9510910-3
IPOS copy

HALFTONE DITHER CELL WITH INTEGRATED PREFERRED COLOR MATCHING

Field of the Invention

This invention relates to digital color printers which produce binary level color dots to create a full-color representation of an image and, more particularly, to a digital color printer performing color digital halftoning using dither matrices with built-in preferred color matching.

Background of the Invention

Currently, both inkjet printers and laser printers are capable of producing full color images with high quality and precision. Such color printers are controlled by a printer driver program which provides an interface between an application program running on a host processor and the printer. Normally, a user creates a document using an application on the host computer and then calls for initiation of the printer driver program. In response to the user's instituting a print command, the host computer transmits a series of page descriptions to the printer driver. The printer driver then proceeds to use built-in functions to rasterize the page description into a pixel map of a predefined resolution (e. g. 300 dots per inch, 600 dots per inch, etc.). The printer driver must also adjust the printed colors to match the screen colors as closely as possible.

Personal computers (PCs) use eight-bit values to designate each primary color. To produce secondary colors, a PC uses combinations of the three eight-bit values to control the computer's display device (e. g. a color CRT). A twenty-four-

bit value can represent 2^{24} different color values, which color values can be reproduced by appropriate control of the CRTs color-associated electron guns. When a color printer is called upon to accurately reproduce that many color values, extensive color processing is required.

Thus, in order to reproduce a received color value, a color printer must convert the color value into a color command that is recognized by the printer engine. It was early realized that it was impractical to produce a color-table map that mapped all 2^{24} possible PC-generated input colors to printer-engine color codes.

Each pixel in the pixel map comprises, for instance, three eight-bit values corresponding to red, green and blue values derived from a display device in the host processor. The printer driver must adjust the color values in accordance with a predetermined calibration function so as to assure that the to-be-printed colors will appear the same as, or related in a desired manner to, the colors displayed on the display device.

This is performed by a color management or mapping system that assures the colors produced by one product (a printer, scanner, monitor, film recorder, etc.) match or are suitably related to those produced on others. Color management systems typically have two components, "profiles" of individual color products that specify the color capabilities of the device, and software that runs on a host computer that uses this information to ensure that the colors produced by one product are thus related to those produced by another. In cases where a particular color is not within the color gamut of a target device (i. e. the target device simply is incapable of reproducing the color), the color management software must provide the closest possible or otherwise-related match. Device-independent color is a term describing a computer system capable of reproducing a color accurately on any attached color device (printer, monitor, scanner, etc.). Device-independent color is usually implemented by developing "device profiles" that describe the colors a product can produce and by developing a color-matching engine that uses the profiles to convert color data to assure a match between devices.

RGB is a color space that uses as its primary colors red, green, and blue. These three colors are the primary "additive" colors. In devices that use projected light to produce an image (for example, televisions or computer monitors), a broad gamut of colors can be reproduced using red, green, and blue. Red and green combine to form yellow, red and blue to form magenta, green and blue to form cyan, and all three to form white. Other shades can be produced by combining different amounts of the three primary colors.

CMYK is a color space that uses as its primary colors cyan, magenta, yellow and black. These four colors are the primary "subtractive" colors; that is, when printed on paper, the CMYK colors subtract some colors while reflecting others. Cyan and magenta combine to form blue, cyan and yellow to form green, magenta and yellow to form red, and in theory, all three to form black; however, it is sometimes difficult to get a satisfying black using a given set of cyan, magenta, and yellow pigments, and therefore many reflective color-based products add a "true" black color, hence CMYK, not CMY. (To avoid confusion with blue, the letter K is used to represent black). The CMYK color set is sometimes called "process color."

In printing, the printer uses the three subtractive primary colors. In each, one of the three additive colors has been subtracted from the white light. When red is subtracted, green and blue are left which combine to form the color cyan. When green is subtracted, red and blue are left which combine to form the color magenta. When blue is subtracted, red and green light combine to form the color yellow. The printer's subtractive primary colors are cyan, magenta, and yellow. In principle the overprinting of all three in solid images yields black. The combination is black because each has subtracted one of the three additive primary components of white light and the complete absence of light is black.

Thus, the printer driver must convert the red, green and blue values to cyan (C), magenta (M), and yellow (Y) values. As a result, each pixel is then represented by three eight-bit values which identify the corresponding levels of C,M,Y that will be used to subsequently control the print mechanism. An additional eight-bit value is supplied for a pixel black (K) dot to be applied at the pixel location.

Color printers can print one of eight colors at a particular pixel (red, green, blue, cyan, magenta, yellow, black, or white). However, the computer can request any one of sixteen million colors. Therefore, it is necessary to produce a translation between twenty-four-bit pixels (sixteen million colors) and three-bit pixels (eight colors). This translation is called digital halftoning. It is an integral part of color printing.

Halftoning refers to any process that creates the illusion of continuous-tone images by judicious arrangement of binary picture elements, such as ink drops in the case of inkjet printers. Thus, halftoning is printing the simulation of a continuous-tone image, such as a shaded drawing or a photograph, with groups or cells of color or black dots. The dots are placed in such a way that they appear to the human eye to be a single color. Digital halftoning is sometimes called spatial dithering.

Printing presses and most printers use halftoning to render images. On printing presses, different-size dots can be used to produce different shades of gray or color. Most color printers are binary in nature, in that they either apply a full color dot or no color dot to a pixel location. Such color printers do not employ a control mechanism to enable adjustment of the intensity of a particularly applied color dot. In binary printers, different patterns of identical dots are used to produce halftone images. As a result, a printer driver for a binary color printer employs a color digital halftoning process which reduces the twenty-four-bit color information to three bits per pixel print position (one bit for each of the C,Y, and M color planes).

Dithering can be used to reproduce gray shades using only black ink, or a fuller gamut of color using only the process colors (cyan, magenta, yellow, black). For example, to produce green, a color printer lays down patterns of small yellow and cyan dots that appear to the eye to be green. There are many halftoning techniques, each with its own method for laying down dots. Examples include pattern dithering and error diffusion. Pattern dithering uses a library of set patterns to reproduce a color (in color printing) or a gray shade (in monochrome printing).

Pattern dithering can be characterized as ordered or random. Ordered dithers generally fall into one of two broad classes, dispersed, or clustered.

In dithering, a lot of work has been done to create the ideal "dither cell". This effort has been put into developing dither cells that have random or "blue noise" characteristics. Such "supersmooth" dither cells produce an image that appears almost as good as error diffused, but with the speed performance of a dither.

Dithers are implemented by use of a dither cell or dither matrix or threshold array, also called a mask, a two-dimensional matrix of thresholds. Pixel values are compared to corresponding entries in the dither cell to determine whether they should be turned on or off. In this way a shade of red for example can be converted to full red or no red. Many different approaches exist that vary the size of the cell and the distribution of the thresholds. Thus halftoning is accomplished by a simple pointwise comparison of the input image against a predetermined threshold array or mask. For every point or pixel in the input image, depending on which point value is larger of the image or the mask, either a one or zero, respectively, is placed at the corresponding location in the binary output image.

This matrix is mapped over the continuous-tone image. For an image that is larger than the space occupied by the dither matrix, the matrix replicates itself to cover the entire image.

The elements of the dither matrix are termed threshold values. The distribution of threshold values in a dither matrix determines the tonal response of the dither matrix. A goal of halftone printer design is accurate and pleasing reproduction of continuous-tone images. To achieve these goals, the dither matrix must be calibrated, i. e., the threshold values which makes up the dither matrix must be adjusted. One goal of calibration is to ensure that when continuous-tone images of particular darkness are printed on a halftone printer they have the same darkness as the original images, respectively. Analogous goals exist for other output devices, such as monitors.

Pattern dithering in general benefits from ease of implementation. Pattern dithering is computationally fast but does not offer the best possible reproduction

quality. Error diffusion is a technique for laying down dots of the three process colors to produce the full spectrum of color. Error-diffusion techniques use complex algorithms to lay down dots of color in a random rather than a repeated pattern, which improves the quality of the image. Error diffusion makes the best approximation it can for a given pixel, calculates how far that approximation is from the ideal and propagates this "error" to neighboring pixels. In this way a given pixel may not be particularly accurate, but the area is. In general, error diffusion generates much better print quality than dithering; however, typically, intense calculation is required to create the random pattern, and therefore printing images using error diffusion is much slower than using pattern dithering.

Thus, halftoning algorithms can generally be evaluated in terms of speed of execution and resulting print quality. Often a tradeoff needs to be made between an algorithm that is fast, but does not produce optimum print quality versus an alternative approach with better print quality that takes longer. So the problem each halftoning technique is trying to address is how to quickly produce a high print-quality image.

Traditional methods of color matching involve 3x3 matrix transformations of colors from the device-specific color space of the input source (e. g., a computer monitor or color scanner) into a device-independent color space (e. g. CIE XYZ). The device-independent color is then transformed into the device-specific color space of the output device (e. g., a printer or computer monitor). Normally in the case of a color printer, a three-dimensional color correction is done with a large color-lookup table instead of 3x3 matrix transformations. The matrix transformations alone do not produce an adequate match to the printer. This is because color reproduction is fundamentally a subtractive color process on a hardcopy device and ink-media interactions are very nonlinear in comparison to the additive color process of a scanner or monitor.

The size of a complete three-dimensional (256 values per dimension) color-lookup table, however, is prohibitive. Therefore, the color space is subsampled and a smaller color-lookup table and interpolation are used to produce the desired results. This color-matching procedure achieves sufficient accuracy, but is very

computation-intensive. For example, when using PostScript™ some printers take fifty percent longer to process and print an 8x10 color image when using color matching versus when color matching is not used. Moreover, customer surveys regarding color matching have led to the discovery that most users are more interested in attractive, vivid, color output than they are in precise colorimetric color matching.

Prior solutions for accurate color matching as described above have been used in existing color printers, but their big disadvantages are poor performance and a less-than-optimal color match. Thus, there is a continuing need to improve the clarity and vividness of color images produced by binary color printers without decreasing speed.

It is therefore advantageous to provide a system and a method for generating dither matrices with vivid color matching embedded during the dither-matrix generation process.

It is further advantageous to have a method and a system for generating dither matrices with embedded vivid color matching without reducing the number of available tone levels.

It is also advantageous to have a dither matrix with embedded vivid color-matching system and method that permits an output device to map the number of distinguishable tone levels in the dynamic range of the source into the same number of levels in its own dynamic range.

Furthermore, it is advantageous to provide a method and a system for enabling fine tuning of the dither matrix to a specific halftone imaging device.

Summary of the Invention

The present invention provides an improved method of color matching by utilizing a vivid-response function that produces clear and vivid output. Moreover, there is no decrease in throughput because the color transformations are embedded

in the dither cell. Also preferably embedded in the dither cell is a second response function that generally lightens the color values established by the first function.

The present invention further provides a system and a method for generating a dither matrix calibrated to that desired vivid color response. The present invention provides a halftoning technique that produces high-quality vivid color output with no decrease in throughput because the color transformations are embedded in the dither cell.

Other very important aspects and advantages of the present invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention, in which:

Brief Description of the Drawings

Fig. 1 is a block diagram of a computer system for generating calibrated dither matrices according to the present invention;

Fig. 2 shows a preferred dither matrix used by the computer system of Fig. 1;

Fig. 3 shows a vivid-response curve in accordance with this invention;

Fig. 4 shows how a vivid-response curve $\underline{y}(\underline{x},\underline{s},\gamma_1,\underline{h})$ is generated using a first gamma function;

Fig. 5 shows a second gamma function $\underline{z}(\underline{x},\gamma_2)$,

Fig. 6 shows a composite function $\underline{z}(\underline{y},\gamma_2)$ of the vivid-response curve and the second gamma function;

Fig.7 shows a histogram of the dither of the present invention;

Fig. 8 shows the preferred steps to generate a dither matrix of the present invention;

Fig. 9 shows an example of an intermediate pattern of the present invention;

Fig. 10 describes the step of generating the intermediate pattern in greater detail;

Fig. 11 describes the steps of generating patterns of the dither matrix of the present invention with fewer ones than the intermediate pattern;

Fig. 12 describes the steps of generating patterns of the dither matrix of the present invention with fewer zeros than the intermediate pattern;

Fig. 13 shows a circle as a model of an individual symbol (mark) printed by the present invention (particularly appropriate when the invention incorporates an inkjet printer that produces essentially discrete inkdrops), and identifies certain elemental overlap areas of such an inkdrop symbol relative to a pixel grid; and

Fig. 14 shows more-complex overlap and nonoverlap areas for such a circle expressed in terms of the Fig. 13 elemental areas.

Detailed Description of the Preferred Embodiment

Color printers typically can print one of eight colors at a particular pixel (red, green, blue, cyan, magenta, yellow, black, or white). However, the computer can request any one of sixteen million colors. Shades of 0-225 for Red, Green, and Blue produce sixteen million combinations. Therefore, it is necessary to produce a translation between twenty-four-bit pixels (sixteen million colors) and three-bit pixels (eight colors). As mentioned above, this translation is called halftoning.

Dither cells are customarily placed initially in a position where the upper left corner of the cell corresponds to the upper left corner of the image being rendered. This placement causes each dither cell location to now correspond to specific pixels in the image area covered by the dither cell. The color printed by the halftoning process is determined by the comparison of the color value of the pixel, and the value of the dither cell that corresponds to that pixel.

Rendering of that area of the image corresponding to the area covered by the dither cells is performed. The dither cell is now repositioned over the image by shifting the cell to the right by the distance which corresponds to the width of the cell. This new placement of the cell covers a new image area immediately adjacent to the area of the image previously covered. This part of the image may then be

rendered. The dither cell is moved again and again until the right hand side of the image is reached. The process continues by returning the dither cell to the left side of the image and shifting it downward by the amount which corresponds to the height of the cell. The process is then repeated until the entire area of the image has been covered and rendered.

Fig. 1 is a block diagram of a computer system 100 for generating calibrated dither matrices according to the present invention. A Central Processing Unit (CPU) 101 is connected to a computer memory 103. The computer memory 103 contains a control program 105, a dither matrix generation program 107, and a cumulative histogram generator 109. The CPU 101 may also be attached to a plurality of secondary storage devices (not shown). The programs 105, 107 and 109 may be stored on these secondary storage devices and loaded into the computer memory prior to the execution of these programs.

The control program 105 is logically connected to the dither matrix generation program 107 and cumulative histogram generator 109.

The CPU 101 is further connected to a halftone output device 111. The halftone output device 111 is operable to output continuous-tone images using a dither matrix generated by the dither matrix generator 107. The halftone output device 111 may be either a black-and-white or a color halftone output device.

In operation, the computer system 100 converts a gray-scale image 102, stored, for example, in the computer memory 103, into a halftone image 104. The halftone image is represented by symbols output on the halftone output device 111. Fig. 2 shows a preferred dither matrix 110 used by the computer system 100 to generate the halftone image 104. The halftone image 104, the continuous-tone image 102, and the dither matrix 110, each occupy an area; for example, the dither matrix occupies an area 112. The three areas are substantially equal to each other. The size of the dither matrix is small, and depends on the resolution of the output device. A dither matrix is a dimensionless array of integers. The size of the dither matrix is defined in terms of the numbers of rows and columns it contains. So long as the numbers of rows and columns of the dither matrix are smaller than the numbers of rows and columns in the image, the matrix will be replicated.

Both the halftone image 104 and the continuous-tone image 102 have many pixels such as the continuous-tone image pixel 114 and the halftone image pixel 116. Each pixel has a value; for example, the continuous-tone image pixel 114 has a value of 200 and the halftone image pixel 116 has a value of 0. The dither matrix 110 has many elements; for example the element 118. In one preferred embodiment, the matrix has 128 rows by 128 columns of elements. Each element has a value and occupies a position in the dither matrix area; for example, the element 118 has a value 188, and occupies a position 122.

Both the continuous-tone and the halftone image may be in black-and-white or color. One preferred way to represent color is to have three components or symbols for each pixel, each component or symbol having a different color.

The method of generating the halftone image includes the steps of comparing the value of each pixel of the continuous-tone image with the value of an element in the dither matrix. Based on the result of the comparison, the value of a corresponding pixel of the halftone image is determined. For example, the value 200 of the continuous-tone image pixel 114 is compared to the value 188 of the element 118. Based on the comparison, the value of the pixel 116 is determined to be a minimum value, such as 0, which means that a symbol will not be output at that pixel.

The present invention has two key components: (1) creation and use of a vivid-response curve that produces clear and vivid output, and (2) embedding of the color transformations in the dither cell. Of course it also encompasses printing with the completed cell, as indicated elsewhere in this document.

As shown in Fig. 3, the vivid-response curve has an S shape with a plateau both at the top and the bottom. The vivid-response curve is applied independently to the red, green, and blue channels or dimensions of the input image data. While this vivid-response curve correction is applied to the three color channels, red, green, and blue, it is considered a "one-dimensional color correction," because it is applied to each color dimension independently. This is in contrast to the conventional "three-dimensional" approach which considers all three color dimensions simultaneously when color matching in each dimension.

Applying the vivid-response curve of Fig.3 stretches the midtones of an image while compressing the highlight and shadow detail. This causes the colors to become more vivid. For example, consider an off-yellow which would have high values for red and green, and a low value for blue. Applying the vivid-response curve, the red and green values would be increased, while the blue value would be reduced. This results in more vividness and more contrast and as a result a sharper appearance. The plateau at the top of the vivid-response curve also causes "snapping to primaries," i. e., values close to 255 (on a 0-255 intensity scale) become 255. Thus, slightly-off white becomes true white and color triplets close to pure primary or secondary colors are pushed closer or all the way to those primary and secondary colors.

This vivid-response curve is produced by obtaining the input range between the plateaus and cutting it in half. The first half of the range is a first gamma function, $\underline{f}(\underline{x}) = \underline{x}^{\gamma_1}$ as shown in Fig. 4B. The second half of the range uses the same gamma function flipped over upward, and flipped over to the right. As a result, the three parameters used to specify a particular vivid-response curve are the lower plateau limit $\underline{s}$, the first gamma value γ_1 that specifies the degree of "S" shape, and the upper plateau limit $\underline{h}$.

The vivid-response function $\underline{y}(\underline{x}, \underline{s}, \gamma_1, \underline{h})$ where $\underline{F} = 255$, $\underline{H} = 127.5$, and $\underline{m} = \frac{1}{2}(\underline{h} - \underline{s})$, is as follows.

$$\underline{y} = \underline{H} \cdot \left(\frac{\underline{x} - \underline{s}}{\underline{m}} \right)^{\gamma_1} \qquad \text{if } \underline{x} < \underline{s} + \underline{m}, \text{ and}$$

$$\underline{y} = \underline{F} - \underline{H} \cdot \left(\frac{\underline{h} - \underline{x}}{\underline{m}} \right)^{\gamma_1} \qquad \text{otherwise.}$$

The second gamma function $\underline{z}(\underline{y}, \gamma_2)$—which is to say, with the result $\underline{y}$ of the first function inserted in lieu of $\underline{x}$ as an argument of the second function to form the "composite curve" in Fig. 6—is as follows.

$$\underline{z} = \underline{F} \cdot \left(\frac{\underline{y}}{\underline{F}} \right)^{\gamma_2}$$

For example, values for ($\underline{s}$,γ_1, $\underline{h}$) of (0, 1.6, 245) work well on the Hewlett-Packard DeskJet 1200C and 1600C color inkjet printers, making $\underline{m}$ = 245/2 = 122½.

For a printer with well-behaved inks like the Hewlett-Packard DeskJet 1200C and DeskJet 1600C color inkjet printers, this one-dimensional color correction produces excellent printed color output when combined with a second gamma function (also a one-dimensional color correction) that lightens the output. As shown in Fig. 5, this second gamma function—graphed here as a noncomposite function $\underline{z}(\underline{x},\gamma_2)$ of $\underline{x}$ rather than $\underline{y}$—pushes all values up (with $\gamma_2 < 1.0$), thus lightening the image. The curves of Figs. 5 and 6 were developed, as can be verified from the data as graphed, using a second gamma value $\gamma_2 = 0.6$. Combining the vivid-response curve with the second gamma curve produces an attractive image that is actually preferred over an accurate color match. Furthermore, computationally, this processing is much faster than using a traditional three-dimensional color-lookup table.

A second aspect of the present invention involves incorporating these vivid-response curves into the dither cell. U .S. patent application serial 08/308,321, entitled SYSTEM AND METHOD FOR GENERATING CALIBRATED DITHER MATRICES, filed September 19, 1994, discloses the use of a measured tonal-response curve in producing a histogram for a blue-noise dither cell that results in visually linear output. This is accomplished by using the measured tonal response to determine how many incremental dots to turn on at each level in the dither mask. This is in lieu of turning on a constant number of dots at each level. The motivation for this is to aid the color matching by making the output device more linear.

The present invention incorporates into the dither cell a composite function which is a concatenation of the vivid-response gamma function and a second gamma function. This composite function is shown in Fig. 6. The word "concatenation" here means that preferably the result of the first (vivid-response) gamma function is used as the argument of the second.

By employing this composite function it is possible to fabricate a dither (blue-noise or conventional dither), that incorporates all of the color correction

within it. This composite function is used to determine how many dots to turn on at each level. This dither differs from a normal dither in that its histogram (the number of dots turned on at each level) reflects the composite function rather than a uniform number of dots. Fig. 7 shows a typical histogram for the dither matrix of the present invention.

Results indicate that for an inkjet printer with well-balanced cyan, magenta, and yellow inks, halftoning with the new matrix alone can produce vivid color images without the need of doing additional color correction. With the resultant "vivid dither" matrix, all color correction is contained within the halftone dither cell. Since all of the additional processing is off-line in the construction of the dither cell, there is no increase in processing time over unmatched color output. This can significantly reduce the amount of processing time involved in color-table interpolation.

Suppose the input optimization curve is $\underline{f}(\underline{x})$, and the histogram of the original matrix is $\underline{g}(\underline{x})$, then the new matrix threshold can be obtained from the mapping:

$$\underline{x}_n = \underline{f}^{-1}(\underline{x}).$$

Therefore

$$\underline{x} = \underline{f}(\underline{x}_n),$$

and the cumulative histogram of the original matrix is—

$$\underline{c}(\underline{x}) = \underline{g}(\underline{x})\, d\underline{x}.$$

After the mapping, the cumulative histogram of the new matrix is—

$$\underline{c}_n(\underline{x}) = \underline{g}(\underline{f}(\underline{x}_n))\, d(\underline{f}(\underline{x}_n)),$$

and with some simple mathematics one obtains the new cumulative histogram as

$$\underline{c}_n(\underline{x}) = \underline{g}(\underline{f}(\underline{x}))\, \underline{f}'(\underline{x})\, d\underline{x}.$$

If the histogram of the original matrix is uniform, i. e. $\underline{g}(\underline{x}) = \underline{K}$, a constant, then the new cumulative histogram is $\underline{c}_n(\underline{x}) = \underline{K}\, \underline{f}(\underline{x})$. Hence, the new cumulative histogram is proportional to the optimization curve.

The above analysis clarifies the relationship between the mapping and the histogram of the new matrix. In practical implementation, one can start with generating a matrix that has a uniform histogram. One can then obtain the optimization curve either through measurement or through image analysis. After that, one can

derive the new histogram from the optimization curve, and generate a dither matrix with that cumulative histogram.

In case one needs to do further adjustment of the tonal response, one can repeat the process to obtain the optimization curve; however, the matrix no longer has a uniform histogram. One can perform the threshold mapping first to obtain a new threshold matrix. The new matrix will have fewer gray levels since many levels are lost in the mapping process due to quantization errors. To make up for the lost levels, one can then calculate the cumulative histogram of the mapped matrix, smooth out the discrete steps, and generate a new strictly monotonic cumulative histogram. This new histogram is input to the matrix generation program to generate the final matrix.

The calibrated dither matrix is generated by inputting the cumulative histogram to the dither-matrix generation program 107 to generate a dither matrix based on the cumulative histogram.

Fig. 8 shows the steps 200 to generate the dither matrix 110. The first step, 202, is to generate an intermediate pattern with elements (generally speaking, pixel positions) having associated values of either one or zero. The elements with ones are substantially uniformly distributed within the pattern.

As will be understood, however, the effective spatial coverage of each element that is assigned a value of one is not actually exactly one, because a unit mark (or "symbol") made by a printing machine typically does not exactly match the shape of a pixel, and in any event does not cover exactly the area of a single pixel. The actual amount of area or space covered by a nominally one-pixel mark depends upon the particular printing mechanism—for example, inkjet printing or laser-imaged dry-toner printing—that is employed to make the mark.

This is true even when tonal shades are relatively light, so that marks ("symbols") are scattered rather sparsely. These effects become even far more important, however, for darker shadings in which it is more frequent to find marks assigned to adjacent pixels—inasmuch as overlapping of adjacent marks can very strongly reduce the incremental area coverage for each added mark. Rather than

the nominal one-unit added pixel, addressing one incremental pixel may only add a rather small fraction of a unit added pixel, thus severely degrading tonal linearity.. These observations are considered to be well known in the art—having been published for example in U. S. Patent 5,469,515.

It shows that the amount of areal coverage accurately attributed to each pixel can be calculated from knowledge of a specific mask considered together with a spatial model for the particular marking technology that is in use. Thus the values of the elements are dependent on the model that is used for the symbols.

That patent, entitled "HALFTONE IMAGE FORMATION USING DITHER MATRIX GENERATED BASED UPON PRINTED SYMBOL MODELS", suggests for instance that inkjet technology produces generally circular inkdrops. In such a modelling each position of an element (pixel) is represented as a square with a dimension T, 350 (Fig. 13), and each circular mark is taken as having a definite radius r, 352.

The area of overlap of a circle to its nondiagonal neighbor is c, 354. The area of overlap of such a circle to its diagonal neighbor is a, 358. The area of intersection of two circles in a common diagonal neighbor is b, 356. Defining $m = r/T$, the values for a, b and c are:

$$a = \frac{\pi m^2}{4} - m^2 \sin^{-1}\left(\frac{1}{2m}\right) - \frac{\sqrt{4m^2-1}}{2} + \frac{1}{4}$$

$$b = m^2 \sin^{-1}\sqrt{\frac{2m^2-1}{2m^2}} - \frac{\sqrt{2m^2-1}}{2} - a$$

$$c = \frac{\sqrt{4m^2-1}}{4} + m^2 \sin^{-1}\left(\frac{1}{2m}\right) - \frac{1}{2}$$

Fig. 14 shows the values of different elements when more of them overlap. For instance, the value of the element in the position 362 is (3c - 2b). Taking into

account the actual positions of the nominally unit-valued "symbols" in a specified matrix pattern, all the values of the elements are to be modified by the above model, or another suitable model if a different kind of printing mechanism is in use. For example U. S. Patent 5,469,515 also shows how to make analogous calculations where the marking mechanism is a laser-imaged dry-toner system that produces an approximately Gaussian probability distribution of toner particles, rather than a relatively well defined circular inkdrop.

Fig. 9 shows an example of an intermediate pattern 300, with 128 rows by 128 columns of elements. The pattern shown in Fig. 9 was printed by a 600 dpi printer and enhanced by magnifying it nine times, with the image duplicated four times, once along the horizontal direction, and then along the vertical direction. A symbol output implies that the element at that position has a value of one, and a void implies that the element at that position has a value of zero.

In the second step, 204, patterns of the dither matrix 110 with fewer elements having ones than the intermediate pattern 300 are generated. This is done through modifying the values of the elements by the model of the symbols and through replacing a plurality of ones with zeros from the intermediate pattern 300. The ones to be replaced are in regions with elements having ones clustered together as identified by a filter.

The difference in the number of elements having ones, from one pattern to its next pattern in succession, is dependent on a "quantization number", also now well known in the art through introduction in the above-mentioned U. S. Patent 5,469,515. As there explained, the quantization number expresses the condition that inking levels should be spaced apart, in perceptual space, in some rational manner.

This may be accomplished for example by holding substantially constant, as between successive patterns, either the incremental number or the incremental fraction of ones or zeros. The former alternative of course corresponds to constant steps on a linear scale, whereas the latter alternative may be preferable as it

amounts to constant steps on a logarithmic scale—which may more closely track human capability for perceptual discrimination between levels.

The linear case is readily understood: since a representative dither matrix 110 has 128 rows and 128 columns of elements, the total number of elements in the matrix is 128^2. The gray-scale image 102 is taken as having 256 levels of lightness, including the two end levels (one with no symbols and the other entirely covered with symbols). The levels of lightness determine the number of patterns for the dither matrix 110; thus there are 255 steps to take, from one end level to the other, and there are 128^2 elements to be covered progressively by these 255 steps, so that each step should add (or subtract, depending on the direction one has in mind) $128^2/255 \approx 64$ elements per step. It is this value, for the linear case, which is denominated the quantization number.

For the logarithmic case, which is taken as corresponding generally to perceived changes in tone reflectance between successive patterns, again there are substantially 255 transitions, but now to be taken in steps that are equal multiplicatively. The above-mentioned patent suggests a constant step of 1/255 (to be applied, of course, in the logarithmic domain, to the ratio between maximum and minimum values).

It also shows how to model the above-mentioned printing mechanism of laser printers, in which toner particles follow a probability density function that is a Gaussian curve through a pixel center. In addition it explores various other models including an ad hoc analysis for arbitrary symbol shapes or density functions, in which the technique is to digitize an actual symbol and to store the digitized values as a bit map in the computer system.

In the third step, 206, patterns of the dither matrix 110 with fewer zeros than the intermediate pattern 300 are generated. This is done through modifying the values of the elements, once again by the model of the symbols, and through replacing a plurality of zeros with ones from the intermediate pattern 300. The zeros to be replaced are in regions with elements having zeros clustered together as

identified by the filter. The difference in the number of elements having zeros from one pattern to its next pattern is dependent on the quantization number.

Finally, in step four, 208, the dither matrix, 110, is formed by adding all the patterns to the intermediate pattern.

Steps two to four are not restricted to that sequence. Step three, 206, can be done before step two 204. Step four, 208, the summing step, can be performed as the patterns are generated. For example, after the formation of the intermediate pattern, the pattern is copied to the dither matrix. Then, as each additional pattern is generated, it is added to the dither matrix by matrix addition. So, when all the patterns are generated, the dither matrix 110 is also formed.

In the present example, the dither matrix, 110, has 128 rows by 128 columns of elements, and the gray-scale image, 102, has 256 levels of lightness, including the two end levels, the level with no symbols and the level entirely covered with symbols. The levels of lightness determine the number of patterns for the dither matrix, 110. In one preferred embodiment, the total number of patterns including the intermediate pattern is 256. The difference in the number of elements having values equal to one between a pattern and its next pattern is dependent on the quantization number. For the purpose of generating a dither matrix which conforms to a cumulative histogram, the quantization number depends on the histogram and is defined as the increment in value of the histogram between successive tone level values.

Fig. 10 describes the step, 202, of generating the intermediate pattern 300, in more detail. The intermediate pattern is set, 425, to be a random pattern, with ones and zeros randomly distributed. Preferably, a fraction of at least 100/255 of the elements in the random pattern has values equal to one; and at least the same fraction 100/255 has values equal to zero. Then, the intermediate pattern 300 is copied to a dummy pattern, 426, and the values of the elements in the dummy pattern are modified, 427, by the model of the symbols. Another preferred method for modifying the values of the elements is to find the probability of having toner particles as a function of position. Some printers, such as laser printers, uses toner particles to generate a symbol. In one embodiment, the optical beam to generate

the toner particles is Gaussian in shape, and accordingly one preferred probability of having toner particles follows a Gaussian shape.

Referring back to Fig. 10, after the values of the elements in the dummy pattern have been modified 427, the next step is to pass the dummy pattern through the filter. One preferred filter is described in U.S. Patent No. 5,317,418, which is incorporated herein by reference. The dummy pattern is duplicated two dimensionally before it is filtered. This is known as circular convolution. For the modified dummy pattern, with subelements and bit maps, the filter has to sample more frequently, such as at every subelement or at every bit. The filtering process will not be described in this application. A general discussion on this type of spatial filtering can be found in Discrete Time Signal Processing, written by A. V. Opponheim and R. W. Schafer, Prentice Hall, 1989.

The filtered output would have the largest value at the location where the ones are most clustered together, and the smallest value at the location where there is the largest void. The position with the largest value is identified, 430, as the maximum position, and the position with the smallest value is identified as the minimum position, 356. If there are locations that are equally clustered or voided, then there will be more than one maximum or minimum position. In one preferred embodiment, the first position found to be the maximum and having a value of one in its corresponding position in the intermediate pattern is selected to be the maximum position. Similarly, the first position found to be the minimum and having the value of zero in its corresponding position in the intermediate pattern is selected to be the minimum position.

After the maximum and minimum position are identified, 430, their elements in the intermediate pattern, 300, are exchanged, 432.

The intermediate pattern, 300, is then again copied to the dummy pattern to be modified by the model of the symbols and filtered. These steps are repeated numerous times, 434, until the intermediate pattern reaches an equilibrium state, such as the example shown in Fig. 9. At that point, the ones and the zeros are substantially uniformly distributed within the pattern.

The dither matrix requires one or more patterns with elements having fewer ones. Fig. 11 describes the step, 204, of generating patterns of the dither matrix 110 with fewer ones than the intermediate pattern in more detail.

In generating the many patterns with fewer ones, first, the intermediate pattern, 300, is copied, 521, to a temporary-intermediate pattern, and is also copied, 522, into a dummy pattern. Then, the values of the elements in the dummy pattern are modified, 524, by the model of the symbols as described above, and a start sum is generated, 526.

In one embodiment, the start sum is equal to the sum of the values of the modified pattern. The way to calculate the tone reflectance from the coverage of the symbols on the pattern by the Yule-Nielsen equation is as described above. The dummy pattern is then filtered to identify, 527, the position of the dummy pattern where ones are substantially most clustered as the maximum position. The identification step is similar to the methods described above by the normalized Gaussian filter with a sigma of 1.5 and will not be further described here. The value of the element in the maximum position of the temporary-intermediate pattern is replaced, 528, with a zero.

Then, the values of the temporary-intermediate pattern are copied, 529, to the dummy pattern. The dummy pattern is again modified by the model of the symbols, 530, and an end sum is generated, 532. The end sum is generated the same way as the start sum.

The invented method then goes back, 534, to the "identify" step 527 until the difference between the start sum and end sum is less than or equal to the quantization number to generate one pattern of the dither matrix. The start sum then is set, 535, to be the end sum, and the generated pattern is stored in the computer 100. The quantization number is determined from the cumulative histogram.

The method then again goes back, 536, to the "identify" step 527, until the end sum is less than or equal to the quantization number to generate numerous patterns of the dither matrix 110 with fewer ones. The method will stop when the pattern above the most sparse pattern is generated because the most sparse pattern is the pattern with no symbols.

The dither matrix 110 also requires one or more patterns with fewer zeros. Fig. 12 describes the step, 206, of generating patterns of the dither matrix 110 with fewer zeros than the intermediate pattern in more detail. The steps shown in Fig. 12 are similar to the steps shown in Fig. 11. First, the intermediate pattern is copied, 551, into a temporary-intermediate pattern, and is also copied, 552, into a dummy pattern. Then, the values of the elements in the dummy pattern are modified, 554, by the model of the symbols. After the modification, a start sum is generated, 556. The element in the position of the dummy pattern where elements with zeros are substantially most clustered together is identified, 557, to be the minimum position. The element in the minimum position of the temporary-intermediate pattern is replaced, 558, with a one. After that, the temporary-intermediate pattern is copied, 559, to the dummy pattern, and the values of the elements in the dummy pattern are modified, 560, by the model of the symbols and an end sum is generated, 562. The above steps are repeated, 564, from the "identify" step, 557, until the difference between the start sum and end sum is less than or equal to the quantization number to generate one pattern of the dither matrix. At that point, the start sum is set, 565, to be the end sum, and the generated pattern is stored in the computer 100. Finally, the above steps are repeated from the "identify" step, 557, until the end sum is greater than or equal to:

$$\left(\frac{\text{the total number of}}{\text{patterns of the dither matrix}} - 1 \right) \cdot (\text{the quantization number}),$$

to generate numerous patterns for the dither matrix. Note that the total number of patterns is 256 in the present example. The method will stop when the pattern below the most dense pattern is generated because the most dense pattern is the pattern totally filled with symbols.

It should be understood that the foregoing description is only illustrative of the invention. Various alternatives and modifications can be devised by those skilled in the art without departing from the invention. In addition, while the invention has been described with respect to a color inkjet printer, it is equally

applicable to other color output apparatus which is required to perform a color transformation. Accordingly, the present invention is intended to embrace all such alternatives, modifications and variances which fall within the scope of the appended claims.

WHAT IS CLAIMED IS:

1. A method of reproducing a multilevel plural-color input image (102); said method comprising the steps of:

applying to the image (102) a dither matrix (110) that incorporates a vivid response function (y, 11-12; 15-16), to obtain, for each of plural output colors, a bit plane (104) having fewer levels than the input image; and

printing (111) the image using the obtained bit planes.

2. The method of claim 1, wherein:

the vivid response function (y) comprises a generally S-shaped response curve (11-12).

3. The method of claim 2, wherein:

the vivid response function (y) comprises two sections (11, 12) blended together at a generally central point;

each of the two sections is defined by a respective function (y: 11, and Fig. 4B; 12, and right end of Fig. 4C) having a respective argument that is proportional to an input color signal (x); and

each said function comprises raising its respective argument to a power γ_1 greater than unity and preferably about 1.6.

4. The method of claim 3, wherein::

each said argument is a fraction substantially equal to a range-normalized version (respectively [x - s] /m and [h - x] /m) of the input color signal (x).

5. The method of claim 4, wherein:

the response further incorporates a second function (z[x], Fig. 5) that in general lightens the response of the first-mentioned function (y).

6. The method of claim 1, wherein:

the response further incorporates a second function (z[x], Fig. 5) that in general lightens the response of the first-mentioned function (y);

said second function preferably comprises raising an argument of the second function to a power smaller than unity, ideally about 0.6, said argument of the second function preferably being proportional to the first-mentioned function.

7. The method of claim 1, further comprising the step of:

before the applying step, developing (109) a dither matrix (110, Fig. 9) that incorporates said vivid response function;

then measuring actual response;

then manipulating (Figs. 8 and 10-12) a cumulative histogram (Fig. 7) for said dither matrix in response to said actual response; and

then inputting said cumulative histogram to a dither generation program (105, 107) to generate a dither matrix that conforms to said cumulative histogram;

wherein said applying step employs said generated matrix.

8. The method of claim 1, wherein the vivid response function is a composite response function that includes:

a first response function having two separate sections (11, 12) blended together at a general middle point;

said separate sections of the response function having generally opposite trends, respectively; and

a second response function (14) having a general trend that opposes the trend of one (11) of said two separate sections and reenforces the trend of the other (12) of said two separate sections.

9. The method of claim 8, wherein, when presented as curves on a linear graph, the trend of:

a first (11) of the two sections of the first response function appears upward concave;

a second (12) of the two sections of the first response function appears upward convex; and

the second response function (14) appears upward convex.

10. The method of claim 1, wherein:

the first response function y, given that F and H (see marked in Fig. 6) are the full-scale and half-scale values of the output range, h and s (see marked in Fig. 6) are respectively the highest and lowest input-range values for which the function is defined, and the input midrange value is m (see marked in Fig. 6) = ½(h - s), is:

$$y = H \cdot \left(\frac{x-s}{m} \right)^{\gamma_1} \qquad \text{if } x < s+m, \text{ and}$$

$$y = F - H \cdot \left(\frac{h-x}{m} \right)^{\gamma_1} \qquad \text{otherwise; and}$$

the second response function z is:

$$z = F \cdot \left(\frac{y}{F} \right)^{\gamma_2}.$$

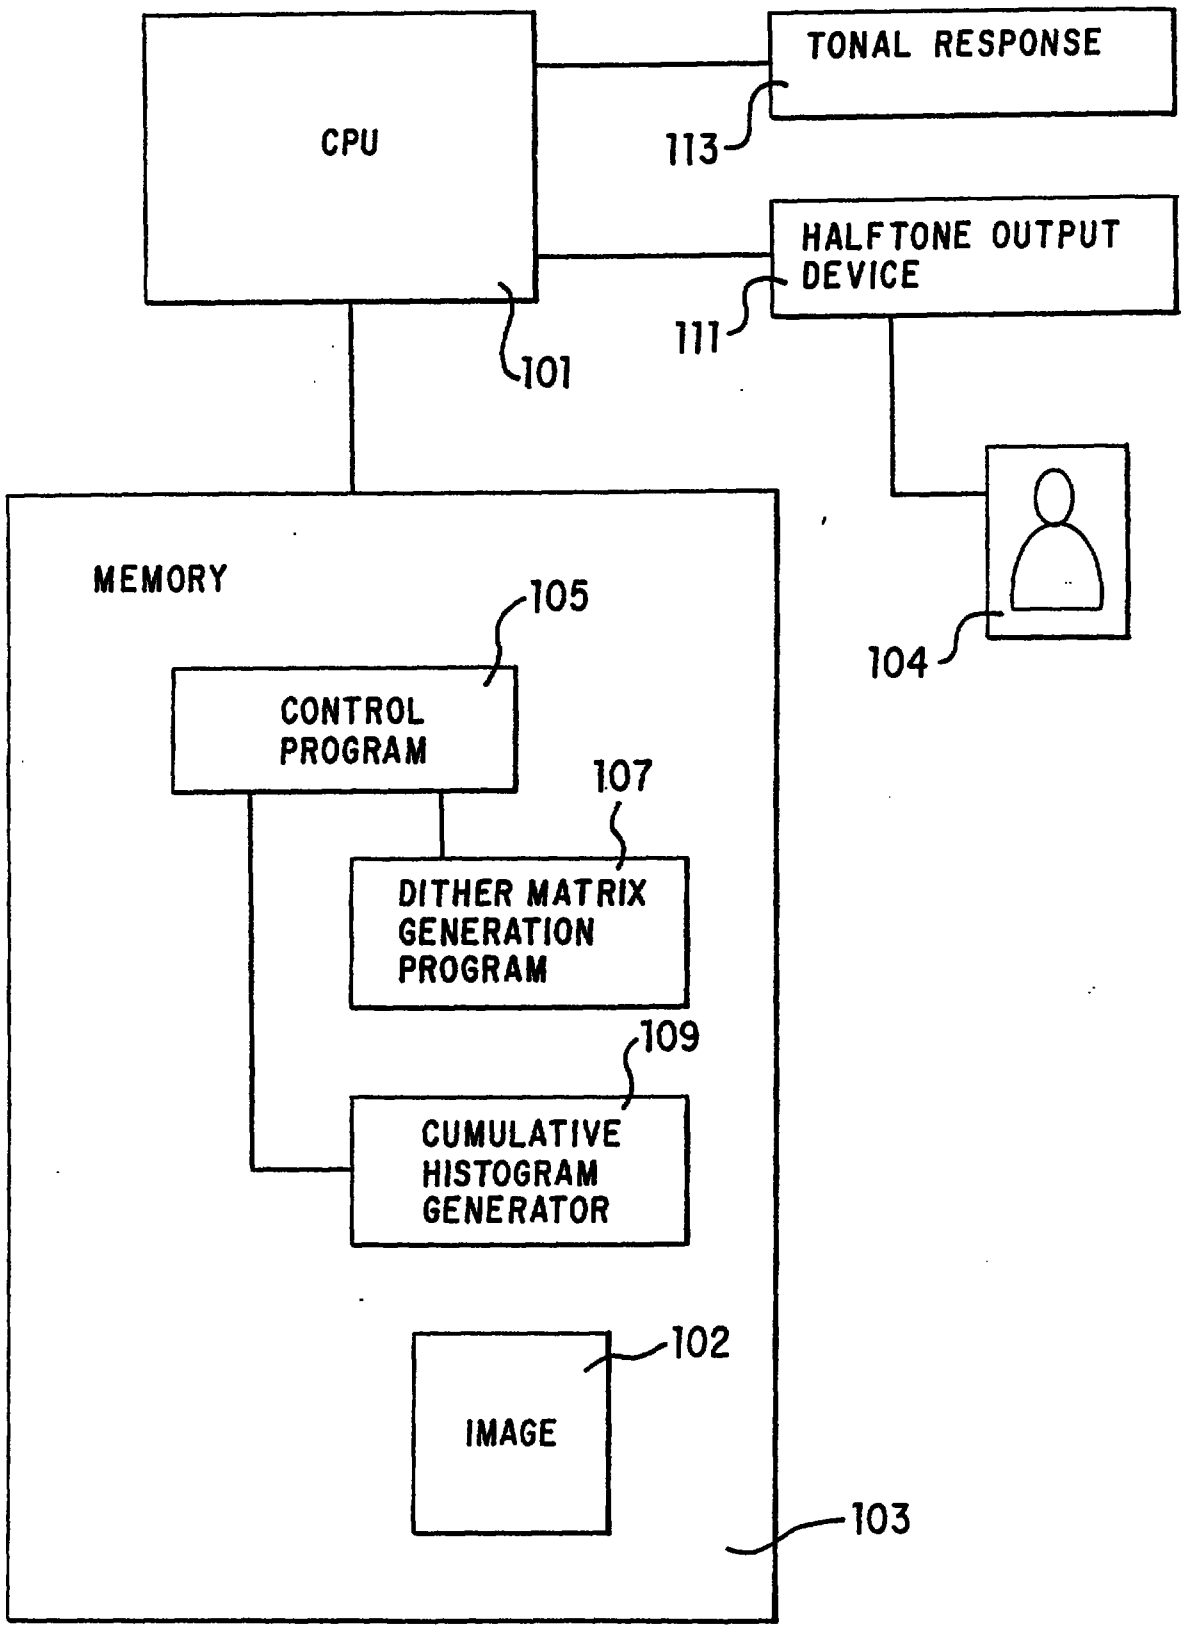

CPU
TONAL RESPONSE
113
HALFTONE OUTPUT DEVICE
111
101
104
MEMORY
105
CONTROL PROGRAM
107
DITHER MATRIX GENERATION PROGRAM
109
CUMULATIVE HISTOGRAM GENERATOR
102
IMAGE
103

FIG. 1

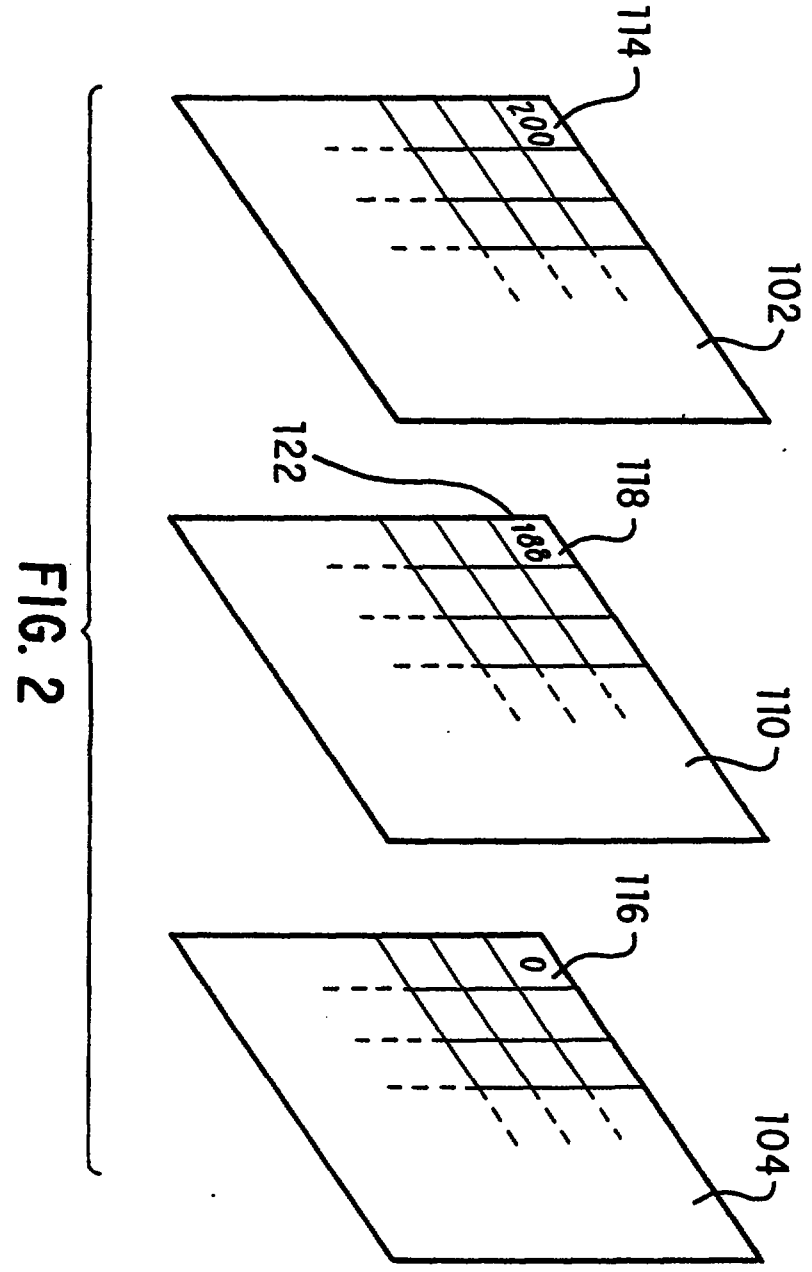


FIG. 2

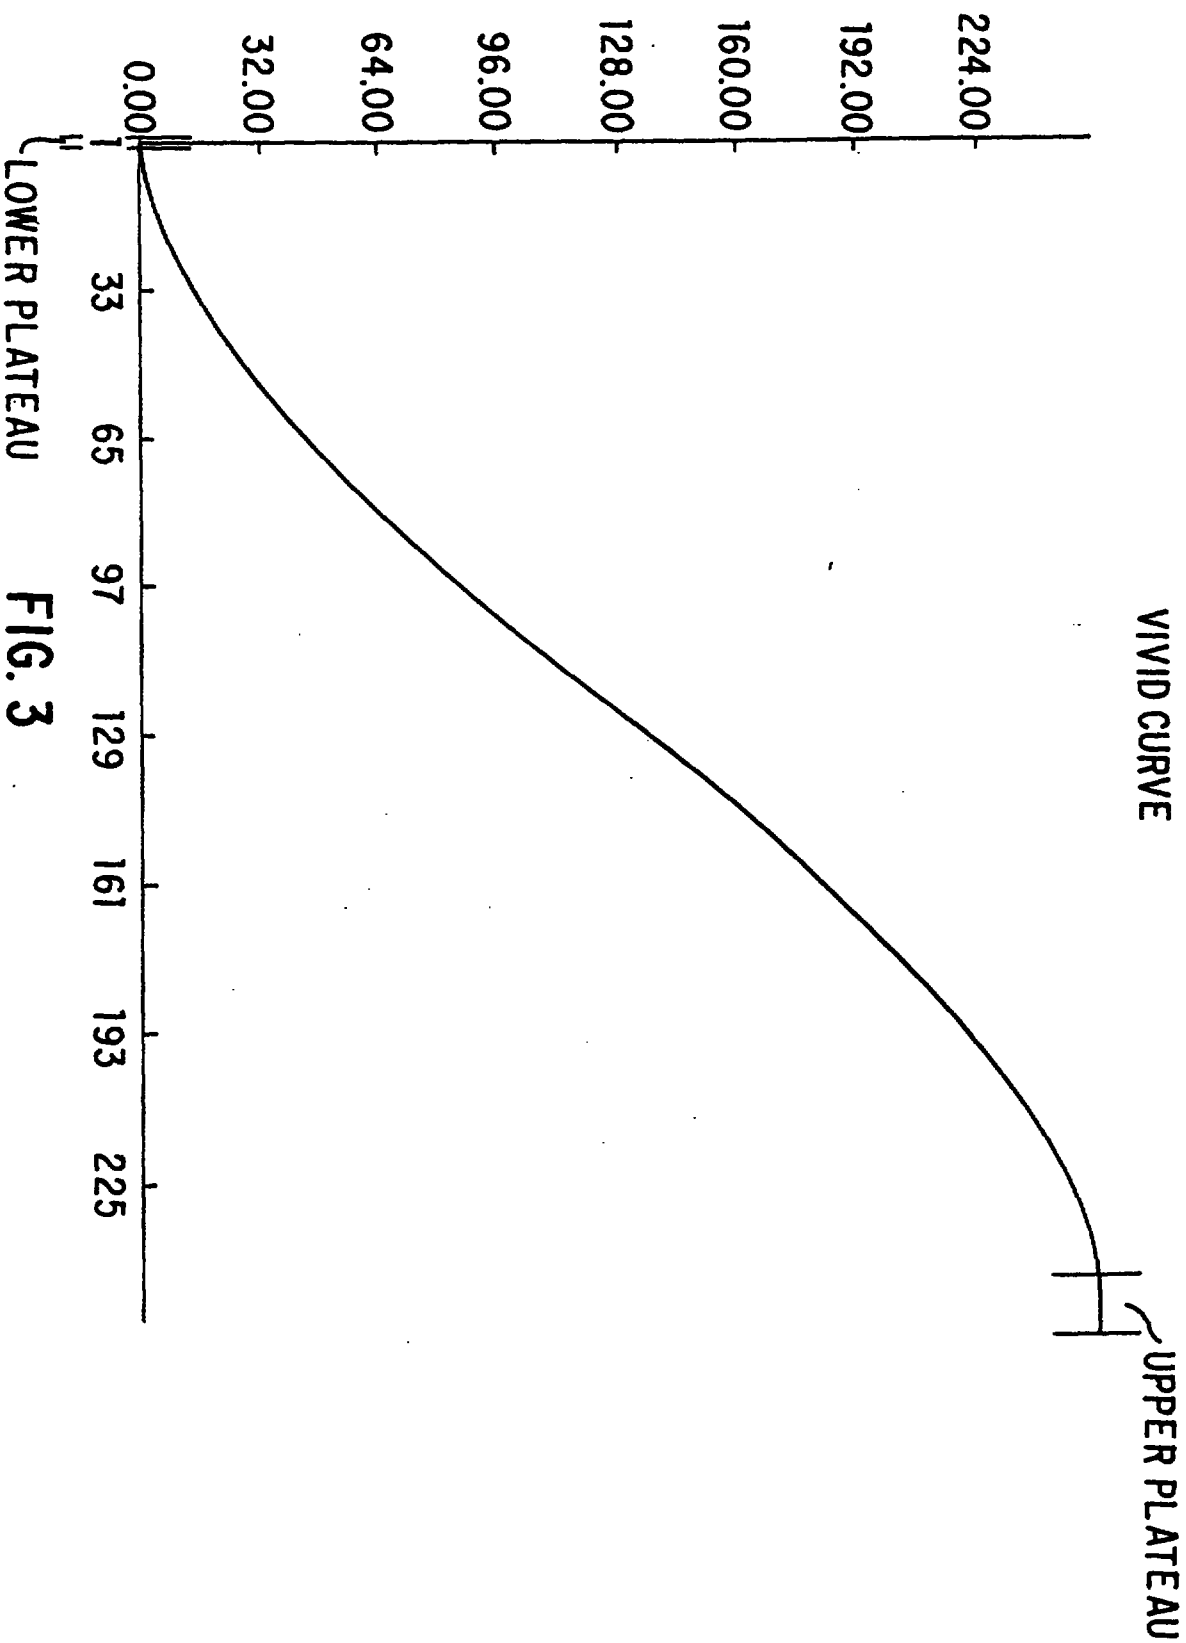

0.00
32.00
64.00
96.00
128.00
160.00
192.00
224.00
1
33
65
97
129
161
193
225
LOWER PLATEAU
VIVID CURVE
UPPER PLATEAU

FIG. 3

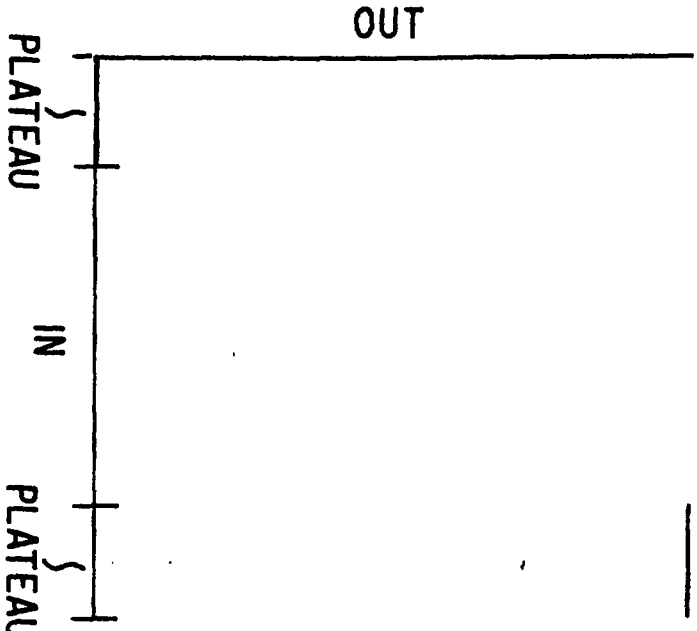


FIG. 4A

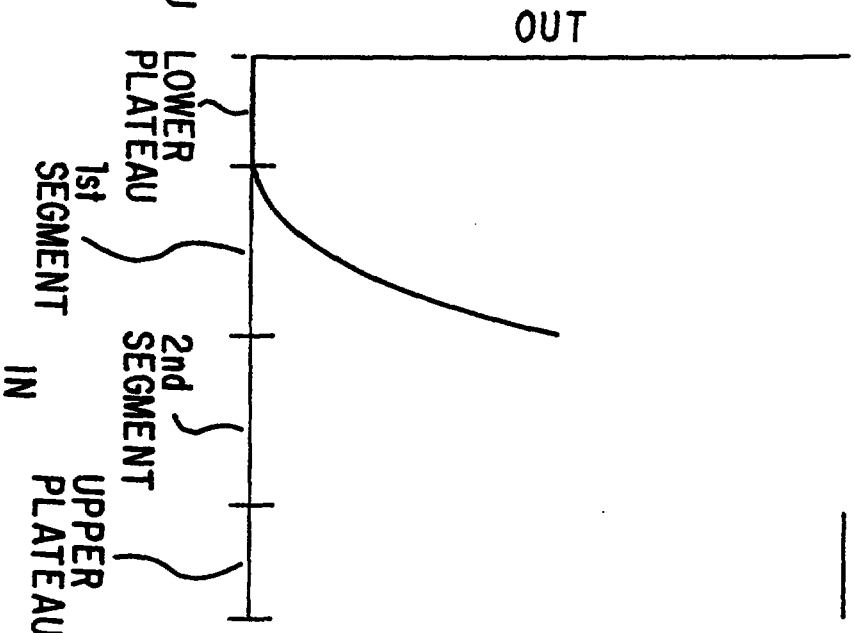


FIG. 4B

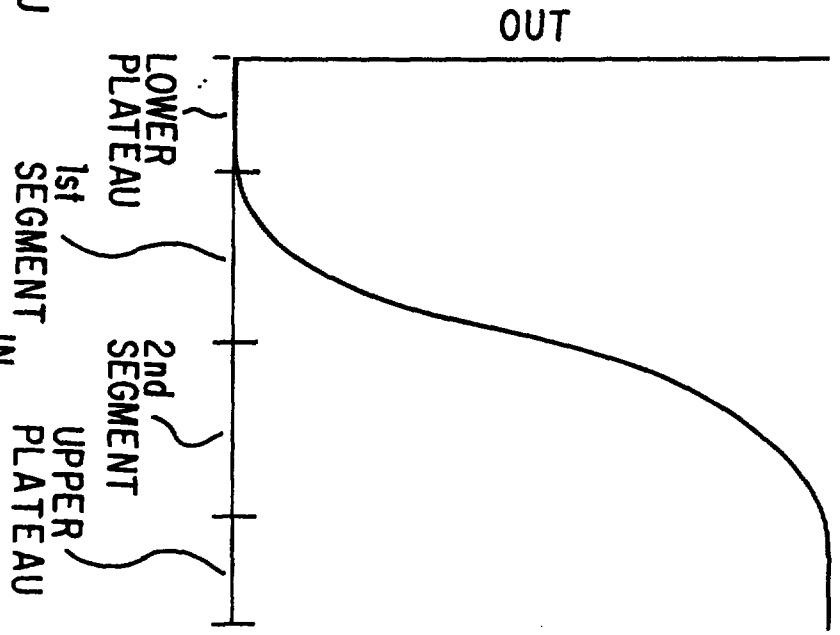


FIG. 4C

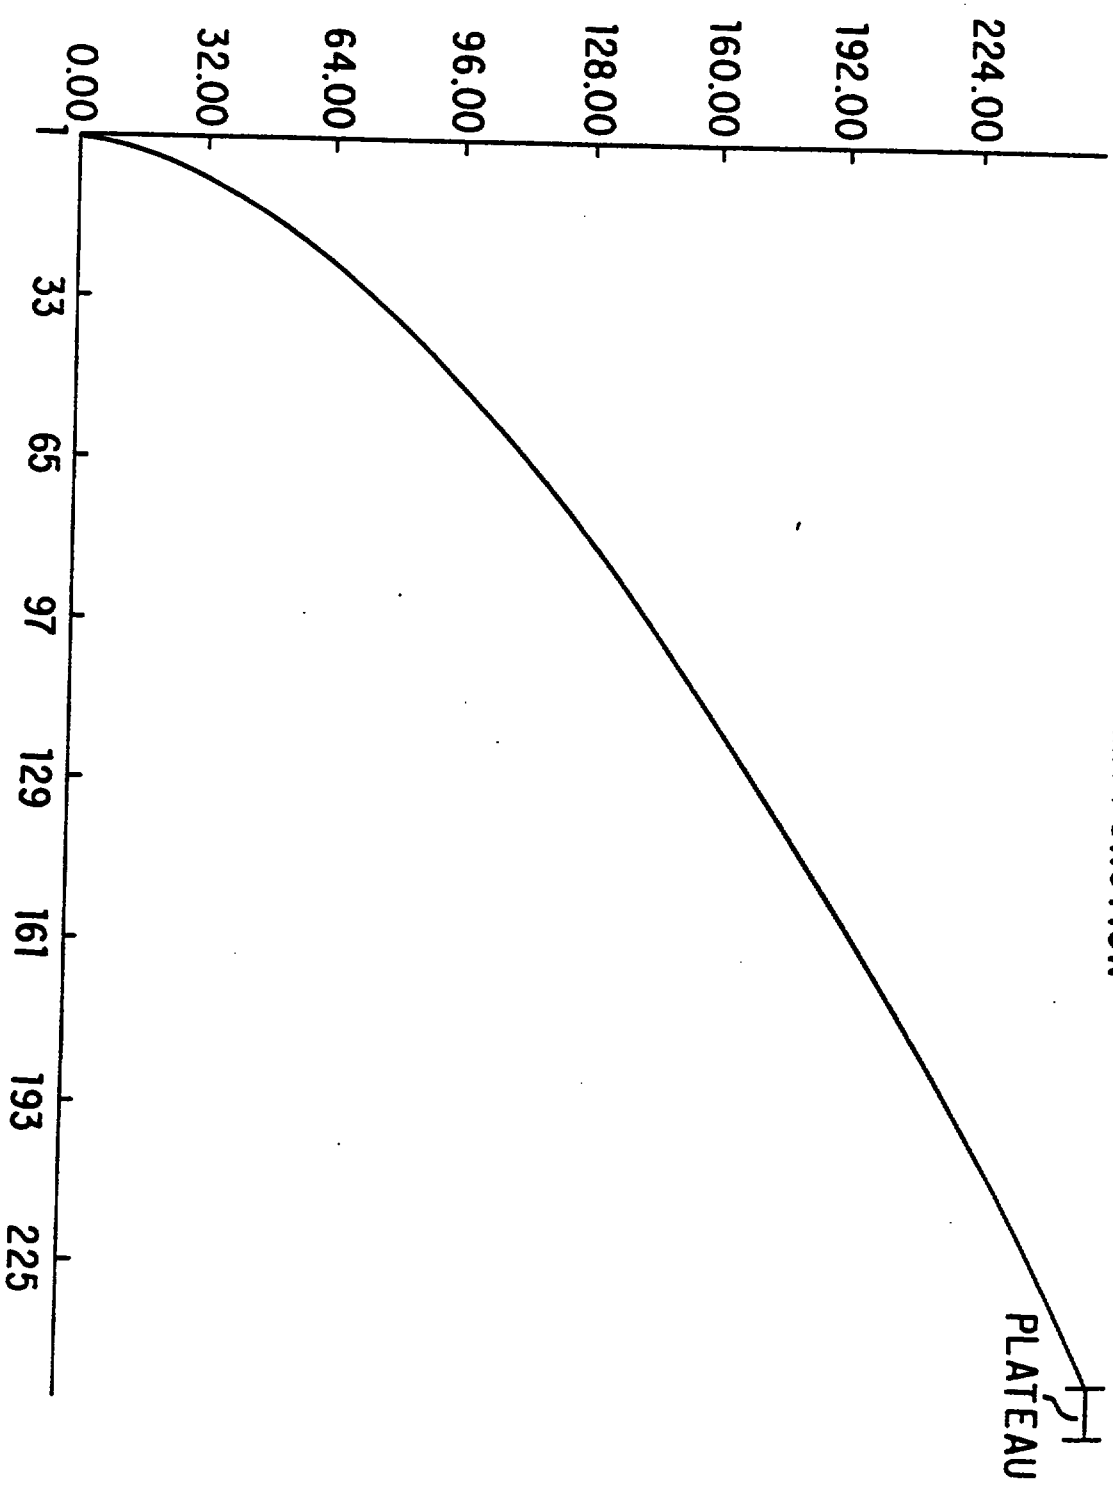

GAMMA FUNCTION
0.00
32.00
64.00
96.00
128.00
160.00
192.00
224.00
33
65
97
129
161
193
225
PLATEAU

FIG. 5

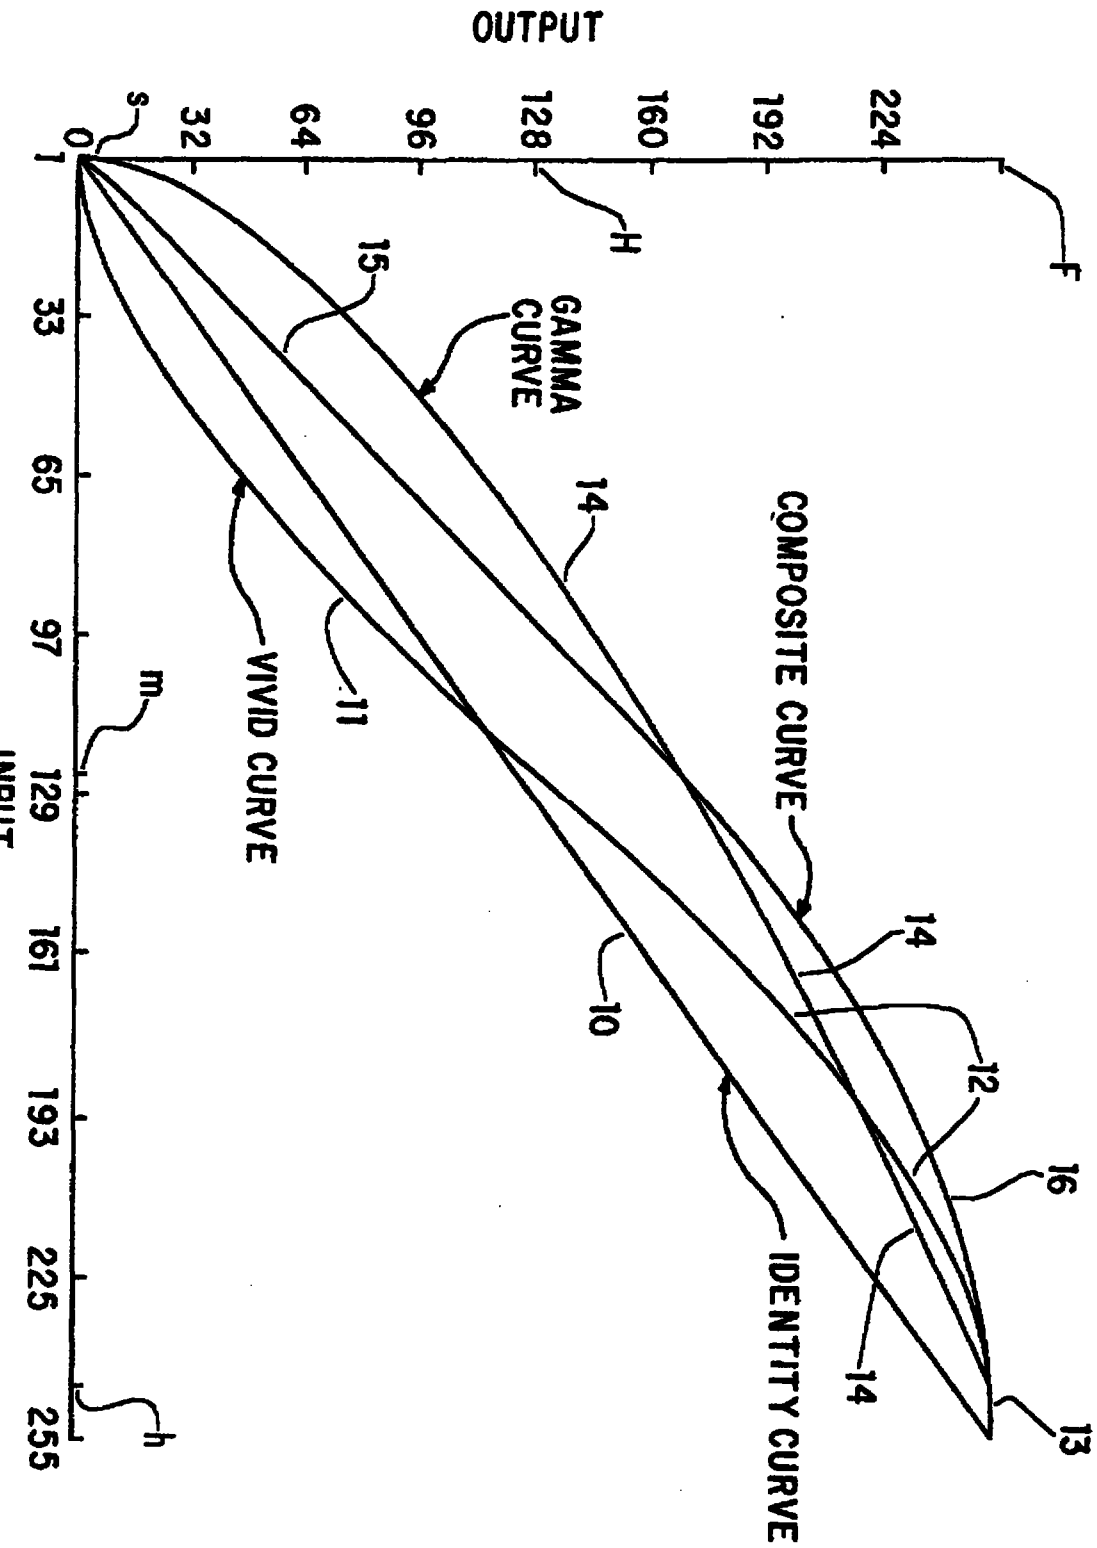

OUTPUT
0
32
64
96
128
160
192
224
s
H
F
1
33
65
97
129
161
193
225
255
m
h
INPUT
GAMMA CURVE
COMPOSITE CURVE
VIVID CURVE
IDENTITY CURVE
10
11
12
13
14
15
16

FIG. 6

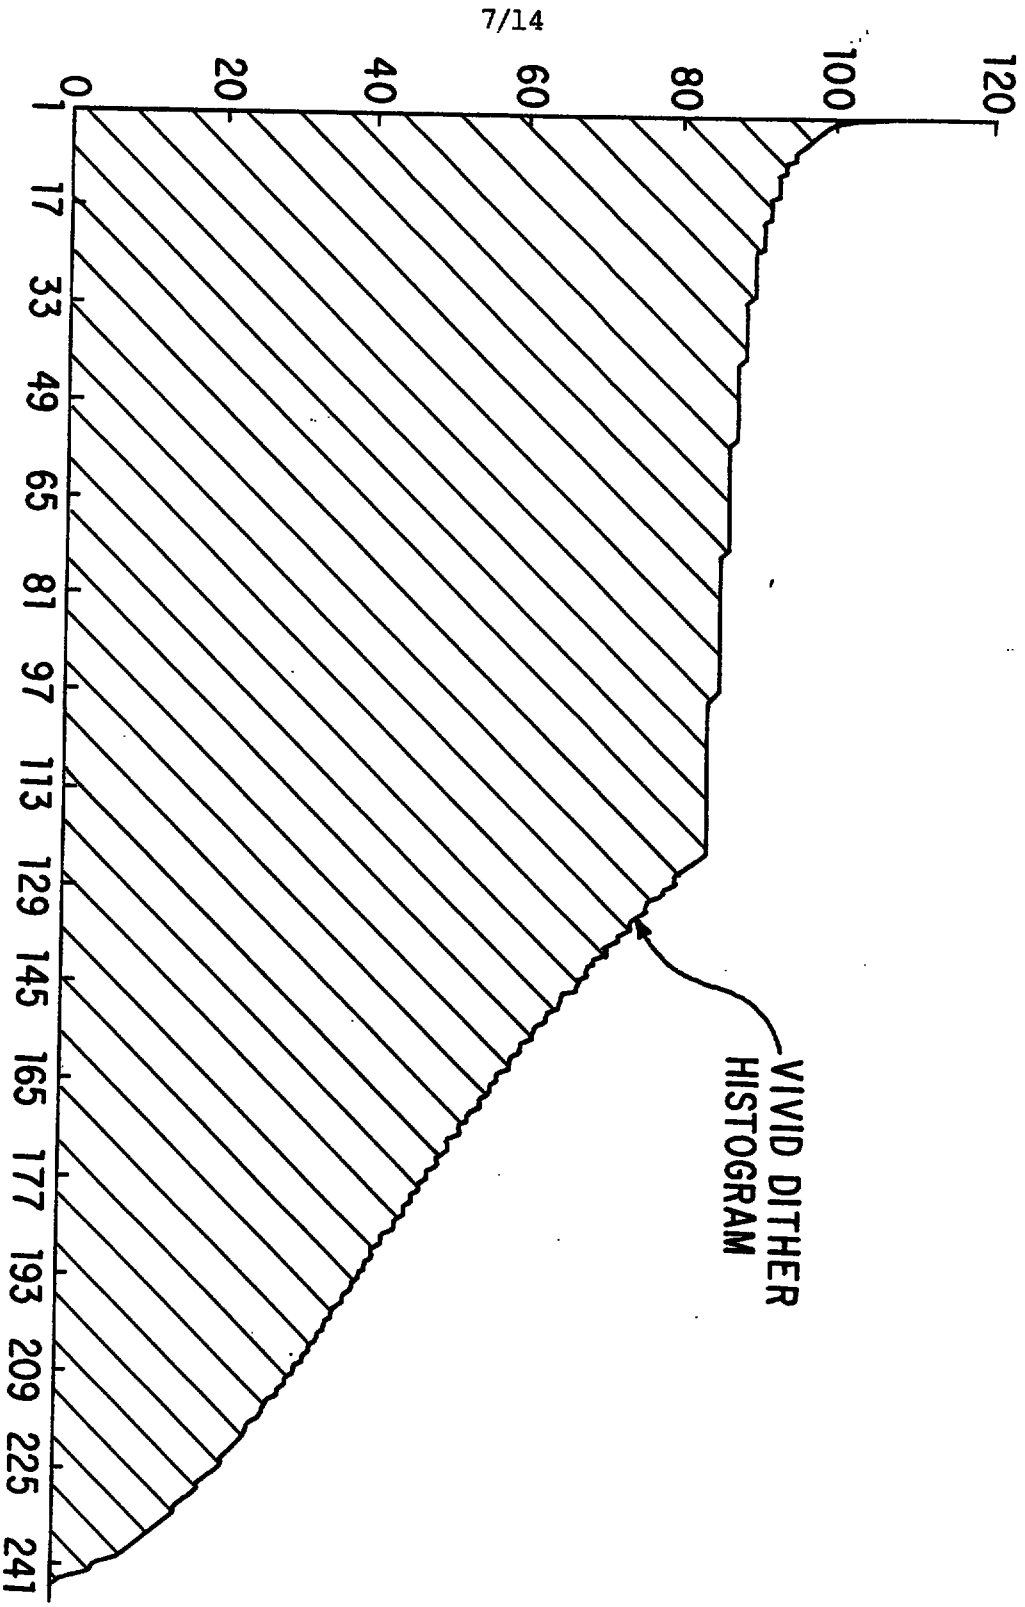

0
20
40
60
80
100
120
1
17
33
49
65
81
97
113
129
145
165
177
193
209
225
241
VIVID DITHER
HISTOGRAM

FIG. 7

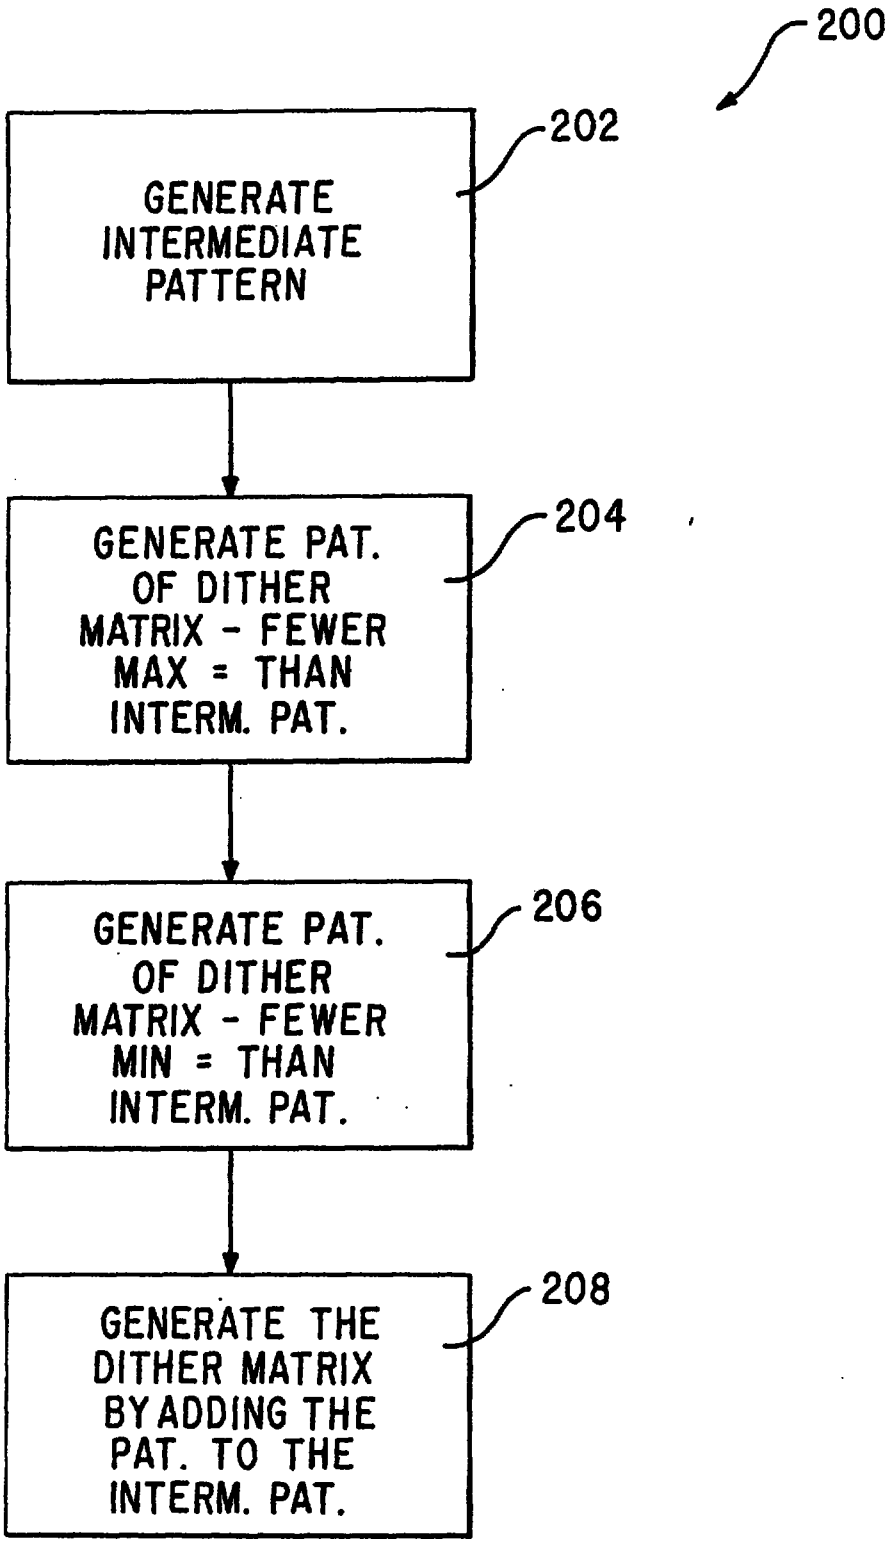

200
202
GENERATE INTERMEDIATE PATTERN
204
GENERATE PAT. OF DITHER MATRIX - FEWER MAX = THAN INTERM. PAT.
206
GENERATE PAT. OF DITHER MATRIX - FEWER MIN = THAN INTERM. PAT.
208
GENERATE THE DITHER MATRIX BY ADDING THE PAT. TO THE INTERM. PAT.

FIG. 8

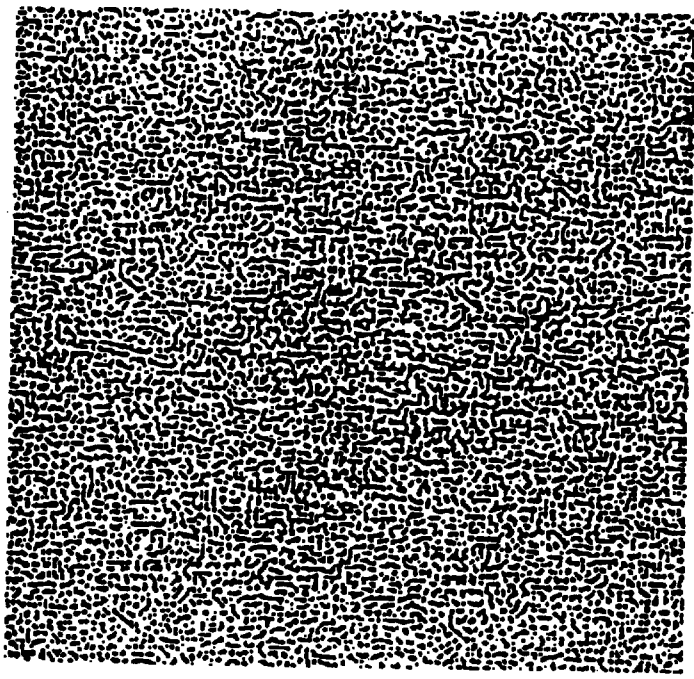


FIG. 9

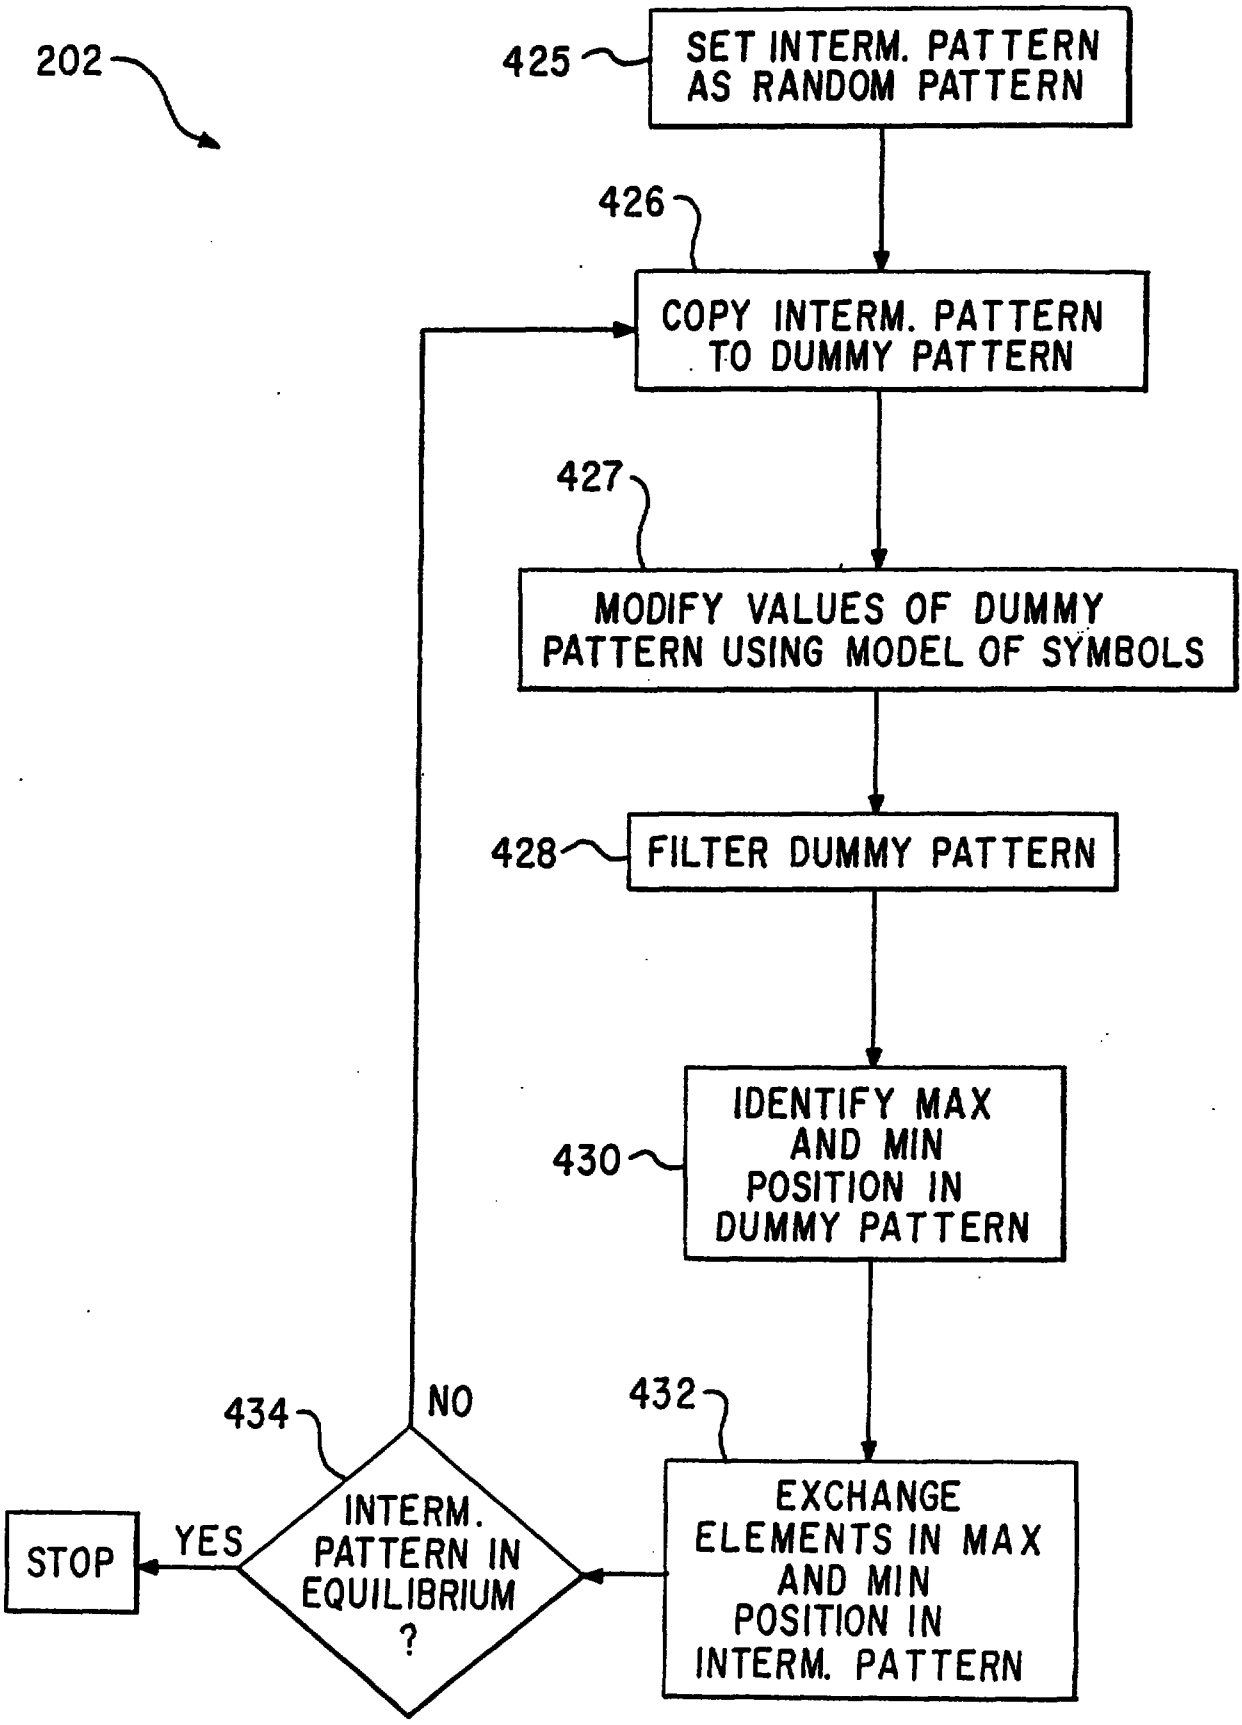

202
425
SET INTERM. PATTERN AS RANDOM PATTERN
426
COPY INTERM. PATTERN TO DUMMY PATTERN
427
MODIFY VALUES OF DUMMY PATTERN USING MODEL OF SYMBOLS
428
FILTER DUMMY PATTERN
430
IDENTIFY MAX AND MIN POSITION IN DUMMY PATTERN
432
EXCHANGE ELEMENTS IN MAX AND MIN POSITION IN INTERM. PATTERN
434
INTERM. PATTERN IN EQUILIBRIUM ?
NO
YES
STOP

FIG. 10

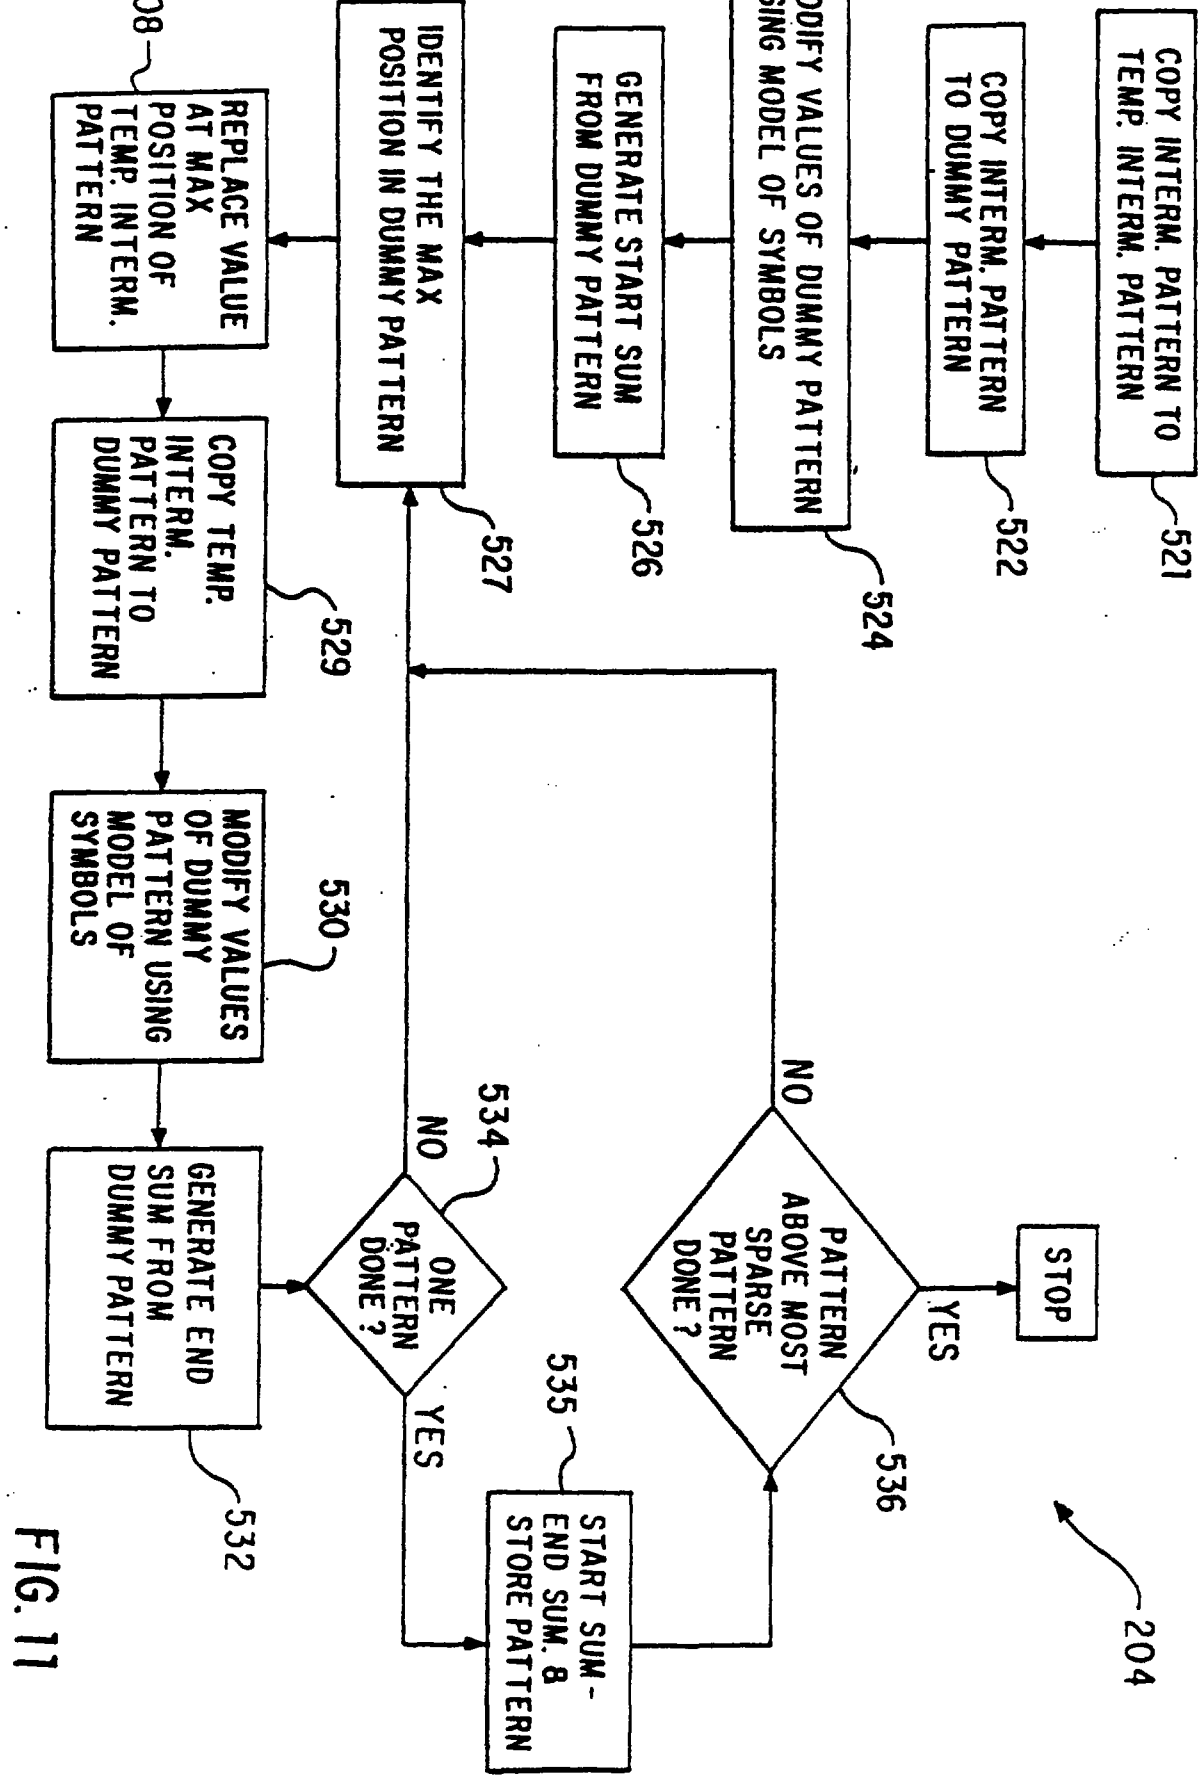

COPY INTERM. PATTERN TO TEMP. INTERM. PATTERN
521
COPY INTERM. PATTERN TO DUMMY PATTERN
522
MODIFY VALUES OF DUMMY PATTERN USING MODEL OF SYMBOLS
524
GENERATE START SUM FROM DUMMY PATTERN
526
IDENTIFY THE MAX POSITION IN DUMMY PATTERN
527
508
REPLACE VALUE AT MAX POSITION OF TEMP. INTERM. PATTERN
COPY TEMP. INTERM. PATTERN TO DUMMY PATTERN
529
MODIFY VALUES OF DUMMY PATTERN USING MODEL OF SYMBOLS
530
GENERATE END SUM FROM DUMMY PATTERN
532
534
ONE PATTERN DONE ?
NO
YES
535
START SUM - END SUM. & STORE PATTERN
PATTERN ABOVE MOST SPARSE PATTERN DONE ?
NO
YES
536
STOP
204

FIG. 11

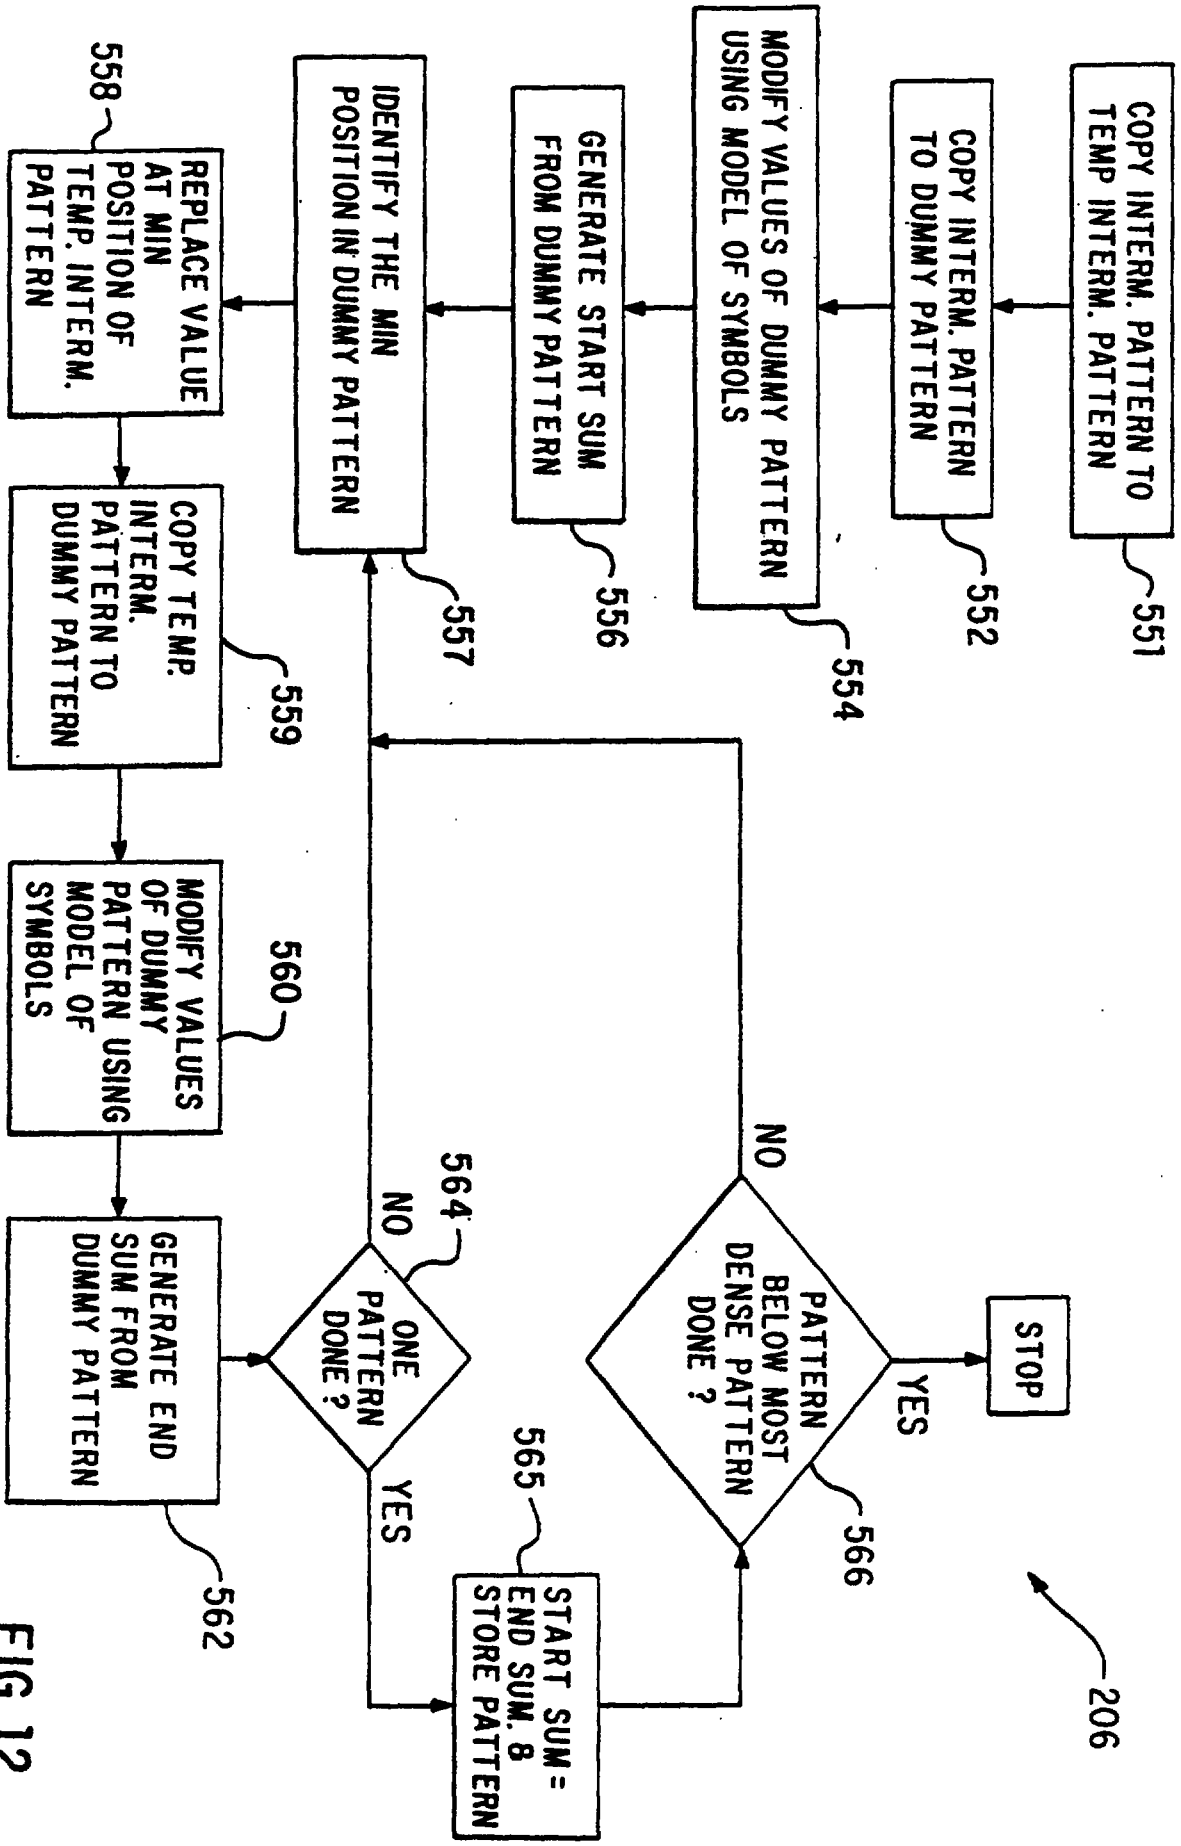

COPY INTERM. PATTERN TO TEMP INTERM. PATTERN
551
COPY INTERM. PATTERN TO DUMMY PATTERN
552
MODIFY VALUES OF DUMMY PATTERN USING MODEL OF SYMBOLS
554
GENERATE START SUM FROM DUMMY PATTERN
556
IDENTIFY THE MIN POSITION IN DUMMY PATTERN
557
REPLACE VALUE AT MIN POSITION OF TEMP. INTERM. PATTERN
558
COPY TEMP. INTERM. PATTERN TO DUMMY PATTERN
559
MODIFY VALUES OF DUMMY PATTERN USING MODEL OF SYMBOLS
560
GENERATE END SUM FROM DUMMY PATTERN
562
ONE PATTERN DONE ?
564
NO
YES
PATTERN BELOW MOST DENSE PATTERN DONE ?
NO
YES
566
START SUM = END SUM. & STORE PATTERN
565
STOP
206

FIG. 12

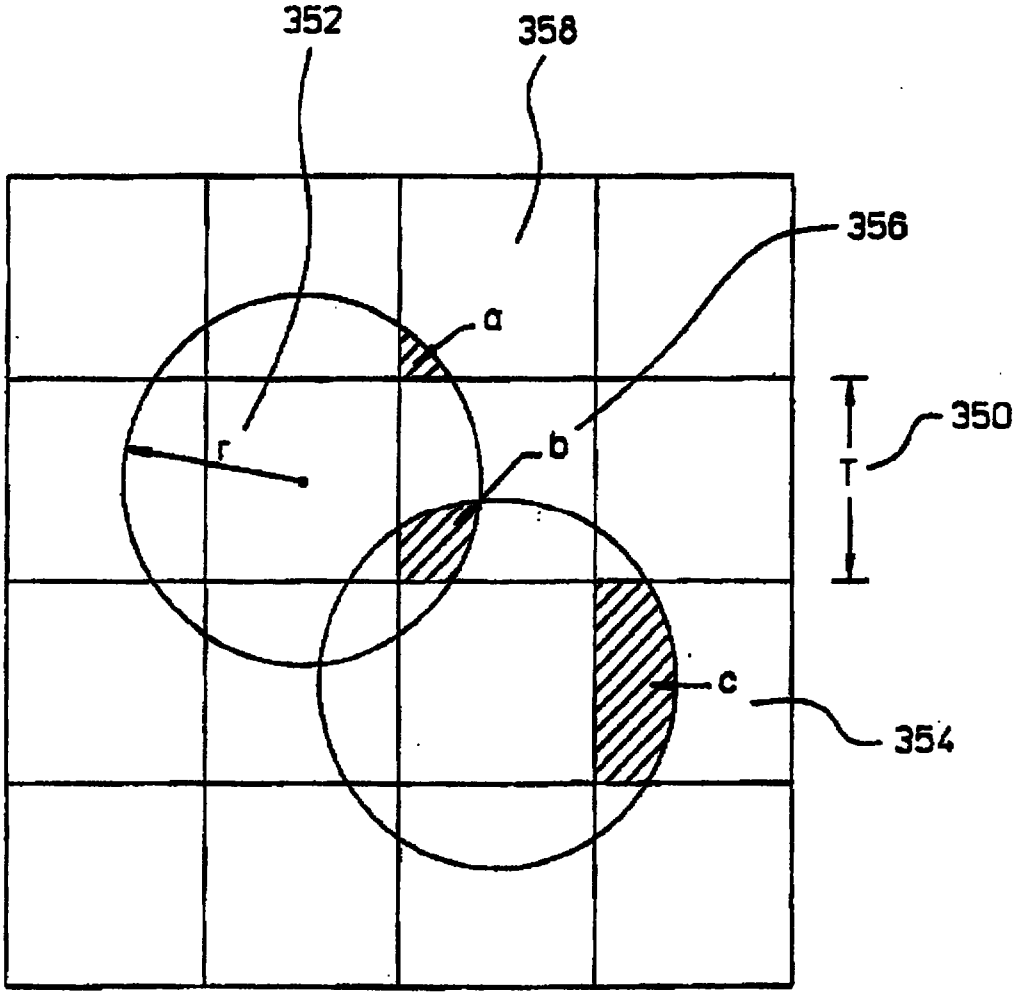


FIG. 13

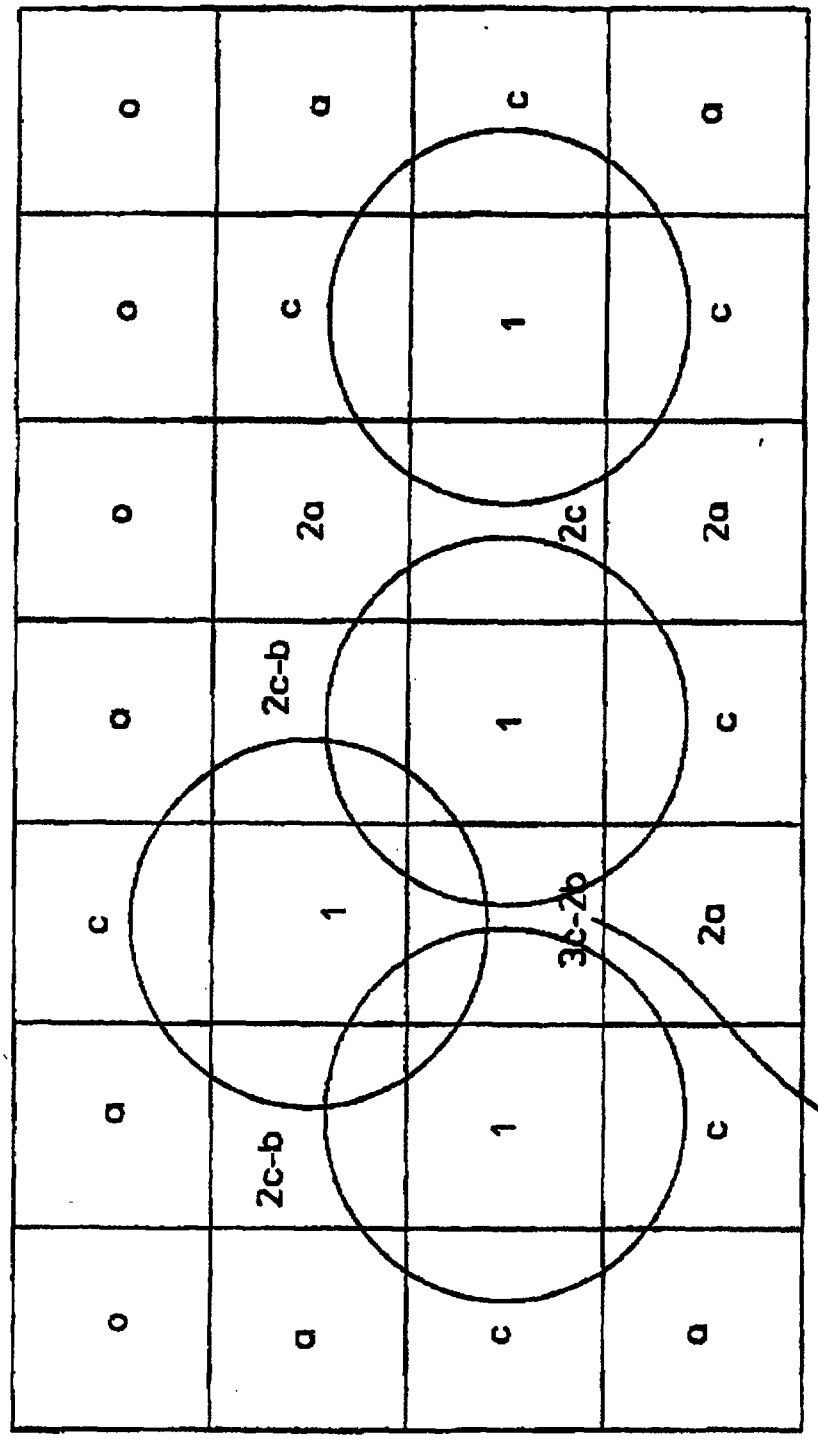


FIG. 14

HALFTONE DITHER CELL WITH INTEGRATED PREFERRED COLOR MATCHING

Abstract of the Disclosure

Methods and apparatus generate, and use in printing, a dither matrix (110) that incorporates a vivid-response function (y, 11-12; 15-16), which follows a generally S-shaped response curve (11-12). Printing with this pregenerated dither cell is just as fast as with a conventional cell, but printed color has a slightly exaggerated vividness, more pleasing to most users than more-accurately reproduced color. Color printed with this curve (11-12) alone, however, is found unsatisfactorily dark; the invention also encompasses including in the overall response (15-16) another function (z[x], Fig.5) to lighten the colors—particularly at the low end of the brightness range. We define both functions as so-called "gamma functions", but with different arguments (x, y) and exponents (γ_1,γ_2), and construct the S-shaped curve as two separate sections (11, and Fig. 4B; vs. 12, and right end of Fig. 4C) blended at a generally central point (H, m). Preferably the argument of the first (low-brightness) section (11) is an expression ([x - s]/m) that is proportional to an input color signal (x); the function raises this argument to a power (γ_1) greater than unity (ideally 1.6). The second (high-brightness) section (12) of the first function is essentially a mirror image of the first, flipped over first upward and then to the "right" (i. e., in both steps toward higher brightness); and is defined by an equation closely analogous to the first but using a different argument ([h - x]/m) to measure the signal from the high end (h) of the input range, and subtracting the curve value from the maximum output F. The second gamma function has as its argument (y) the output value of the first, and an exponent (γ_2) smaller than one—ideally 0.6.

(Figure 6)