

METHOD AND APPARATUS FOR DITHERING
ANTI-ALIASED VECTORS

BACKGROUND OF THE INVENTION

1. FIELD OF THE INVENTION:

The method and apparatus of the present invention relates to the rendering of lines on computer graphic display devices. More specifically, the method and apparatus of the present invention relates to the rendering of high quality lines into a low resolution frame buffer.

2. ART BACKGROUND:

A picture or image may be thought of as a continuous representation of the object projected. It is often desirable to represent images on a digital computer display device such as a monitor or a printer. However, a computer and its output devices are binary in nature. A computer display comprises a grid of display elements, referred to as pixels, which can either be in one of a plurality of on states indicative of its intensity, in which the pixel location is illuminated, or in the off state in which the pixel location is not illuminated. In order to display continuous images, the continuous images must be converted to discrete representations corresponding to the pixel locations on the display device. A digital image is quantized in both the spatial domain and the intensity domain.

The amount of memory available (that is, the size of the frame buffer) dictates the quantized levels available to represent the image. If the resolution of either sample space (i.e. spatial and intensity) drops below a threshold (due to memory limitations), the eye will detect the discrete boundaries between samples. In the intensity domain, insufficient resolution is marked by the presence of artificial edges delimiting the transitions between regions of incremental intensity. Other undesirable visible effects, such as patterning, color shifting and biasing, are introduced due to the visible thresholding between quantized intensity levels.

To minimize the undesirable effects, a technique, referred to as dithering or digital halftoning, is used. Dithering is a technique which permits the simulation of intensity levels between quantized levels by permitting the eye to integrate fine detail within an area and record only the overall intensity of the area. Dithering aims to sacrifice some of an image's spatial resolution for an increase in perceived intensity resolution, accomplished by averaging the intensities of several neighboring pixels to simulate intensities that lie between quantization levels. Typically, this technique works well, since the eye naturally blends individual pixel values such that only the average intensity in the surrounding area or neighborhood is perceived. For more information on dithering, see Ulichney, Digital Halftoning (1987, MIT Press); Foley & Van Dam, Fundamentals of Interactive Computer Graphics, p. 597-602 (Addison-Wesley, 1984).

Several types of dithering techniques exist. The types of dithering algorithms are distinguished by the way the neighborhoods are chosen. In the technique of ordered dither, the neighborhoods are chosen according to a two dimensional set of values, referred to as the dither table or dither matrix, which is tiled into the image's coordinate space. Typically for area dither, the table has the same number of row and column elements and the total number of entries in the table equals the number of simulatable intensity levels between the quantized levels. The

values contained in the table are used to make decisions about which quantized intensity value will be output at each position, that is, should the intensity value be quantized to the quantized value above or below the original intensity value. The dither table values are different at each x,y coordinate location such that when a constant input intensity is dithered over some area, the output values will alternate in some pattern between the upper and lower quantized intensity levels.

The ordered dither process is explained in reference to Figure 1.

Process 1.0 determines the appropriate index into the dither matrix based on the x-y coordinates of the current pixel. Conceptually, the idea is to "tile" the dither matrix into the image coordinate space, thereby replicating each entry once every n pixels, where n is the dimension of the dither table. The indices to the matrix, i - j , are determined according to the following equations:

$$i = x \bmod n$$

$$j = y \bmod n$$

If n is a power of two, it is possible to utilize the least significant bits of the x-y screen coordinates for indexing the dither matrix. For example, if the dimension of the dither matrix is 4×4 , then the 2 least significant bits of the x-y coordinates will provide the required modulo-4 indexing.

At process 2.0, the dither values are retrieved from the matrix according to i - j dither indices determined. By the nature of the problem, the output intensity scale has fewer entries than the input intensity scale. In order to quantize the input value, there must be a predetermined correspondence between each level on the output scale to a value on the input scale. It is the task of process 3.0 to determine the input-scale base value, which is the input scale value nearest to but not greater than the original input intensity and which directly corresponds to a value in the stored range of intensities.

The next process, 4.0, determines the value in the stored range of quantized intensity levels that corresponds to the input - scale base. This value will be referred to as the "stored range base value".

At process 5.0, the difference between the input intensity value and its input-scale base value is determined. The result, always guaranteed to be greater than or equal to zero, is the delta value passed along to process 6.0. Process 6.0 is the pivotal dither process, comparing the delta value to the dither value and outputting either the stored range base value or the next higher value. Specifically, if the delta value is less than or equal to the dither value, the stored range base value is output; if delta is greater than the dither value, then the stored range base plus one is output. The output from process 6.0 is the quantized pixel value that is deposited into the pixel store (e.g. a frame buffer).

Although quantized values are stored, these values are typically mapped back to an intensity scale with a greater intensity range for display viewing. That is the task of process 7.0, which is usually implemented by lookup-tables lying between the frame buffer output and the video display input. (This is not part of dithering per se, but is utilized for reconstruction of the dithered image for viewing.)

By performing processes 1.0 through 6.0 on successive input values, dithering causes values both greater than and less than the input intensity to ultimately be displayed, producing regions over which the overall average intensity is approximately equal to the original input intensity.

Because the dithering process exchanges spatial resolution for a perceived increase in bit (i.e., intensity/color) resolution, it is assumed that the object to be dithered has sufficient area to permit the resolution exchange to be accomplished. However, when dithering a line of single pixel width, very little surface area exists and as a consequence, the quality of dithered lines are reduced. Artifacts, such as

disappearing line segments when the line color is close to the background color and line raggedness ("jaggies") manifest themselves in the digital image during dithering. This is illustrated in Fig. 2 which depicts a series of radial lines originating from a common central point dithered using area dither techniques. A decrease in the bit resolution (that represents the intensity values) also decreases the number of intensities available, making the image appear more ragged along its edges. In low resolution devices, these artifacts become more prevalent because the number of pixels available in any area of the digital image is small compared to the size of the image and therefore there are fewer pixels that may be used to minimize the visual effects of the artifacts by varying the intensities of the pixels.

To remove the jaggies which arise due to the under sampling of a line, anti-aliasing techniques are used. The visual effects of jaggies are decreased by smoothing the edges of the lines into the background colors by increasing the width of the line and using intermediate colors between the actual line color and the background color on the added line edge pixels. For further information on anti-aliasing, see Gupta, Sproull, "Filtering Edges for Gray Scale Displays, Computer Graphics, Vol 15, No. 3, August, 1981.

Compositing is a technique in which a complex graphic image to be rendered can be broken down into a plurality of "sub-images" or elements which are separately rendered and subsequently blended together. However, "hot spots" often arise after anti-aliasing techniques are performed on intersecting or overlapping lines which are composited together into a single image. This is due to the addition of intensity values at the area of intersections and the nature of anti-aliased lines which typically dictates that the center of the line has a greater intensity than the edges of the lines (which are at a lower intensity in order to blend into the background color).

SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide a method and apparatus for the rendering of high quality vectors in low resolution devices which employ dithering.

It is therefore an object of the present invention to provide a method and apparatus to render high quality vectors on low resolution frame buffers or output display devices.

In the method and apparatus of the present invention, vectors can be rendered and stored in low resolution frame buffers or digital display devices with the minimum of distortion. The vector to be rendered is first anti-aliased at a first resolution (for example, 8 bits/color component) to minimize jaggies and then is composited with the corresponding background pixels of the digital image stored in the frame buffer. Once the vector is composited with the corresponding pixels of the digital image, the pixel data output from the compositing process is dithered to a second, lower resolution for storage in the frame buffer and for subsequent output to a graphics display device. Prior to compositing, the data from the stored digital image is back-transformed to the first resolution for compositing the digital image at the same, higher resolution as the vector to be incorporated into the digital image.

According to the present invention there is provided a method for rendering a vector represented at a first resolution of intensity on a low resolution device at a second lower resolution of intensity, a store of said low resolution device storing a background pixel image, said vector to be rendered between first and second coordinates in the screen space of the low

resolution device, said method comprising the steps of:

anti-aliasing the vector at the first resolution to produce a pixel image of the anti-aliased vector;

compositing the said pixel image of the anti-aliased vector with the corresponding pixels of the background image stored in the low resolution device to produce a composited image, said corresponding pixels being those pixels of the background image which form a vector between said first and second coordinates of the screen space, said step of compositing comprising the steps of;

back-transforming the corresponding pixels of the background image from the second resolution to the first resolution;

compositing the pixel image of the anti-aliased vector with the back-transformed pixels of the background image;

dithering the composited image from the first resolution to a second resolution;

storing the dithered image at the second resolution to the store of the low resolution device;

whereby high quality vectors are rendered on the low resolution device.

Also, according to the present invention there is provided an apparatus for rendering a vector represented at a first resolution of intensity on a low resolution device at a second lower resolution of intensity, a store of said low resolution device storing a background pixel image, said vector to be rendered between first and second coordinates in the screen space of the low resolution device, said apparatus comprising:

means for anti-aliasing the vector at the first resolution to produce a pixel image of the anti-aliased vector;

means for compositing the pixel image of the anti-aliased vector with the corresponding pixels of the background image stored in the low resolution device to

produce a composited image, said corresponding pixels being those pixels of the background image which form a vector between said first and second coordinates of the screen space, said means for compositing comprising the steps of;

means for back-transforming the corresponding pixels of the background image from the second resolution to the first resolution;

means for compositing the pixel image of the anti-aliased vector with the back-transformed pixels of the background image;

means for dithering the composited image from the first resolution to a second resolution;

means for storing the dithered image at the second resolution to the store of the low resolution device;

whereby high quality vectors are rendered on the low resolution device.

NOTATION AND NOMENCLATURE

The detailed descriptions which follow are presented largely in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art.

An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. These steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It proves convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities.

Further, the manipulations performed are often referred to in terms, such as adding or comparing, which are commonly associated with mental operations performed by a human operator. No such capability of a human operator is necessary, or desirable in most cases, in any of the operations described herein which form part of the present invention; the operations are machine operations. Useful machines for performing the operations of the present invention include general purpose digital computers or other similar devices. In all cases there should be borne in mind the distinction between the method operations in operating a computer and the method of computation itself. The present invention relates to method steps for operating a computer in processing electrical or other (e.g., mechanical, chemical) physical signals to generate other desired physical signals.

The present invention also relates to apparatus for performing these operations. This apparatus may be specially constructed for the required purposes or it may comprise a general purpose computer as selectively activated or reconfigured by a computer program stored in the computer. The algorithms presented herein are not inherently related to a particular computer or other apparatus. In particular, various general purpose machines may be used with programs written in accordance with the teachings herein, or it may prove more convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these machines will appear from the description given below.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features and advantages of the method and apparatus of the present invention will be apparent from the following detailed description of the invention in which:

FIG. 1 is a flow diagram illustrative of the prior art ordered dithered process.

FIG. 2 illustrates the artifacts which occur when area dither is performed on vectors.

FIG. 3 depicts an illustrative computer system employed in the system of the present invention.

FIG. 4 is a block diagram of a preferred embodiment of the system of the present invention.

FIGS. 5a, 5b, 5c and 5d illustrate the Gupta-Sproull anti-aliasing method employed in the preferred embodiment of the present invention.

FIGS. 6a, 6b and 6c illustrate the dither model and back-transform model employed in the preferred embodiment of the present invention.

FIGS. 7a, 7b, 7c and 7d set forth an error analysis with respect to the execution order of the steps of compositing and dithering.

FIGS. 8a, 8b, 8c and 8d illustrate the vector and corresponding antialiased vector used in the numerical examples.

FIGS. 9a, 9b, 9c, 9d and 9e illustrate a first numerical example utilizing the preferred embodiment of the present invention.

FIGS. 10a, 10b, 10c, 10d and 10e illustrate a second numerical example utilizing the preferred embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

GENERAL SYSTEM CONFIGURATION

Fig. 3 shows a typical computer-based system for the dithering of digital images according to the present invention. Shown there is a computer 1 which comprises three major components. The first of these is the input/output (I/O) circuit 2 which is used to communicate information in appropriately structured form to and from the other parts of the computer 1. Also shown as a part of computer 1 is the central processing unit (CPU) 3 and memory 4. These latter two elements are those typically found in most general purpose computers and almost all special purpose computers. In fact, the several elements contained within computer 1 are intended to be representative of this broad category of data processors. Particular examples of suitable data processors to fill the role of computer 1 include machines manufactured by Sun Microsystems, Inc., Mountain View, California. Other computers having like capabilities may of course be adapted in a straightforward manner to perform the functions described below.

Also shown in Fig. 3 is an input device 5, shown in typical embodiment as a keyboard. It should be understood, however, that the input device may actually be a card reader, magnetic or paper tape reader, or other well-known input device (including, of course, another computer). A mass memory device 6 is coupled to the I/O circuit 2 and provides additional storage capability for the computer 1. The mass memory may include other programs and the like and may take the form of a magnetic or paper tape reader or other well known device. It will be appreciated that the data retained within mass memory 6, may, in appropriate cases, be incorporated in standard fashion into computer 1 as part of memory 4.

In addition, a display monitor 7 is illustrated which is used to display messages or other communications to the user. Such a display monitor may take the

form of any of several well-known varieties of CRT displays. Preferably, the display monitor 7 displays the graphic images generated according to the process of the present invention. A cursor control 8 is used to select command modes and provides a more convenient means to input information into the system.

PROCESS DESCRIPTION

A preferred embodiment of the present invention is illustrated by the block diagram of Fig. 4. In the following example used to describe a preferred embodiment of the present invention, a 4 bit/color component per pixel frame buffer is used. It will be apparent from the following description that the present invention is not limited to specific frame buffer resolution. In addition, the present invention is equally applicable to greyscale frame buffers and color frame buffers.

Although the quality of a digital image improves when higher resolution frame buffers (e.g., 8 or 12 bits/color component) are used, the dramatic increase in cost for the additional memory is generally prohibitive. Therefore in order to keep memory costs within acceptable limits, the digital image is initially rendered at a higher resolution (in the present example, 8 bits/color component) and dithered to a lower resolution (e.g., 4 bits/color component).

Referring to Fig. 4, At block 100, the vector data to be displayed is modified to produce the pixel image of the corresponding anti-aliased vector. The anti-aliased vector is rendered with an 8 bit/color component intensity resolution. Although there is no limitation as to the anti-aliasing technique to be employed, it is preferred that Gupta-Sproull method is used. In this method, a pixel is viewed as the center of a conic filter of both unit radius and height and therefore of unit volume (See Fig. 5a). When a unit wide line is drawn within a perpendicular distance of 1 pixel from its edge to any pixel center (or 1.5 pixels from the pixel center to the line center), it carves out a volume from the conic filter centered at that pixel. The final intensity of that pixel is

then proportional to the volume intersected. Fig. 5b illustrates that for lines of unit thickness in the first octant (i.e., $0 \leq y \leq x$), three pixels in each column are utilized. Figs. 5c and 5d sets forth illustrative pseudo code which implements the Gupta-Sproull method. Instead of calculating the volume intersection for each pixel, it is preferred that the calculation of the volume intersection is replaced by a look-up table, whereing the table is generated by the pseudo code of Fig. 5c. The index into the table is the perpendicular distance from the center of the pixel to the center of the line. Fig. 5d illustrates the pseudo code for the algorithm. For further information, see Gupta, Sproull, "Filtering Edges for Gray Scale Displays", Computer Graphics, Vol 15, No. 3, August, 1981.

The pixel data representing the anti-aliased vector is, at block 110, composited with the corresponding pixel data contained in the frame buffer. The compositing process may be any known compositing process that blends a foreground color and intensity into the background color and intensity to provide a smoothing effect or transition between the foreground and background. The corresponding pixel data retrieved from the frame buffer are the pixels at the coordinate locations the anti-aliased vector is to be placed in the frame buffer. It is preferred that the compositing process uses an additional component, referred to as the alpha component. The alpha component contains a value determinative of the proportion of the incoming color that is to be blended into the color of pixel retrieved from the frame buffer (See, Carpenter, "The A-Buffer, an Anti-aliased Hidden Surface Method", Computer Graphics, Volume 18, No.3, July 1984). Preferably one of the following compositing equations is utilized:

$$C_{new} = (\alpha * C_{in}) + ((1 - \alpha) * C_{buff}) \quad [1]$$

$$C_{new} = \min (255, (\max (0, (\alpha * (C_{in} - C_{bg})) + C_{buff}))) \quad [2]$$

where: C_{new} == resultant color to be placed in the frame buffer
 C_{in} == color of incoming vector

C_{bg} == a constant user-specified background color

C_{buff} == color read back from frame buffer

α == proportion of incoming color to blend

Equation [1] is an alpha blend equation as described in Duff, "Compositing 3-D Rendered Images", Computer Graphics, Vol 19, No. 3, July, 1985. Equation [2] is an additive model which historically originates from stroke vector displays in which the background color of digital image in the frame buffer is considered to be a constant value.

It has been determined the steps of compositing and dithering must be performed in a predetermined sequence in order to render images of the highest possible quality. Specifically, it has been found that compositing should be performed prior to dithering such that the compositing process is executed at the higher resolution. By compositing at the higher resolution, the additional margin of error which occurs when compositing at a lower resolution is avoided. Because the compositing error is performed in floating point format, a round-off error is introduced when the data is converted back to integer format. The fractional error introduced during round-off is significantly less when the floating point compositing calculation is rounded to the nearest integer in the 8 bit domain than in the 4 bit domain. Figs. 7a, 7b, 7c and 7d set forth an error analysis performed with respect to the sequence of execution. In Figs. 7a and 7b, the foreground and background values are maintained constant and the alpha value is varied. The foreground and background colors in the example set forth in Fig. 7a are low intensity values and the foreground and background colors in the example set forth in Fig. 7b are high intensity values. In Figs. 7c and 7d, the intensity values of the foreground pixels are varied while the intensity value of the background value and the alpha value are maintained constant. As can be seen by examination of the graphs reflecting the % error as a function of the non-constant variable, the order of compositing first and then dithering (identified respectively as 300, 320, 340 and 360) consistently produced resulting values at or about 0% error. These results are

significantly better than the results (In Fig. 7a, 7b, 7c and 7d, lines 310, 330, 350 and 370) of % error for the order of dithering prior to compositing.

However, the quality of the digital image would not be maintained during the compositing process if the vector, rendered and anti-aliased at the first higher resolution, was composited with the lower resolution image stored in the frame buffer. Therefore, at block 140, it is necessary that the data read from the frame buffer is back-transformed to the higher resolution representation prior to compositing.

The back-transformation process corresponds to the specific dithering process utilized in order that the dithered values are accurately transformed back to their original intensity values.

The back-transformation process for the example is illustrated by Figs. 6a, 6b and 6c. Fig. 6a represents a functional definition of a 8 bit to 4 bit non-linear dither quantization model. The corresponding back-transformation function may simply be a table look-up function wherein the pixel's quantized intensity value is the index to the table. In the present example, the table is generated according to the equations set forth in Fig. 6b. Alternatively, to avoid the need of extra memory that would be required to store the back-transformation table, the back-transformation may be implemented in hardware and the back-transformed value computed when required. For example, the back-transformation may be computed, according to the functions set forth in Fig. 6c, using comparators, adders and shift registers.

At block 120, the composited image is dithered to a 12 bit/color representation for subsequent storage in frame buffer 130. Any dither process may be used, such as

the ordered dithered process earlier described.

In the present illustration, the digital image stored in the frame buffer 130 is subsequently output to a display device such as a computer monitor through a lookup table 150 which outputs the corresponding 8 bits/color component pixel data to the circuitry which controls the color intensities on the display device.

The following discussion describes two numerical examples which employ the preferred embodiment of the present invention. In the tables that are referenced, each row corresponds to one rendered pixel for the vector. The rows (and hence pixels) are presented in the order that would be drawn. The tables' column headings are defined as follows:

X, Y	The x,y coordinate of the pixel.
INDEX	The index (or indices into the dither matrix).
D(INDEX)	The dither matrix value at the specified index.
G_{in}	The 8-bit grey level (0-255) of the pixel. This is the value input to the dithering process.
BASE_{in}	The base value to which G _{in} is compared to determine which quantization range it falls into.
BASE_{out}	The corresponding 4-bit base value for generating dithered output. Depending upon the dither comparison result, either BASE _{out} or BASE _{out} + 1 is placed in the frame buffer.

Gdisplay The 8-bit grey level (0-255) which actually observed, produced by mapping the 4-bit frame buffer value to 8 bits via an output lookup table.

In dithering 24-bit RGB data (8 bits/color component) to 12-bit RGB data (4 bits/color component), each color component (R, G and B) is dithered identically but independently. For simplicity, the examples described illustrate the process with respect to only one component (or grey level), but the process is equally applicable to all of the RGB components. Similarly, the following examples illustrate the process with respect to dithering from 8 bits/color component to 4 bits/color component, but it is apparent to one in the art that the invention is not limited to a specific bit resolution or order of dither.

A first numerical example is discussed with respect to Figs. 8a, 8b, 8c and 8d and Figs. 9a, 9b, 9c and 9d. In this example a vector is drawn from (x=5, y=25) to (x=25, y=15). The grey (intensity) level is interpolated along the span of the vector, beginning at 0 and ending at 255. Fig. 8a is a table describing the vector as it would be drawn if a simple Bresenham's line drawing algorithm was utilized and Fig. 8b graphically illustrates the vector (For information on line drawing methods, including Bresenham's algorithm, See: Foley, Van Dam, Fundamentals Of Interactive Computer Graphics, pp 432-436 (Addison, Wesley 1984).

Figs. 8c and 8d illustrate the same vector anti-aliased using the Gupta-Sproull anti-aliasing technique and, for purposes of illustration, initially composited to a black background.

The initial background intensity is illustrated in Fig. 9a. The background intensity value of 40 is dithered using the screen-aligned linear dither process described in the Art Background section of this Specification in conjunction with the dither matrix illustrated in Fig. 9b to produce a pattern consisting of the values 34 and 51 (22 and 33 hexadecimal).

The anti-aliasing algorithm employs the alpha blend rule for compositing intensities (Equation [1]). Since this example composites with background values which were linearly dithered, the back-transformation equation applied to the background values prior to compositing is: $BT = i * 17$, where BT is the 8-bit back-transformed value and i is the 4-bit quantized value read back from the dithered background.

The pixels of the composited vector are then dithered using the vector aligned ditherin process (described in copending US Patent Application Serial No. 07/413,984, entitled "Method and Apparatus for Vector Aligned Dithering"), using the dither matrix illustrated in Fig. 9c, resulting in the values set forth in table of Fig. 9d and graphically depicted in Fig 9e.

A second numerical example is discussed with respect to Figs. 8a, 8b, 8c and 8d and Figs. 10a, 10b, 10c, 10d and 10e. As with the previous numerical example, a vector is drawn from (x=5, y=25) to (x=25, y=15). The grey (intensity) level is interpolated along the span of the vector, beginning at 0 and ending at 255. Fig. 8a is a table describing the vector as it would be drawn if a simple Bresenham's line drawing algorithm was utilized and Fig. 8b graphically illustrates the vector

Figs. 8c and 8d illustrate the same vector anti-aliased using the Gupta-Sproull anti-aliasing technique and, for purposes of illustration, initially composited to a black background.

The initial background intensity is illustrated in Fig. 10a. In this example, the background intensity value of 40 is dithered using a non-linear screen-aligned dither process (the non-linear dither process is described in co-pending U.K. Patent Application GB2,236,464A, entitled "Method and Apparatus for Non Linear Dithering of Digital Images") in conjunction with the dither matrices depicted in Fig. 10b to produce a pattern consisting of the values 36 and 45 (24 and 2d hexadecimal).

The anti-aliasing algorithm employs the alpha blend rule for compositing intensities (Equation [1]). Since this example composites with background values which were non-linearly dithered, the back-transformation equation applied to the background values prior to compositing is:

$$BT = \begin{array}{ll} i \cdot 9 & \text{for } 0 \leq i \leq 7 \\ ((i - 6) \cdot 17) + 54 & \text{for } 7 \leq i \leq 13 \\ ((i - 12) \cdot 33) + 156 & \text{for } 13 \leq i \leq 16 \end{array}$$

where BT is the 8-bit back-transformed value and i is the 4-bit quantized value read back from the dithered background.

The pixels of the composited vector are then dithered using a non-linear vector aligned ditherin process (described in copending UK Patent Application GB 2,236,464A, entitled 'Method and Apparatus for Non-Linear Dithering of Digital Images'). using the dither matrices illustrated in Fig. 10c. Each of the nonlinear dither matrices is applied only to a limited range of intensities. The first dither matrix is used in the range 0-54, the second dither matrix is used for input intensity values in the range of 157-255. The resulting values are set forth in table of Fig. 10d and graphically depicted in Fig. 10e.

While the invention has been described in conjunction with the preferred embodiment, it is evident that numerous alternatives, modifications, variations and uses will be apparent to those skilled in the art in light of the foregoing description.

CLAIMS

1. A method for rendering a vector represented at a first resolution of intensity on a low resolution device at a second lower resolution of intensity, a store of said low resolution device storing a background pixel image, said vector to be rendered between first and second coordinates in the screen space of the low resolution device, said method comprising the steps of:

anti-aliasing the vector at the first resolution to produce a pixel image of the anti-aliased vector;

compositing the said pixel image of the anti-aliased vector with the corresponding pixels of the background image stored in the low resolution device to produce a composited image, said corresponding pixels being those pixels of the background image which form a vector between said first and second coordinates of the screen space, said step of compositing comprising the steps of;

back-transforming the corresponding pixels of the background image from the second resolution to the first resolution;

compositing the pixel image of the anti-aliased vector with the back-transformed pixels of the background image;

dithering the composited image from the first resolution to a second resolution;

storing the dithered image at the second resolution to the store of the low resolution device.

2. The method of claim 1 wherein the vector is a single pixel in width and is anti-aliased to a vector three pixels in width.

3. The method of claim 1 wherein the store of the low resolution device is a frame buffer and the vector is



composited with the background image according to the following equation:

$$C_{new} = (\alpha * C_{in}) + ((1-\alpha) * C_{buff})$$

where: C_{new} is the resultant color to be placed in the frame buffer, C_{in} is color of the vector to be rendered, C_{buff} is the colour read back from frame buffer and α is the proportion of the vector color to blend into the background color.

4. The method of claim 1 wherein the store of the low resolution device is a frame buffer and the vector is composited with the background image according to the following equation:

$$C_{new} = \min (255, (\max (0, (\alpha * (C_{in} - C_{bg})) + C_{buff})))$$

where C_{new} is the resultant color to be placed in the frame buffer, C_{in} is the color of incoming vector, C_{bg} is a constant value user-specified background color, C_{buff} is the color read back from frame buffer and α is the proportion of vector color to blend into the background color.

5. The method of claim 1 wherein the composited image is dithered using a vector aligned dithering process.

6. The method of claim 5 wherein the composited image is dithered using a non-linear dithering process.

7. The method of claim 6 wherein the store of the low resolution device is a frame buffer and the first resolution is 8 bits/color component and the second resolution is 4 bits per color component and the composited image is dithered such that:



a low range dither matrix is established for input intensity values between the values of zero and 54 and having output values between zero and 6;

a mid range dither matrix is established for input intensity values between the values of 54 and 156 and having output values between 6 and 12.

a high range dither matrix is established for input intensity values between the values of 156 and 255 and having output values between 12 and 15; and

the intensity values of the corresponding pixels of the background image are back-transformed according to the following equation:

$$\begin{aligned} \text{Back-Transform}(i) = & q(i)*9 && \text{for } 0 \leq q(i) < 7 \\ & ((q(i)-6)*17)+54 && \text{for } 7 \leq q(i) < 13 \\ & ((q(i)-12)*33)+156 && \text{for } 13 \leq q(i) < 16 \end{aligned}$$

where $q(i)$ is the intensity values of a corresponding pixel stored in the frame buffer.

8. An apparatus for rendering a vector represented at a first resolution of intensity on a low resolution device at a second lower resolution of intensity, a store of said low resolution device storing a background pixel image, said vector to be rendered between first and second coordinates in the screen space of the low resolution device, said apparatus comprising:

means for anti-aliasing the vector at the first resolution to produce a pixel image of the anti-aliased vector;

means for compositing the pixel image of the anti-aliased vector with the corresponding pixels of the background image stored in the low resolution device to produce a composited image, said corresponding pixels being those pixels of the background image which form a vector between said first and second coordinates of the screen space, said means for compositing comprising the steps of;

means for back-transforming the corresponding pixels of the background image from the second resolution



to the first resolution;

means for compositing the pixel image of the anti-aliased vector with the back-transformed pixels of the background image;

means for dithering the composited image from the first resolution to a second resolution;

means for storing the dithered image at the second resolution to the store of the low resolution device.

9. The apparatus of claim 8 wherein the vector is a single pixel in width and is anti-aliased to a vector three pixels in width.

10. The apparatus of claim 8 wherein the store of the low resolution device is a frame buffer and the vector is composited with the background image according to the following equation:

$$C_{new} = (\alpha * C_{in}) + ((1 - \alpha) * C_{buff})$$

where: C_{new} is the resultant color to be placed in the frame buffer, C_{in} is color of the vector to be rendered, C_{buff} is the colour read back from frame buffer and α is the proportion of the vector color to blend into the background color.

11. The apparatus of claim 8 wherein the store of the low resolution device is a frame buffer and the vector is composited with the background image according to the following equation:

$$C_{new} = \min(255, (\max(0, (\alpha * (C_{in} - C_{bg})) + C_{buff})))$$

where: C_{new} is the resultant color to be placed in the frame buffer, C_{in} is the color of incoming vector, C_{bg} is



a constant value user-specified background color, C_{buff} is the color read back from frame buffer and alpha is the proportion of vector color to blend into the background color.

12. The apparatus of claim 8 wherein the composited image is dithered using a vector aligned dithering process.

13. The apparatus of claim 12 wherein the composited image is dithered using a non-linear dithering process.

14. The apparatus of claim 13 wherein the store of the low resolution device is a frame buffer, the first resolution is 8 bits/color component and the second resolution is 4 bits color component and the composited image is dithered such that:

a low range dither matrix is established for input intensity values between the values of zero and 54 and having output values between zero and 6;

a mid range dither matrix is established for input intensity values between the values of 54 and 156 and having output values between 6 and 12.

a high range dither matrix is established for input intensity values between the values of 156 and 255 and having output values between 12 and 15; and

the intensity values of the corresponding pixels of the background image are back-transformed according to the following equation:

$$\text{Back-Transform}(i) = \begin{cases} q(i)*9 & \text{for } 0 \leq q(i) < 7 \\ ((q(i)-6)*17)+54 & \text{for } 7 \leq q(i) < 13 \\ ((q(i)-12)*33)+156 & \text{for } 13 \leq q(i) < 16 \end{cases}$$

where $q(i)$ is the intensity values of a corresponding pixel stored in the frame buffer.

15. A method for rendering a vector represented at a first resolution of intensity on a low resolution device at a second lower resolution of intensity, a store of



said low resolution device storing a background pixel image, said vector to be rendered between a first and second coordinates in the screen space of the low resolution device, substantially as hereinbefore described with reference to the accompanying drawings.

16. An apparatus for rendering a vector represented at a first resolution of intensity on a low resolution device at a second lower resolution of intensity, a store of said low resolution device storing a background pixel image, said vector to be rendered between a first and second coordinates in the screen space of the low resolution device, substantially as hereinbefore described with reference to the accompanying drawings.

