of effort has been devoted to developing sophisticated graphics processing, rendering, and sampling methods for improving image quality. Rendering sampled images, however, produces unwanted artifacts. This production is called “aliasing.” An “anti-aliased” rendering method aims to reduce objectionable artifacts.

[0003] In addition to higher quality static images, rendering high-quality animation depicting visual effects of motion is also desirable. The illusion of motion is created by quickly displaying related images in a sequence in which an object changes appearance slightly. Although the animation of the object may appear to be continuous, each frame of the sequence is a separate image that is displayed momentarily. Thus, the quality of each frame impacts the quality of the overall animation or illusion of motion. Depending upon the method by which the sequence of images is sampled, the resulting representation of motion may result in the creation of motion artifacts. The more noticeable effects include “crawling” and “popping” of pixels on an edge of an object having a color that contrasts sharply with a different colored background.

[0004] For example, consider an object of a first color having a straight edge, that appears to be moving across a background of a second color. As the edge of the object moves across the boundary of a pixel, at some point there must be a determination as to how and when the color of the pixel changes from the color of the background to the color of the object. If the value of the pixel is taken from a single sample point, then when the edge of the object passes the sample point of the pixel, the color of the pixel is changed. The location of a single sample point is typically in the center of the pixel, and consequently, the color of a pixel is determined by the color at the center of the pixel. As a result the edge of the object may be well within the region defined by the pixel,

but the pixel may still have the color of the background because the edge has not reached the center (i.e., the sample point) of the pixel.

[0005] When multiple frames of images are displayed in sequence to provide the illusion of an object in motion, the effect of pixels along the edge of the object “popping” from one color or value to another can be distracting for a viewer. The relative motion of an object with respect to the orientation of the sampling grid of the image may be such that pixels along the edge of the object pop values in a manner and with a regularity that create a visual crawling effect along the edge of the object.

[0006] There have been many different approaches to addressing the issue of such aliasing artifacts, including motion artifacts and “staircasing.” One such approach is to increase the resolution, or the number of pixels used to represent an image. For example, the available resolution of computer graphics has increased dramatically since the first monochrome computer monitors were developed, and undoubtedly will continue to do so. Similarly, printing devices have ever increasing resolution as well. However, there are practical limitations on the manufacturing of displays and other devices used for recreating an image. Computational limitations must also be considered because there are practical limitations with respect to the amount of data that can be processed in rendering an image. Moreover, no matter how high a resolution is used to represent a graphics image, the image is still nevertheless sampled at discrete intervals, and consequently, subject to some degree of aliasing, including motion artifacts.

[0007] Another approach has been to increase the number of values computed for each pixel. More specifically in constructing an image, multi-sampling anti-aliasing systems might use sample values from multiple samples taken at different locations within a pixel in determining the final value of a pixel. As a result, as the edge of a moving object first passes over one of a pixel’s sample points, the pixel may take on a different value that is a compromise between the value of the object and that of the disparate background. For example, the pixel might change its value each time the object covers another sample point of the pixel, until all of the sample points share the color of the object, the pixel finally taking on the color of the object, assuming the size, shape, and motion of the object’s result in the object gradually and completely covering the pixel.

[0008] The multiple samples for each pixel are typically arranged in an ordered pattern, such as in a rectangular grid where the sides of a rectangle defined by the positioned samples are parallel to the axes of the pixels. Previous approaches have also positioned the samples in a (pseudo)-random fashion. Random multi-sampling per pixel generally requires a considerable amount of computational power. As a result, some anti-aliasing algorithms employ sampling patterns which attempt to simulate a random sampling pattern. However, these patterns may require nearly the same computing power as for random multiple sampling per pixel.

[0009] Although conventional, multi-sample, full-scene anti-aliasing methods generally do reduce motion artifacts and edge effects, in some situations motion artifacts remain. Full-scene anti-aliasing is antialiased rendering of substantially a full frame. For example, when an object has an edge parallel to an axis of the sampling pattern and is moving in a direction perpendicular to the edge, a popping effect may be created when the edge of the object passes through multiple

sample points simultaneously. Taking more samples per pixel may reduce the severity of multi-sampling motion artifacts, but processing overhead and economic considerations will place practical limits on the maximum number of samples that can be reasonably taken per pixel.

[0010] Therefore, there is a need for an alternative system and method for rendering an image from a representation of a scene, while reducing aliasing, such as motion artifacts and edge effects.

SUMMARY OF THE INVENTION

[0011] The present invention relates to a method and system for rendering a two- or three-dimensional scene, producing a sampled image, while minimizing aliasing. A scene is rasterized or sampled at a first angle relative to a first coordinate system. The resultant raster, comprising a plurality of pixels, is then rotated at a second angle equal to the inverse of the first angle. An image of the scene is then formed by sampling the pixels in the raster, producing an image that is properly oriented with respect to the first coordinate system. The second rotation may be performed by first applying the pixels to a polygon surface. The pixels may be first converted to a texture map and applied to the polygon surface as a texture. Then this polygon is rotated and rendered. The anti-aliasing technique may also include a super-sampling process, rasterizing the scene to produce a temporary surplus of pixels. To down-sample the number of pixels applied to the polygon surface, bilinear filtering may also be applied, when converting pixels to texels to form a texture map to be applied to the polygon surface. Such anti-aliased rendering may also be repeated on the scene, using two or more different initial angles, rendering two or more separate images that are ultimately combined, to produce a final resultant image. The result of multiple steps described above may be computed in a single step.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] FIG. 1 is a block diagram of a system in which embodiments of the present invention are implemented.

[0013] FIG. 2 is a block diagram of a graphics processing system in the system of FIG. 1 according to an embodiment of the present invention.

[0014] FIG. 3 is a block diagram of a portion of a pixel engine according to an embodiment of the present invention in the graphics processing system of FIG. 2.

[0015] FIG. 4 is a flow diagram of an anti-aliasing technique according to an embodiment of the present invention.

[0016] FIG. 5 is a flow diagram of an anti-aliasing technique according to another embodiment of the present invention.

[0017] FIG. 6 is a flow diagram of an anti-aliasing technique according to another embodiment of the present invention.

[0018] FIG. 7 is a flow diagram of an anti-aliasing technique according to another embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0019] Embodiments of the present invention render anti-aliased images while reducing artifacts in static images and/or in motion/animation. The systems and methods described herein rasterize a scene at an angle and apply the pixel values from the rasterization to a polygon rendered at the inverse of

the angle at which the scene was first rasterized. Embodiments of the invention may be called “rotate-unrotate” rendering. It will be appreciated that the terms “pixel,” “pixel value,” and “value” will be used herein interchangeably. When the terms are to convey different meanings, it will be apparent from the context. Certain details are set forth below to provide a sufficient understanding of the invention. However, it will be clear to one skilled in the art that the invention may be practiced without these particular details. In other instances, well-known circuits, control signals, timing protocols, and software operations have not been shown in detail in order to avoid unnecessarily obscuring the invention.

[0020] FIG. 1 illustrates a computer system 100 in which embodiments of the present invention are implemented. The computer system 100 includes a processor 104 coupled to a host memory 108 through a memory/bus interface 112. The memory/bus interface 112 is coupled to an expansion bus 116, such as an industry standard architecture (ISA) bus or a peripheral component interconnect (PCI) bus. The computer system 100 also includes one or more input devices 120, such as a keypad or a mouse, coupled to the processor 104 through the expansion bus 116 and the memory/bus interface 112. The input devices 120 allow an operator or an electronic device to input data to the computer system 100. One or more output devices 124 are coupled to the processor 104 to provide output data generated by the processor 104. The output devices 124 are coupled to the processor 104 through the expansion bus 116 and memory/bus interface 112. Examples of output devices 124 include a printer and a sound card driving audio speakers. One or more data storage devices 128 are coupled to the processor 104 through the memory/bus interface 112 and the expansion bus 116 to store data in, or retrieve data from, storage media (not shown). Examples of storage devices 128 and storage media include fixed disk drives, floppy disk drives, tape cassettes and compact-disc read-only memory drives.

[0021] The computer system 100 further includes a graphics processing system 132 coupled to the processor 104 and the host memory 108 through the expansion bus 116 and memory/bus interface 112. Optionally, the graphics processing system 132 may be coupled to the processor 104 and the host memory 108 through other types of architectures. For example, the graphics processing system 132 may be coupled through the memory/bus interface 112 and a high speed bus 136, such as an accelerated graphics port (AGP), to provide the graphics processing system 132 with direct memory access (DMA) to the host memory 108. That is, the high speed bus 136 and memory bus interface 112 allow the graphics processing system 132 to read and write host memory 108 without the intervention of the processor 104. Thus, data may be transferred to, and from, the host memory 108 at transfer rates much greater than over the expansion bus 116. A display 140 is coupled to the graphics processing system 132 to display graphics images. The display 140 may be any type of display, such as those commonly used for desktop computers, portable computers, and workstation or server applications, for example, a cathode ray tube (CRT), a field emission display (FED), a liquid crystal display (LCD), or the like.

[0022] FIG. 2 illustrates circuitry included within the graphics processing system 132, which performs various three-dimensional (3D) graphics functions. As shown in FIG. 2, a bus interface 200 couples the graphics processing system 132 to the expansion bus 116. In the case where the graphics processing system 132 is coupled to the processor 104 and the

host memory 108 through the high speed data bus 136 and the memory/bus interface 112, the bus interface 200 will include a DMA controller (not shown) to coordinate transfer of data to and from the host memory 108 and the processor 104. A graphics processor 204 is coupled to the bus interface 200 and is designed to perform various graphics and video processing functions, such as, but not limited to, generating vertex data and performing vertex transformations for polygon graphics primitives that are used to model 3D objects. The graphics processor 204 is coupled to a triangle engine 208 that includes circuitry for performing various graphics functions, such as scissoring, attribute transformations, rendering of graphics primitives, and interpolation of texture coordinates for a texture map.

[0023] A pixel engine 212 is coupled to receive the graphics data generated by the triangle engine 208. The pixel engine 212 contains circuitry for performing various graphics functions, such as, but not limited to, texture application or mapping, bilinear filtering, fog, blending, and color space conversion. The pixel engine further includes texture engine circuitry 214. As will be described in more detail below, the texture engine circuitry 214 is employed in rendering a polygon surface rotated at the inverse of the angle at which a scene was first rasterized.

[0024] A memory controller 216 coupled to the pixel engine 212 and the graphics processor 204 handles memory requests to and from a local memory 220. The local memory 220 stores graphics data, such as source pixel values and destination pixel values. A display controller 224 is coupled to the memory controller 216 to receive processed destination values for pixels that are to be displayed. The output values from the display controller 224 are subsequently provided to a display driver 232 that includes circuitry to provide digital color signals, or to convert digital color signals to red, green, and blue analog color signals, to drive the display 140 (FIG. 1).

[0025] Although the texture engine circuitry 214 is illustrated as being included within the pixel engine 212, it will be appreciated that the texture engine circuitry 214 may also be a circuit separate from the pixel engine 212, or included in one of the aforementioned circuit blocks of the graphics processing system 132. Therefore, the particular location of the texture engine circuitry 214 is a detail that may be modified without deviating from the subject matter of the invention and should not be used in limiting the scope of the present invention.

[0026] A portion of a pixel processing pipeline 300 that may be used for the texture engine circuitry 214 of the pixel engine 212 (FIG. 2) is illustrated in FIG. 3. A texture map buffer 302 stores texel data that are provided to and applied by a texture application engine 304 to surfaces defined by geometry data the texture application engine 304 receives. Following application of a texture map to a surface, a texture rendering engine 306 calculates, from the texture-mapped surface, pixel values that will be used to represent an image. A blending unit 308 may be included as part of the pixel processing pipeline 300 and is used for combining or blending multiple pixel values for each pixel of an image. In this manner, the pixels of multiple images can be overlaid and produce a resultant image where each pixel of the resultant image has a value that is a combination of the multiple values of corresponding, overlying pixels. It may be appreciated that a texture may be "applied" to a surface, simply by defining the geometric relationship between a defined surface and a defined texture map. Operation of the texture engine circuitry 214 in embodiments of the present invention will be explained in more detail below.

[0027] As mentioned previously, well-known circuits, control signals, timing protocols, and software operations have not been shown in detail in order to avoid unnecessarily obscuring the invention. However, it will be appreciated that the texture map buffer 302, the texture application engine 304, the texture rendering engine 306, and the blending unit 308 are conventional in design and operation, and the description provided herein is sufficient to enable those of ordinary skill in the art to practice the present invention. In summary, in embodiments of the present invention these elements may be used to render an image of a scene at a particular angle, as part of the operation of embodiments of the invention, or other elements may be used to render the image of a given scene at a particular angle, in the operation of embodiments of the invention. Commonly, these elements are used to render scenes, and these elements could be used to do so, at a selected angle. In the operation of embodiments of the invention, in any case, whatever employment of these elements in the rendering of the scene at a particular angle, they or their analogues are then used to render an image rotated at the inverse angle, of the image rendered at the first angle. It will be appreciated by those of ordinary skill in the art that appropriate control circuits, data paths, and memory, etc. are employed, as necessary.

[0028] It will be further appreciated that the sample values calculated during rasterization and applied to the polygon surface may represent values such as color or grayscale at a sample location, and that the sample values for the samples may consist of several different components. For example, a sample value may represent colors which are the combination of red, green, and blue color components. Another example includes sample values representing colors which are the combination of luma and chroma components. Consequently, because it is well understood in the art, although circuitry to perform graphics operation for each of the components is not expressly shown or described herein, embodiments of the present invention include circuitry, control signals, and the like necessary to perform anti-aliased rendering of each component for multi-component sample values.

[0029] Embodiments of the present invention can use conventional rendering circuitry/software to perform its operation. The operation of the graphics processing system 132 will be described with respect to the flow diagram of FIG. 4. In operation, a scene is rasterized at an angle such that the pixel values that are stored into a back-buffer, within local memory 220, for example, form an image rotated at an angle θ to a first coordinate system. Typically, the first coordinate system is the world space coordinate system where two axes are oriented along the horizontal and vertical directions. Thus, when the rasterization occurs at the angle θ , the image of the scene appears tilted. Were the values stored in the back-buffer to be displayed, the image would appear at an angle. So far, the texture engine may not have been used. However, in alternative embodiments of the present invention, the texture engine may have been used to render any texture-mapped polygons in the scene.

[0030] Possibly after format conversion and/or up sampling, the rasterized pixel values are then treated as texel values of a texture map such that the rasterized scene is a texture that may be applied to the surface of a polygon. The texture map is provided to the texture map buffer 302 (FIG. 3) for storage. Geometry data is provided to the texture application engine 304 defining a polygon surface onto which the resulting texture stored in the texture map buffer 302 is applied.

[0031] The texture application engine 304 applies the texture map, representing the rasterized scene at an angle θ , to

the surface of the polygon, and the texture rendering engine 306 renders the resulting polygon at an angle $-\theta$. That is, the polygon is effectively rotated by the inverse of the initial rasterization angle to correctly reorient the image with respect to the world space coordinates. The rendered polygon, which is now oriented correctly, may be displayed in a conventional visual medium, such as a computer display, a printer, or the like. Typically, the size of the polygon surface will be such that the display window will be filled by the interior of the polygon. However, it will be appreciated that the size, shape, and resolution of the polygon surface may be modified without departing from the scope of the present invention. Additionally, if the surface portion rendered is larger than the final image, the portion of the surface comprising the pixels needed to produce the final image will be extracted. The polygon onto which the texture is applied may have any desired resolution, however, as will be explained in more detail below, the desired resolution of the polygon is typically at least that of the desired output resolution.

[0032] It will be further appreciated that the angle θ can be selected in various manners without departing from the scope of the present invention. For example, the angle θ can be preprogrammed to a certain value. Alternatively, the angle θ can be selected from a preselected set of angles. The selection may be made in an ordered fashion or in a random or pseudo-random fashion as well. The preselected angles can be a fixed set of angles or positions calculated in advance, so that choosing an angle from a group of preselected angles will mean that a table of angles is constructed in advance, and then selections from the table are made during anti-aliased rendering. However, it will be further appreciated that, as an alternative, angles or positions could be generated in a constrained or controlled manner, to accomplish a similar result. For example, angles could be generated randomly or pseudo-randomly, and then transformed such that a useful subset of these become the set of preselected angles.

[0033] The anti-aliased rendering by embodiments of the present invention reduce polygon popping and other aliasing which might result along the vertical and horizontal axis. When the scene is initially rasterized at the angle θ , the sampling is oriented at an angle with respect to the desired image and the pixels from the rasterization represent an image of the rotated scene. Thus, horizontal and vertical lines in the original scene are now de-emphasized rather than being accentuated by the orientation of the scanlines. The scene is resampled during the rendering of the polygon rotated at an inverse angle of the initial rasterization angle to reorient the image.

[0034] Another embodiment of the present invention is described in the flow diagram of FIG. 5. The anti-aliasing algorithm illustrated in FIG. 5 is similar to that illustrated in FIG. 4, however, rather than constructing the texture map using point-sampling techniques, bilinear filtering is used in converting pixels into a texture map. It will be appreciated by those of ordinary skill in the art, however, that in the embodiment described with respect to FIG. 4, where rasterization is to desired texel locations, point-sampling is not necessary. In the embodiment shown in FIG. 5, as a result of bilinear filtering, artifacts that are introduced when texels of the texture map do not fall on pixel centers are reduced, and the overall appearance of edges in the resulting texture is smoothed further.

[0035] In a similar vein, illustrated in FIG. 6 is another embodiment of the present invention which relies on super-sampling during rasterization of the scene at an angle θ . The surplus of samples calculated during the rasterization of the scene are subsequently down-sampled when rendering the

polygon surface. The combination of super-sampling and down-sampling enhances the subjective anti-aliasing quality, and can be efficiently integrated with the initial rasterization and the reorienting rasterization.

[0036] Although not shown in FIG. 6, the rotate-unrotate technique that includes super-sampling of a scene may be modified to further include down-or up-sampling when converting the pixels to a texture map. For example, the number of pixels calculated through super-sampling can again be doubled along each dimension of the resulting image to produce approximately four times the number of texels for the texture map or, a bilinear blending technique can be used during application to the surface of the polygon to down-sample the number of texels to the appropriate output resolution.

[0037] Another embodiment of the present invention is shown in FIG. 7. The anti-aliased rendering illustrated in FIG. 7 is similar to that illustrated in FIG. 4; however, the scene is rasterized at two different angles, θ_1 and θ_2 , thus producing two frames of pixels that represent images of the scene, rotated, at the respective angles. Each of the frames of pixels is converted into, or treated as, a texture map of respective texels, which is subsequently applied to a respective polygon. Each polygon is then rendered as if rotated by the inverse of the respective initial rasterization angle. Both images now appear in the same orientation. The resulting images are then aligned and combined to produce a final image.

[0038] The values of the two respective rendered polygons can be combined in a variety of methods. For example, in one method of combining the values of the rendered polygons, the values of the first and second renderings are averaged to provide values for the pixels of the resultant image. Thus, an aliasing artifact at one location in the first image may be due to the angle in that the second image at the second angle may contain less aliasing at that same location. The averaging minimizes the degree of aliasing at such locations.

[0039] Another approach to combining the rendered images is to use an alpha-blending technique, where the pixels of the rendered polygons are overlaid and the respective opacity or alpha values are adjusted to yield composites for the pixels of the resultant image. In the case where the pixels are alpha-blended, the values of correspondingly located pixels are scaled by their respective alpha values, and then the scaled values are summed together to yield a resulting value for each pixel. However, other alpha- or opacity-blending techniques may be applied as well.

[0040] Still another method of combining the rendered images is to interlace or display in an alternating sequence the rendered images. Video interlacing can be used to effect a blending of multiple frame buffers. In conventional NTSC television, each interlaced video frame consists of two images called "fields." Each field is the even or odd lines of the image. When displaying video, the television updates each field every $1/60^{\text{th}}$ of a second. If the pixels from one frame buffer are used as the first field, while pixels from a different frame buffer are used for the second field, the interlacing apparatus will construct an output image which is a combination of both input images. Another similar technique for blending multiple frames involves displaying the unblended frames at a rate sufficient for the human eye to fuse images together.

[0041] Combining pixels of n framebuffers: the values are combined by summing the respective values, each having been multiplied by a respective weight. The set of weights comprises one weight for each of the images, each typically rendered at a different angle. This sum can be expressed by the following formula:

$$P_{\text{combination}} = w_0 * P_0 + w_1 * P_1 + w_2 * P_2 + \dots + w_n * P_n$$

where w_i is the i th filter weight and the P_i are corresponding pixels, each from its own image.

[0042] In general, it is expected that the filter is normalized. That is, the sum of the set of weights should be equal to one (sometimes referred to as “unity”). It is possible, however, that some applications will use weights that sum to less than one, or greater than one, as in a special effect. Although this technique appears very similar to the mathematical construction of alpha blending, with alpha blending, the term alpha is conceptually used to refer to the “coverage” of the pixel.

[0043] Another blending or combining technique uses logic operators. In this blending method, values are combined by the use of a logical expression. The logical expression is applied to each bit of the value of a pixel. This has the effect that individual portions of the pixel value may be selected or masked on a per-bit basis. This particular method will have applications where image mattes are used. For example,

[0044] blended output= P_1 AND P_2 ; or

[0045] blended output= P_1 AND NOT P_2

[0046] It will be appreciated that the techniques described herein for combining the multiple rendered images are exemplary, and do not provide a comprehensive recitation of all possible combining alternatives. Consequently, the particular technique for combining the multiple rendered images should not be interpreted as limiting the scope of the present invention, except to the extent the limitation is recited in the claims appended hereto.

[0047] As previously described, in the embodiment of the invention illustrated in FIG. 7, a scene is rasterized at both a first and a second angle. It should be noted that the time at which the scene is rasterized at the first and second angles may be at the same instant, or at two different instants. That is, the “same instant” is not intended to mean that the scene is simultaneously rasterized at the two different angles, but rather, the scene has not advanced or changed from the time rasterization at the first angle occurs and when rasterization at the second angle occurs. Thus, rasterization at the first and second angles capture the scene at identical moments in time.

[0048] In contrast, rasterizing at the two angles at different instants is intended to mean that the scene has changed from when the rasterization at the first angle occurs and when the rasterization at the second angle occurs. In this manner, a sequence of frames can be created and blended together. After two initial frames are rasterized and blended, a subsequent rasterized frame can be blended with the previous, single, non-composite image, thus, two consecutive frames are blended together for each output frame. As a result, when scene geometry moves at high speeds, the motion will tend to appear smoother, while motion at lower speeds will yield more anti-aliasing.

[0049] Although the embodiments of the invention described with respect to FIG. 7 have been described as combining only two different renderings resulting from the rasterization of a scene at two different angles, it will be appreciated that any plurality of renderings can be combined to form the resultant image. Combining more than two renderings may result in a resultant image having higher visual quality; however, processing capabilities and resources of the graphics processing system should be considered.

[0050] It will be further appreciated that multiple frames of texels and pixels may be combined using various alternative techniques to form a resultant image and still remain within the scope of the present invention. For example, the embodiment shown in FIG. 7 combined the frames of pixels follow-

ing the rendering of the polygon surfaces. However, the multiple frames could alternatively be combined by averaging the values of corresponding texels in the multiple texture maps prior to being applied to the polygon surface. That is, the texel values of a texel applied to the polygon is the average of the texel values for correspondingly located texels of the multiple texture maps.

[0051] Embodiments of the present invention have been described to this point as separate variations of embodiments of the present invention. It will be appreciated, however, that the various embodiments so described can be combined to provide additional alternative embodiments of the present invention. For example, anti-aliased rendering according to embodiments of the present invention can include super-sampling during the initial rasterization of the scene as well as bilinear filtering when converting pixels to a texture map. Similarly, bilinear filtering and rasterization at multiple angles can be combined to produce another embodiment. Thus, embodiments resulting from the combination of the embodiments particularly described herein remain within the scope of the present invention.

[0052] From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. For example, embodiments of the present invention have been described as being implemented in portions of the pixel engine 212 (FIG. 2) and performed by using conventional processing stages typically included in graphic processing systems. The ability to implement the present full-scene anti-aliased rendering invention with conventional processing stages is an advantage which the present invention provides. However, it will be appreciated that embodiments of the present invention may also be implemented in circuitry dedicated to performing full-scene anti-aliasing as described herein. Additionally, although the present invention has been described as being implemented in hardware embodiments, it will be appreciated that the present invention may also be implemented in software embodiments, or combinations of hardware and software. It will be further appreciated by those of ordinary skill in the art that embodiments of the present invention are not limited to implementation in a computer system. For example, embodiments of the present invention can be implemented in game machines, cellular phones, and other systems where rendering of anti-aliased images would be desirable. Accordingly, the invention is not limited except as by the appended claims.

1. A method for full-scene anti-aliased rendering of a scene, comprising:

- rasterizing the scene into a raster with the raster coordinate system rotated by a first angle relative to a viewing coordinate system;
- converting pixels of the raster to texels of a texture map;
- applying the texture map to a surface; and
- rendering at least a portion of the surface with a surface coordinate system at a second angle relative to the viewing coordinate system, the first and second angles being related such that a region of the rendered surface portion represents a final image of the scene in a desired orientation of the scene relative to the viewing coordinate system.

2-90. (canceled)

* * * * *