

- 1 -

METHOD AND APPARATUS FOR NON-LINEAR DITHERING OF DIGITAL IMAGESBACKGROUND OF THE INVENTION1. Field of the invention:

The method and apparatus of the present invention relates to dithering techniques used to represent digital images. More particularly, the method and apparatus of the present invention relates to the non-linear dithering of digital images.

5

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
10 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 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
5 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
10 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
15 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:

$$i = x \bmod n$$

$$j = y \bmod n$$

20

If n is a power of two, then 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.

25

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
30 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 that is 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.

The dithering process is applied to all portions of the digital image for all intensities. However, it is known that the eye is more sensitive to variations at lower intensities than at higher intensities. Therefore results of the dithering process at lower intensities are not as effective as the results are at
5 higher intensities.

SUMMARY OF THE INVENTION

It is therefor an object of the present invention to provide a method and apparatus for dithering digital images using non-linear quantization.

5

It is an object of the present invention to provide a method and apparatus for dithering in which the artificial variations in intensity the eye perceives are minimized.

10

It is further an object of the present invention to provide a method and apparatus for dithering digital images in which smaller quantization levels are used at lower intensities than at higher intensities, thereby accommodating for the greater sensitivity of the eye at lower intensities than at higher intensities.

15

The method and apparatus of the present invention takes into consideration that the eye is more sensitive to intensity variations at low levels than at high levels and produces higher quality dithered images by employing a sectioned or piecewise linear quantization strategy which sacrifices the fineness of interval spacing at the high end of the intensity scale in order to achieve greater accuracy of reproduction at the low end.

20

In a preferred embodiment, the input intensity values are divided into three ranges, a low range, medium range and high range. The spacing between output intensity levels at the low range is significantly smaller than the spacing between intensity levels at the medium range and high range. The medium range preferably has smaller spacing between intensity values than in the high range. In conjunction with the different interval spacing in the different ranges, separate dither matrices are employed providing the range of intermediate intensity values to be simulated.

25

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
5 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
10 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
5 description of the invention in which:

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

10 FIG. 2 depicts the illustrative computer system employed in the system of the present invention.

FIG. 3 illustrates a logarithmic function representing the sensitivity of the eye to varying intensity levels.

15 FIG. 4 is a flowchart depicting the preferred embodiment of the process of the present invention.

FIGS. 5a, 5b and 5c illustrate the sub-ranges of a preferred
20 embodiment of the present invention.

FIGS. 6a, 6b and 6c illustrate exemplary dither matrices to be used with the preferred embodiment of the present invention.

25 FIGS. 7a, 7b, 7c and 7d illustrate a numerical example of an aliased vector employing a preferred embodiment of the present invention.

FIGS. 8a, 8b, 8c and 8d illustrate a numerical example of an anti-aliased vector employing a preferred embodiment of the present invention.

30

DETAILED DESCRIPTION OF THE INVENTION

GENERAL SYSTEM CONFIGURATION

5 Fig. 2 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
10 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
15 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.

20 Also shown in Fig. 2 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
25 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.

30

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

10

In the method and apparatus of the present invention a non-linear quantization model provides for smaller intensity interval spacing at the lower intensities at the cost of an increase in interval spacing at the higher intensities.

15

The range of input intensity values is divided into sections representing low intensity ranges, medium intensity ranges and high intensity ranges. The method of selecting the number of ranges to be employed is implementation dependent. In the following example described in which an 8 bit/color representation is dithered to 4 bit/color, 3 linear ranges have been used to approximate a logarithmic function. The choice of 3 ranges was based on the fact that the number of stored dithered values in an 8 to 4 bit dither per color component scheme is only 16. Using more than 3 ranges would have resulted in a better approximation to the desired logarithmic function but would also have required a greater amount of hardware and computational overhead with little significant improvement in image quality. However, dithering from an 8 to 5 bits per color component or higher might have justified the increase in hardware complexity and overhead for a greater number of ranges.

The number of ranges and the number of intensity levels in each range are selected such that the totality of the ranges encompasses the

30

complete range of input intensity levels. Preferably the sections are selected to closely emulate the logarithmic function which describes the eyes' sensitivity to changes in intensity. For information on eye sensitivity to intensity see Pratt, Digital Image Processing, pp 32-47 (Wiley), Gonzalez and Wintz, Digital Image
5 Processing, pp 13-21 (Addison-Wesley 1977); Netravali and Prasada, "Adaptive Quantization of Picture Signals Using Spatial Masking", Proceedings of the IEEE, Vol. 65, No. 4, (April 1977). An exemplary logarithmic function of the eyes' sensitivity to changes in intensity is shown in Fig. 3. Fig. 3 shows the functions of the ratio of changes in intensity ΔI to intensity values I for a variety of
10 background intensity values I_0 .

Thus, for example, a logarithmic function of the eye sensitivity to changes in intensity is applied to an input range values which span from 0 to 255. Two intermediate points close to the logarithmic function are selected. A
15 first straight line between 0 and the first point selected and is determined representative of one intensity range. A second straight line is determined between the two intermediate points which is representative of a second intensity range. A third intensity range corresponds to the straight line drawn between the second selected point and 255. Ideally, the position and length of
20 the lines are selected so as to minimize the error between the desired logarithmic function and the piecewise (straight-line) approximation. Any straight line approximation technique, such as linear regression, may be used to minimize the approximation error.

25 The ranges may also be divided according to the locality of concentration of intensity values in the digital image. For example, it has been found that a great proportion of intensity values when the image is depth cued into a black background reside in the range of 0-100. Thus it is desirable to increase the number of quantized intensity values in the low range at the cost of
30 decreasing the number of quantized intensity values in the high range.

Optimally, the intensity range should be divided into a large number of sub-ranges because the larger the number of sub-ranges, the more closely the eye sensitivity function may be emulated. However, the more sub-ranges utilized, the greater the CPU overhead and degradation in processing speed. Therefore, a balance between the number of sub-ranges and processing speed must be achieved. For instance, in the following example illustrated by Figs. 5 and 6, the input range of intensities is divided into three sub-ranges - a low intensity range, a medium intensity range and a high intensity range, wherein input intensity values 0-54 are allocated to the low intensity sub-range, values 55-156 are allocated to the medium intensity sub-range and values 157-255 are allocated to the high intensity sub-range.

A dither matrix is then created for each of the sub-ranges, wherein the values in the dither matrix represent the intermediate intensity values between quantized output intensity levels. The dither matrix is a small two dimensional set of values which is used to determine which quantized intensity value will be output at each position. The values contained in the dither matrix comprise the intermediate intensity values between quantized levels. Thus, for example, if the quantized levels are separated by an interval spacing of 15, the dither matrix would contain the values 0 to 14. However, depending on such considerations as hardware configurations and memory allocations, the dither matrix can be modified slightly to not include all intermediate intensity values or to include more than one occurrence of each of the intermediate intensity values. For further information on the formation of dither matrices see Ulichney, Digital Halftoning (MIT Press, 1988); Foley and Van Dam, Principles of Interactive Graphics, pp. 597-602 (Addison-Wesley 1984).

The dithering process is then initiated for each of the pixels wherein the input intensity values are quantized to the quantized output

intensity levels in such a manner that the eye integrates over areas of quantized intensity levels and perceives the intermediate intensity levels desired.

A preferred embodiment of the present invention is described with
5 respect to the flow chart of Fig. 4. At block 200 the range of input intensity values are divided into sub-ranges indicative of low intensities, mid-range intensities and high intensities. Preferably the number of sub-ranges used is limited to three thereby minimizing the complexity of the process while
10 maintaining the advantages of providing different interval spacing at different sub-ranges of intensities. The size of the ranges may be equal in size to cover the full range of input intensity ranges but it is preferred that the lower intensity range(s) is smaller than the high and medium intensity ranges to provide for the smaller interval spacing between quantized output levels.

15 At block 210 the total number of quantized output levels available are divided among the sub-ranges. Preferably the quantized levels are divided in such a manner that a larger number of quantized levels are associated with the low intensity range than the high intensity range. The larger number of quantized levels in the low intensity range, the smaller the interval spacing and
20 therefore the greater the sensitivity to variations in intensity. The number of quantized levels per sub-range and the size of the sub-range determines the interval spacing between quantized output intensity levels and therefore the amount of sensitivity to variations in intensity levels.

25 At block 220 the quantized levels assigned to each sub-range are arranged within each sub-range. Preferably, the quantized levels assigned to each sub-range are linearly arranged within each sub-range such that within each range the quantized intensity levels are equally spaced apart to simplify the dither matrices utilized.

At block 230 a separate dither matrix is prepared for each of the sub-ranges and at block 240 the input intensity values are dithered to the quantized output intensity levels. The process is not limited to the type of dither process; however, the dither matrix used in the process must correspond to the sub-range the input intensity level is located within. An exemplary dither process is the ordered dither process described with respect to Fig. 1 in the art background section of the Specification.

The system of the present invention may best be described by an example of dithering an 8 bit intensity representation to a 4 bit intensity representation. It is preferred that the full input value range is mapped into the entire output value range, implying that input level 0 must map to output level 0, and input level 255 must map to output level 15. In the example illustrated by Figs. 5 and 6b output intensity levels are allocated to input intensities in the range 0 to 54, 6 levels to the mid-range 54-156 and 3 levels to the high range 156-255. The quantization function is illustrated in tabular form in Fig. 5a graphically in Fig. 5c and functionally in Fig. 5b.

The dither matrices used in the present example are shown in Figs. 6a, 6b and 6c. As stated previously, because the number of input intensities between each quantization level is different for each of the three ranges - 9 for the first range, 17 for the second range and 33 for the third range - a different dither matrix must be used for dithering values that fall into each of the three ranges.

Typically, a dither matrix should contain one element per intermediate intensity to be dithered, suggesting that the three required matrices should possess 9, 17 and 33 entries respectively. However, it is preferred, for purposes of optimal hardware design of this example, that a matrix comprises an even number of entries, more preferably a number of entries that are powers

of two. In the present example each of the values are off a power of two by one. Thus if one is subtracted from each, the number of intermediate intensity levels between quantized levels are powers of two: $9-1 = 8 = 2^3$, $17-1 = 16 = 2^4$, $33-1 = 32 = 2^5$.

5

Preferably the matrices employed for each of the ranges is of the same size such as the matrices illustrated in Figs. 6a, 6b and 6c. Fig. 6a shows a third order matrix for the low intensity range of the present example containing values in the range of 0-7. For ease of indexing, there are two periods of the pattern replicated within the 4x4 matrix. Fig. 6b shows a dispersed dot matrix
10 having 16 dither values in the range of 0-15 used for the mid-range intensity values.

The matrix for the high range of input values, shown in Fig. 6c, is the middle range matrix multiplied by two. The ideal matrix is a fifth order diamond matrix containing values 0-31, and would require a 64 entry (8x8)
15 table for storage. In order to keep the storage requirements low, and also to maintain uniform indexing for all matrices, the smaller 4x4 table is preferred. Although the increment of two between dither matrix values causes pairs of input values to be mapped to identical output values, the visual effect is not perceived because the eye is much less sensitive to intensity variations at high
20 intensities.

Once the sub-ranges and corresponding matrices have been established, the input intensity levels are dithered. An exemplary dithering process is the area dithering process set forth in Fig. 1.

Figs. 7a, 7b, 7c and 7d illustrate a first numerical example employing the method and apparatus of the present invention. Figure 7a shows a shaded aliased vector drawn from (x,y) coordinates (5, 25) to (25, 15) wherein the intensity level is interpolated along the span of the vector, beginning at 0
5 and ending at 255.

The dithering process is performed using the dither matrices shown on Fig. 7b. Each of the non-linear dither matrices is applied only to a limited range of intensities. More particularly, the first dither matrix (low intensity range matrix) is used for intensity values in the range 0-54, the second matrix
10 (medium intensity range matrix) is used for intensity values in the range 55-156 and the third matrix (high intensity range matrix) is used for values in the range 157-255.

The results of the computations performed according to the process of the present invention are set forth in the table of Fig. 7c, where
15 $G_{display}$ is the dithered output value, and graphically represented by the illustration set forth in Fig. 7d.

Figs. 8a, 8b, 8c and 8d set forth a second numerical example in which an aliased vector three pixels in width is dithered using the method and apparatus of the present invention. The vector align dither process is again used in the present example on the vector illustrated in Fig. 8a, which is a shaded anti-aliased vector drawn from (x,y) coordinates (5, 25) to (25, 15). The intensity levels of the vector are interpolated along the span of the vector, beginning at the value of 0 and ending at 255.

The dithering process is performed using the dither matrices shown in Fig. 8b. Each dither matrix is a 3x16 matrix comprising the number of intermediate intensity values between quantized intensity levels. The 3x16 matrices are formed by three separate lines each containing a single occurrence of the number of intermediate intensity values wherein the emphasis on the arrangement of the values is to ensure that each of the three-entry columns sums up to approximately the same value. More particularly, the first matrix, the low intensity range matrix, consists of the values 0-7 reflective of the small interval spacing (8) in the low intensity range. The values are arranged such that the average intensity is maintained over a local area because there are only eight intermediate values and the width of the matrix is 16, the values are repeated to fill the matrix. Similarly, the second matrix, the medium range intensity matrix, consists of values 0-15 reflective of a slightly larger interval spacing (16).

The third dither matrix, the high intensity range matrix, consists of intermediate values 0 to 31 indicative of the large interval spacing (32). In order to fit the range of values into the 16 column matrix, the even values between 0 and 31 are utilized. As noted earlier, the increment of two between dither matrix values will cause pairs of input values to be mapped to identical output values, but such drawback is minimal because the eye sensitivity to intensity variations is greatly diminished at the higher intensities.

These dithered values may then be stored until the corresponding pixels are output to the look-up table for display on a display device. An output mapping function is generated which maps back the dithered intensity values to their original value. For example, the following functions are employed to map back the examples described above:

$$\text{Out}(q(i)) = \begin{cases} q(i) \times 9 & \text{for } 0 \leq q(i) \leq 6 \\ \left((q(i) - 6) \times 17 \right) + 54 & \text{for } 6 < q(i) \leq 12 \\ \left((q(i) - 12) \times 33 \right) + 156 & \text{for } 12 < q(i) \leq 15 \end{cases}$$

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. Particularly it is evident that the invention is equally applicable to color images as it is to greyscale images. Furthermore, it is evident that the invention described is not limited to an 8 bit to 4 bit dithering process but is applicable to any dithering scheme.



CLAIMS

1. A method for dithering a digital image from an input intensity range to an output intensity range, said digital image comprising pixels, each pixel
5 having an intensity value associated with it selected from the input intensity range, wherein the digital image is dithered from a p-bit intensity value per pixel to a q-bit intensity value per pixel, where p is a value greater than q, said method comprising the steps of:
 - dividing the range of input intensity values into sections;
 - 10 assigning a section of the output intensity range to a corresponding section of the input intensity range whereby the elements in each output intensity section correspond to a quantized intensity value in the corresponding input intensity section in each section, the number of values between quantized intensity values being an interval spacing
15 for the section;
 - generating a dither matrix for each section of the input intensity range, the number of values in the dither matrix equal to the count of intermediate intensity values between quantized intensity values in the respective sections, the entries in each dither matrix equal to the
20 values in the range of 0 and the interval spacing for the section which are representative of the number of intermediate intensity values between quantized intensity levels, said elements organized to provide the desired visual perception of the intermediate intensity values;
 - retrieving a dither value from the dither matrix for each pixel value;
 - 25 determining the input-scale base value for a pixel value, said input-scale base value being the value less than or equal to the pixel value and which directly corresponds to a quantized intensity value;
 - determining a quantized intensity value which corresponds to the input-scale base value;



- 21 -

if the difference between the input intensity value and the input-scale base value is less than or equal to the retrieved dither value, setting the output intensity level for the pixel to be equal to the quantized intensity value;

5 if the difference between the input intensity value and the input-scale base value is greater than the retrieved dither value, setting the output intensity level for the pixel to be equal to a quantized intensity value which corresponds to the input-scale base value +1;

whereby the sections are divided and assigned to provide greater
10 accuracy in one section at the cost of accuracy in another section.

2. The method of claim 1 wherein the step of dividing the range of input intensity values into sections comprises:

15 selecting a function of eye sensitivity to changes in intensity;
 dividing the function into portions, the number of portions equaling the number of sections of the input intensity range;

 computing a straight line approximation for each of the portions;
 and

20 utilizing the straight line approximation to determine the elements of the output intensity range sections and the correspondence to the sections of the input intensity range.



- 22 -

3. The method of claim 2 wherein the output range of intensity values are selected according to the function of the eyes' sensitivity to changes in light intensity, the step of dividing the input range of intensity values comprises:

5 selecting two intermediate points along a curve representing the function of the eyes' sensitivity to changes in intensity;

determining the low intensity section to be the values along a straight line approximation between the beginning of the curve and the first intermediate point selected;

10 determining the mid intensity section of the output intensity range to be the values along a straight line approximation between the first intermediate value and the second intermediate value selected;

determining the high intensity section to be the values along a straight line approximation between the second intermediate value and the end of the curve.

15 whereby the output range of intensity values are proportionately assigned to the input intensity range.

4. The method of claim 1 wherein the step of dividing the range of input intensity value into sections comprises:

20 determining the sections of the range having a greater concentration of occurrences of intensity values; and

dividing the input intensity range such that the sections having the greater concentration of intensity values are smaller than the sections having a lesser concentration of intensity values.

25



- 23 -

5. The method of claim 4 wherein:

the digital image to be dithered is depth cued to a dark background; and

5 the greatest concentration of occurrence of intensity values is at the low intensity values;

whereby a section around the low intensity values is smaller than a section around the high intensity values.

6. The method of claim 5 wherein:

10 p equals 8 thereby providing an input intensity range of 256 values;

q equals 4 thereby providing an output intensity range of 16 values; and

15 the range of input intensity values is divided into a low intensity section mid-intensity section and high intensity section .

7. The method of claim 6 wherein:

the low intensity section of input intensity values extends from 0 to 54;

20 the mid intensity section of input intensity values extends from 55 to 156; and

the high intensity section of input intensity values extends from 157 to 255.



- 24 -

8. The method of claim 7 wherein:

the step of assigning a section of the output intensity range to a corresponding section of the input intensity range comprises;

5 assigning output intensity values 0 to 6 to the low intensity section of input intensity values;

 assigning output intensity values 7 to 12 to the mid intensity section of input intensity values;

 assigning output intensity values 13 to 15 to the high intensity section of input intensity values; and

10 the step of generating a dither matrix for each section comprises:

 generating a dither matrix corresponding to the low intensity section having elements 0 to 7;

15 generating a dither matrix corresponding to the mid intensity section having elements 0 to 15;

 generating a dither matrix corresponding to the high intensity section having elements 0 to 31.



- 25 -

9. The method of claim 8 wherein the digital image is dithered using
20 area dither and the dither matrices comprise the values:

0	7	1	6
4	2	5	3
1	6	0	7
5	3	4	2

for the matrix which corresponds to the low intensity section;

1	15	2	12
9	5	10	6
3	13	0	14
11	7	8	4



- 26 -

for the matrix which corresponds to the mid intensity section;

2	30	4	24
18	10	20	12
6	26	0	28
22	14	16	8

for the matrix which corresponds to the high intensity section.

5

10. The method of claim 8 wherein the digital image is dithered using vector aligned dither and the dither matrices comprise the values:

| 0 4 2 7 1 5 3 6 0 4 2 7 1 5 3 6 |

10

for the matrix which corresponds to the low intensity section;

| 0 8 4 10 2 12 6 14 1 9 5 11 3 13 7 15 |

for the matrix which corresponds to the mid intensity section;

15

| 0 16 8 20 4 24 12 28 2 18 10 22 6 25 14 30 |

for the matrix which corresponds to the high intensity section.



11. A method for dithering a digital image from an input range of intensity values to an output range of intensity values, the range of values determined by the number of bits representing each intensity value, said digital image comprising pixels, each pixel having an intensity value of p-bits
- 5 associated with it within the first range of intensity values, wherein the digital image is dithered from a p-bit representation to a q-bit representation, the value of p being greater than the value of q, said method comprising the steps of:
- dividing the input range of intensity values into sections representative of a low intensity section, mid intensity section and high intensity section;
- 10 dividing the output range of intensity values into sections representative of a low intensity section, mid intensity section and high intensity section to correspond respectively to the low intensity section, mid intensity section and high intensity section of the input intensity
- 15 range, wherein the elements in each output intensity section correspond to a quantized intensity value in the corresponding input intensity section;
- generating a dither matrix for each intensity section, the elements in each dither matrix equal to the values in the range between 0 and the interval spacing for the section, which are representative of the number of
- 20 intermediate intensity values between quantized intensity values in each section, said elements organized to provide the desired visual perception of the intermediate intensity values;
- retrieving a dither value from a dither matrix for each pixel value the dither matrix the dither value is retrieved from being dependent upon
- 25 the input intensity value of the pixel;
- determining the input-scale base value for a pixel value, said input-scale base value being the value less than or equal to the pixel value;
- determining a quantized intensity value which corresponds to the
- 30 input-scale base value;



- 28 -

if the difference between the input intensity value and the input-scale base value is less than or equal to the retrieved dither value, setting the output intensity level for the pixel to be equal to the quantized intensity value; and

5 if the difference between the input intensity value and the input-scale base value is greater than the retrieved dither value, setting the output intensity level for the pixel to be equal to a quantized intensity value which corresponds to the input-scale base value +1.

10 12. The method of claim 11 wherein a greater number output intensity values are located in the low intensity section than in the high intensity section to provide a higher precision of dithering for input intensity values within the low intensity section.

15 13. The method of claim 12 wherein the input range of intensity values are selected according to the function of the eyes' sensitivity to changes in light intensity, the step of dividing the input range of intensity values comprises:

 selecting a function of eye sensitivity to changes in intensity;

 selecting two intermediate points along a curve representing the
20 function of the eyes' sensitivity to changes in intensity;

 determining the low intensity section to be the values along a straight line approximation between the beginning of the curve and the first intermediate point selected;

 determining the mid intensity section to be the values along a
25 straight line approximation between the first intermediate value and the second intermediate value selected; and

 determining the high intensity section to be the values along a straight line approximation between the second intermediate value and the end of the curve.

30



- 29 -

14. The method of claim 11 wherein:

the digital image to be dithered is depth cued to a dark background;

the greatest concentration of occurrence of intensity values occurs within the low intensity section;

whereby the low intensity section is smaller than the mid and high intensity sections.

15. The method of claim 14 wherein:

p equals 8 thereby providing an input intensity range of 256 values;

q equals 4 thereby providing an output intensity range of 16 values;

the low intensity section of input intensity values extends from 0 to 54;

the mid intensity section of input intensity values extends from 55 to 156; and

the high intensity section of input intensity values extends from 157 to 255;



16. The method of claim 15 wherein:

the low section of output intensity values ranges from 0 to 6,
inclusive;

the mid section of output intensity values ranges from 7 to 12,

5 inclusive; and

the high section of output intensity values ranges from 13 to 15,
inclusive;

the dither matrix corresponding to the low intensity section having
elements in the range of 0 to 7;

10 the dither matrix corresponding to the mid intensity section having
elements in the range of 0 to 15; and

the dither matrix corresponding to the high intensity section having
elements in the range of 0 to 31.

15 17. The method of claim 16 wherein the digital image is dithered using
area dither and the dither matrices comprise the values:

0	7	1	6
4	2	5	3
1	6	0	7
5	3	4	2

for the matrix which corresponds to the low intensity section;

20

1	15	2	12
9	5	10	6
3	13	0	14
11	7	8	4

for the matrix which corresponds to the mid intensity section;



- 31 -

2	30	4	24
18	10	20	12
6	26	0	28
22	14	16	8

for the matrix which corresponds to the high intensity section.

18. The method of claim 16 wherein the digital image is dithered using
5 vector aligned dither and the dither matrices comprise the values:

| 0 4 2 7 1 5 3 6 0 4 2 7 1 5 3 6 |

for the matrix which corresponds to the low intensity section;

10

| 0 8 4 10 2 12 6 14 1 9 5 11 3 13 7 15 |

for the matrix which corresponds to the mid intensity section;

| 0 16 8 20 4 24 12 28 2 18 10 22 6 25 14 30 |

for the matrix which corresponds to the high intensity section.

15



19. An apparatus for dithering a digital image from an input intensity range to an output intensity range, said digital image comprising pixels, each pixel having an intensity value associated with it selected from the input intensity range, wherein the digital image is dithered from a p-bit intensity value per pixel to a q-bit intensity value per pixel, where p is a value greater than q, said apparatus comprising:

means for dividing the range of input intensity values into sections;
means for assigning a section of the output intensity range to a corresponding section of the input intensity range whereby the elements in each output intensity section correspond to a quantized intensity value in the corresponding input intensity section;

means for generating a dither matrix for each section of the input intensity range, the number of values in the dither matrix equal to the count of intermediate intensity values between quantized intensity values in the respective sections, the entries in each dither matrix equal to the values in the range between 0 and the interval spacing for the section which are representative of the number of intermediate intensity values between quantized intensity levels, said elements organized to provide the desired visual perception of the intermediate intensity values;

means for retrieving a dither value from the dither matrix for each pixel value;

means for determining the input-scale base value for a pixel value, said input-scale base value being the value less than or equal to the pixel value and which directly corresponds to a quantized intensity value;

means for determining a quantized intensity value which corresponds to the input-scale base value;

if the difference between the input intensity value and the input-scale base value is less than or equal to the retrieved dither value, means for setting the output intensity level for the pixel to be equal to the quantized intensity value;



- 33 -

if the difference between the input intensity value and the input-scale base value is greater than the retrieved dither value, means for setting the output intensity level for the pixel to be equal to a quantized intensity value which corresponds to the input-scale base value +1;

5 whereby the sections are divided and assigned to provide greater accuracy in one section at the cost of accuracy in another section.

20. The apparatus of claim 19 wherein the means for dividing the range of input intensity values into sections comprises:

10 means for selecting a function of eye sensitivity to changes in intensity;

 means for dividing the function into portions, the number of portions equaling the number of sections of the input intensity range;

15 means for computing a straight line approximation for each of the portions; and

 means for utilizing the straight line approximation to determine the elements of the output intensity range sections and the correspondence to the sections of the input intensity range.



21. The apparatus of claim 20 wherein the output range of intensity values are selected according to the function of the eyes' sensitivity to changes in light intensity, the means for dividing the input range of intensity values comprises:

5 means for selecting two intermediate points along a curve representing the function of the eyes' sensitivity to changes in intensity; means for determining the low intensity section to be the values along a straight line approximation between the beginning of the curve and the first intermediate point selected;

10 means for determining the mid intensity section of the output intensity range to be the values along a straight line approximation between the first intermediate value and the second intermediate value selected;

15 means for determining the high intensity section to be the values along a straight line approximation between the second intermediate value and the end of the curve.

whereby the output range of intensity values are proportionately assigned to the input intensity range.

20 22. The apparatus of claim 19 wherein the means for dividing the range of input intensity value into sections comprises:

means for determining the sections of the range having a greater concentration of occurrences of intensity values; and

25 means for dividing the input intensity range such that the sections having the greater concentration of intensity values are smaller than the sections having a lesser concentration of intensity values.



23. The method of claim 22 wherein:

the digital image to be dithered is depth cued to a dark background; and

5 the greatest concentration of occurrence of intensity values is at the low intensity values;

whereby a section around the low intensity values is smaller than a section around the high intensity values.

24. The apparatus of claim 23 wherein:

10 p equals 8 thereby providing an input intensity range of 256 values;

q equals 4 thereby providing an output intensity range of 16 values; and

15 the range of input intensity values is divided into a low intensity section mid-intensity section and high intensity section .

25. The apparatus of claim 24 wherein:

the low intensity section of input intensity values extends from 0 to 54;

20 the mid intensity section of input intensity values extends from 55 to 156; and

the high intensity section of input intensity values extends from 157 to 255.



26. The apparatus of claim 25 wherein:

the means for assigning a section of the output intensity range to a corresponding section of the input intensity range comprises;

5 means for assigning output intensity values 0 to 6 to the low intensity section of input intensity values;

means for assigning output intensity values 7 to 12 to the mid intensity section of input intensity values;

means for assigning output intensity values 13 to 15 to the high intensity section of input intensity values; and

10 the means for generating a dither matrix for each section comprises:

means for generating a dither matrix corresponding to the low intensity section having elements in the range of 0 to 7;

15 means for generating a dither matrix corresponding to the mid intensity section having elements in the range of 0 to 15;

means for generating a dither matrix corresponding to the high intensity section having elements in the range of 0 to 31.



- 37 -

27. The apparatus of claim 26 wherein the digital image is dithered
20 using area dither and the dither matrices comprise the values:

0	7	1	6
4	2	5	3
1	6	0	7
5	3	4	2

for the matrix which corresponds to the low intensity section;

1	15	2	12
9	5	10	6
3	13	0	14
11	7	8	4



- 38 -

for the matrix which corresponds to the mid intensity section;

2	30	4	24
18	10	20	12
6	26	0	28
22	14	16	8

for the matrix which corresponds to the high intensity section.

5

28. The apparatus of claim 27 wherein the digital image is dithered using vector aligned dither and the dither matrices comprise the values:

| 0 4 2 7 1 5 3 6 0 4 2 7 1 5 3 6 |

10

for the matrix which corresponds to the low intensity section;

| 0 8 4 10 2 12 6 14 1 9 5 11 3 13 7 15 |

for the matrix which corresponds to the mid intensity section;

15

| 0 16 8 20 4 24 12 28 2 18 10 22 6 25 14 30 |

for the matrix which corresponds to the high intensity section.



29. An apparatus for dithering a digital image from an input range of intensity values to an output range of intensity values, the range of values determined by the number of bits representing each intensity value, said digital image comprising pixels, each pixel having an intensity value of p-bits associated with it within the first range of intensity values, wherein the digital image is dithered from a p-bit representation to a q-bit representation, the value of p being greater than the value of q, said method comprising the steps of:
- means for dividing the input range of intensity values into sections representative of a low intensity section, mid intensity section and high intensity section;
 - means for dividing the output range of intensity values into sections representative of a low intensity section, mid intensity section and high intensity section to correspond respectively to the low intensity section, mid intensity section and high intensity section of the input intensity range, wherein the elements in each output intensity section correspond to a quantized intensity value in the corresponding input intensity section;
 - means for generating a dither matrix for each intensity section; the elements in each dither matrix equal to the values in the range between 0 and the interval spacing for the section, which are representative of the number of intermediate intensity values between quantized intensity values in each section, said elements organized to provide the desired visual perception of the intermediate intensity values;
 - means for retrieving a dither value from a dither matrix for each pixel value the dither matrix the dither value is retrieved from being dependent upon the input intensity value of the pixel;
 - means for determining the input-scale base value for a pixel value, said input-scale base value being the value less than or equal to the pixel value;



- 40 -

means for determining a quantized intensity value which corresponds to the input-scale base value;

if the difference between the input intensity value and the input-scale base value is less than or equal to the retrieved dither value,

5 means for setting the output intensity level for the pixel to be equal to the quantized intensity value; and

if the difference between the input intensity value and the input-scale base value is greater than the retrieved dither value, means for setting the output intensity level for the pixel to be equal to a quantized
10 intensity value which corresponds to the input-scale base value +1.

30. The apparatus of claim 29 wherein a greater number output intensity values are located in the low intensity section than in the high intensity section to provide a higher precision of dithering for input intensity values within
15 the low intensity section.



31. The method of claim 30 wherein the input range of intensity values are selected according to the function of the eyes' sensitivity to changes in light intensity, the means for dividing the input range of intensity values comprising:

5 means for providing a function of eye sensitivity to changes in intensity;

means for selecting two intermediate points along a curve representing the function of the eyes' sensitivity to changes in intensity;

10 means for determining the low intensity section to be the values along a straight line approximation between the beginning of the curve and the first intermediate point selected;

means for determining the mid intensity section to be the values along a straight line approximation between the first intermediate value and the second intermediate value selected; and

15 means for determining the high intensity section to be the values along a straight line approximation between the second intermediate value and the end of the curve.

32. The method of claim 29 wherein:

20 the digital image to be dithered is depth cued to a dark background;

the greatest concentration of occurrence of intensity values occurs within the low intensity section;

whereby the low intensity section is smaller than the mid and high intensity sections.

25



33. The apparatus of claim 32 wherein:

p equals 8 thereby providing an input intensity range of 256 values;

5 q equals 4 thereby providing an output intensity range of 16 values;

the low intensity section of input intensity values extends from 0 to 54;

the mid intensity section of input intensity values extends from 55 to 156; and

10 the high intensity section of input intensity values extends from 157 to 255;

34. The apparatus of claim 33 wherein:

15 the low section of output intensity values ranges from 0 to 6, inclusive;

the mid section of output intensity values ranges from 7 to 12, inclusive; and

the high section of output intensity values ranges from 13 to 15, inclusive;

20 the dither matrix corresponding to the low intensity section having elements in the range of 0 to 7;

the dither matrix corresponding to the mid intensity section having elements in the range of 0 to 15; and

25 the dither matrix corresponding to the high intensity section having elements in the range of 0 to 31.



- 43 -

35. The apparatus of claim 34 wherein the digital image is dithered using area dither and the dither matrices comprise the values:

0	7	1	6
4	2	5	3
1	6	0	7
5	3	4	2

5

for the matrix which corresponds to the low intensity section;

1	15	2	12
9	5	10	6
3	13	0	14
11	7	8	4

for the matrix which corresponds to the mid intensity section;

2	30	4	24
18	10	20	12
6	26	0	28
22	14	16	8

10

for the matrix which corresponds to the high intensity section.



- 44 -

36. The apparatus of claim 35 wherein the digital image is dithered using vector aligned dither and the dither matrices comprise the values:

| 0 4 2 7 1 5 3 6 0 4 2 7 1 5 3 6 |

5 for the matrix which corresponds to the low intensity section;

| 0 8 4 10 2 12 6 14 1 9 5 11 3 13 7 15 |

for the matrix which corresponds to the mid intensity section;

10 **| 0 16 8 20 4 24 12 28 2 18 10 22 6 25 14 30 |**

for the matrix which corresponds to the high intensity section.

37. A method for dithering a digital image from an input intensity range to an output intensity range substantially as hereinbefore described.

15

38. An apparatus for dithering a digital image from an input intensity range to an output intensity range substantially as hereinbefore described with reference to the accompanying drawings.
