

METHOD AND APPARATUS FOR DATA COMPRESSION

BACKGROUND OF THE INVENTION

5 1. FIELD OF THE INVENTION

This invention relates to the field of compression and decompression of data.

10 2. BACKGROUND ART

Information is represented in a computer system by binary data in the form of 1s and 0s. This binary data is often maintained in a data storage device. In a computer system, data storage is a limited resource. To more efficiently use data storage resources, data is often compressed prior to storage so that less storage area is required. When the data is retrieved, it is decompressed for use. The need for compression can be demonstrated by describing the way that images are represented in a computer system, the transformation of such images into a form suitable for printing, and the storage problems associated with such images. This discussion is followed by descriptions of compression techniques and prior art approaches to compression.

If a person were to look closely at a television screen, computer display, magazine page, etc., he would see that an image is made up of hundreds or thousands of tiny dots, where each dot is a different color. These dots are known as picture elements, or "pixels," when they are on a computer display and as dots when printed on a page. The color of each pixel is represented by a number value. To store an image in a computer memory, the number value of each pixel of the picture is stored. The number value typically represents the color and intensity of the pixel.

The accuracy with which a document can be reproduced is dependent on the "resolution" of the pixels that make up the document. The resolution of a pixel is determined by the range of the number value used to describe that pixel. The range of the number value is limited by the number of "bits" in the memory available to describe each pixel (a bit is a binary number having a value of 1 or 0). The greater the number of bits available per pixel, the greater the resolution of the document. For example, when only one bit per pixel is available for storage, only two values are available for the pixel. If two bits are available, four levels of color or intensity are available. While greater resolution is desirable, it can lead to greater use of data storage. For example, if each pixel is represented by a 32-bit binary number, 320,000 bits of information would be required to represent a 100 X 100 pixel image. Such information is stored in what is referred to as a "Frame Buffer" or gray array ("G array").

A black and white printer has resolution of only one bit per pixel or dot. That is, the printer is only capable of printing a black dot at a location or of leaving the location blank. When an image is to be printed on a black and white printer, the image must be transformed so that its bit resolution matches the bit resolution of the printer. This transformation is known as "thresholding" and consists of determining, for each pixel in the source image, whether the dot to be printed at the corresponding location on the printed page is to be black or white.

Although the printer can only do black and white printing, a printed image can appear to have many different shades of gray depending on the pattern of black and white dots. When every other dot is black, for example, the resulting printed image will appear gray, because the human eye blends the tiny dots together. Many printers are capable of printing 600 dots per inch in the horizontal and vertical directions. Because of the large number of tiny dots, other shades of gray can be simulated by the relative percentage of black and white dots in any region. The more black dots in a region, the darker that region appears.

As noted above, when thresholding, a decision is made for each pixel, based on its original color in the source image, of whether to print a black or white dot on the page for that pixel. Consider a thresholding scheme where each pixel in the stored grayscale image may be represented by 8 bits, for example, giving 256 (2^8) possible values. One thresholding method that does not produce very realistic images is to assign a black value to all image pixels with a value of 128 (out of 256) or above, and a white value to all image pixels with a value of 127 or below. Using thresholding, an entire multi-bit depth frame buffer can be compressed into a one bit per pixel buffer. However, the resulting image is "aliased" (appears like steps or contains jagged edges) and does not approximate the original image. To produce better images, a threshold matrix is generated and used to determine the thresholded value of an image pixel.

A threshold matrix uses different threshold values for an image pixel, depending on the address of the image pixel in the array. Thus, each cell of the frame buffer corresponds to a threshold matrix cell which has an independent threshold level. The threshold matrix need not be the same size and is often smaller than the G array. For example, at one location, an image pixel may be thresholded to black if its value is 128 or above, while an image pixel at another location may be black only if its value is 225 or higher. The result of applying the threshold matrix is an array of 1s and 0s that could be printed to represent the original continuous tone image.

Figure 1A depicts a frame buffer (G array), with indices i and j ($G[i][j]$). Figure 1B depicts a threshold matrix (T array), with indices i' and j' ($T[i'][j']$). Figure 1C depicts the resulting output or pixel array (P array) with i rows and j columns ($P[i][j]$). Thus, for example, if the pixel maintains a value of 123 (G_{11} of Figure 1A) and the threshold level is 128 (T_{11} of Figure 1B), the resulting value is 0 (P_{11} of Figure 1C) due to the fact that the pixel value is less than the threshold level. Hence the resulting pixel array is created by thresholding the G array as follows:

$$P[i][j] = \begin{cases} 1 & \text{if } G[i][j] \geq T[i'][j'] \\ 0 & \text{otherwise} \end{cases}$$

where $G[i][j]$ is an array of the same dimensions as P but takes on many values, typically 0,1,...,255. $T[i'][j']$ is a threshold array, a matrix of dimension n threshold
5 rows by m threshold columns and taking on values in a range like G , typically 1,2,...,255. For the thresholding (i',j') is a function of (i,j) typically:

$$i' = i \text{ modulo } n$$

$$j' = j \text{ modulo } m$$

10

where modulo means the remainder after division.

After this thresholding step, the entire page can be represented in a memory of 1s and 0s by the same number of bits as there are dots on the page. Even at 1 bit per dot, the amount of memory required can be substantial. For a page that is 8.5 by 11
15 inches and has a resolution of 600 dots per inch (dpi), the amount of memory needed is approximately 4.2 megabytes of memory (if monochrome). This memory in a printer is referred to as a "buffer". Memory is an expensive component, and it is an advantage to reduce the amount of memory required in a printer buffer. In the past, this has been accomplished by applying a "compression algorithm" to the data in the buffer. Despite
20 the significant compression which may arise from thresholding, further compression is desired.

There are currently compression schemes for single bit data. Some of these schemes work preferentially better on data that was originally text and some work better on data that was originally an image. Some of these schemes include facsimile
25 standards, such as the ITU standards, using Huffman encoded run lengths and the JBIG (Joint Bi-level Image Group) standard.

There are two distinct families of prior art compression schemes: lossy and lossless. Lossless compression guarantees that no data will be lost upon a compression and decompression sequence. For example, one lossless compression scheme accomplishes this guarantee by searching the data for any repeating sequences such as

5 "001001001001001001". Using this lossless compression scheme, the sequence "001" would be stored along with the number of times it recurs - six. However, lossless compression schemes may not provide a satisfactory level of compression, e.g., due to the absence of repeating sequences in the source data for the above example. Nonetheless, due to its accuracy, lossless compression is used when storing

10 database records, spreadsheets, or word processing files.

A lossy scheme achieves a greater level of compression while risking a loss of a certain amount of accuracy. However, certain types of stored information do not require perfect accuracy, such as graphics images and digitized voice. As a result, lossy compression is often utilized on such types of information.

15 The most interesting prior art method is set forth in the JBIG standard. Pixels are processed in the usual scan order, i.e., row_i is entirely processed before row_{i+1} and in each row, $column_i$ precedes $column_{i+1}$. Encoding the current pixel uses a context, which comprises a set of nearby pixels that have already been encoded. If, e.g., ten pixels were used as a neighborhood, there would be 2^{10} possible contexts. Based on

20 the frequency with which the current context has been previously encountered, the encoder makes a prediction and estimates the probability that the prediction is accurate. If the estimated probability is reliable and near to certainty, then an arithmetic entropy encoder can losslessly encode the prediction error (e.g., 0 for no error, 1 for error) with much less than 1 bit per pixel. The decoder has the same context and frequency

25 information, and is therefore able to interpret the 0 or 1 to determine the true value of the pixel.

See also U.S. Pat. No. 5,442,458, issued to Rabbani et al., which is directed to an encoding method for image bitplanes using conditioning contexts based on pixels from the current and previous bitplanes.

5

SUMMARY OF THE INVENTION

The present invention provides a method and apparatus for compression of data. Whereas the performance of compression schemes of the prior art are often dependent on the type of data being compressed, the invention provides for the application of a plurality of compression schemes to the data such that improved compression ratios are achieved. A first embodiment provides for compression of each pixel by one of a plurality of different entropy-based compression schemes based upon a probability cost analysis. A second embodiment provides for compression of each pixel based on a hybrid context formed using a plurality of compression schemes for improved probability determination, and thus improved entropy encoding.

The first embodiment of the invention provides a method for choosing the most effective compression scheme per pixel. A multiplicity of compression schemes are utilized and a cost value is associated with each scheme on a pixel by pixel basis. After associating a cost with each scheme, the method selects between the most effective or lowest summed cost scheme on a pixel by pixel basis. The selected scheme provides a predicted value of the current pixel and an estimate of the probability that the prediction is correct. The correctness of the estimate, either true or false, is then encoded by an entropy encoder which uses the estimated probability to encode the true or false outcome in less than one bit, provided that the estimated probability is accurate and not in the vicinity of 0.5. Typically, the estimated probability is greater than 0.95. In a preferred embodiment, the entropy encoding is performed using an arithmetic encoding process.

In one specific embodiment, one of the compression schemes utilized is referred to as the "inverse" scheme. The inverse scheme predicts the gray value of a frame buffer pixel. The scheme examines a set of recently scanned pixels to define a range within which the pixel under consideration is likely to fall. The range is determined as follows.

For each previously scanned pixel, the thresholded value and the associated threshold define a range. For example, if the threshold is 150 and the thresholded value is 1, the pre-thresholded value is in the range of 150 to 255. If the thresholded value is 0, the pre-thresholded value is in the range of 0 to 149. The resulting range is intersected with the similarly determined range from the next closest neighbor pixel.

The intersecting of the threshold matrix ranges continues for each previous neighbor pixel until all pixels have been analyzed or until there is a contradiction, namely that a range is encountered which has no common overlap with the most recently calculated range. A value within the most recently calculated range is then selected as an estimate of the value of the pixel under consideration. The difference between the corresponding threshold value and the estimated value is then calculated. Subsequently, a probability table is updated which records the number of times a binary one (1) or zero (0) resulted with that difference calculation. This probability table is used in the encoding stage.

A second scheme utilized in the specific embodiment above is one similar to that of JBIG, referred to as the "context" scheme. A set of pixels in a frame buffer are selected. Subsequent to the pixel set's thresholding, the thresholded values are concatenated to provide a binary number. A frequency table is maintained which records the number of times a binary one (1) or a binary (0) occurs for each sequence of binary numbers for the pixel set. The concatenated binary number is used as an index into the frequency table, from which the probabilities for use in the encoding stage are derived.

In the second embodiment of the invention, a hybrid context is formed using a first number of bits of the hybrid context to store previous pixel information gathered using a first compression scheme, and using further portions of the hybrid context to store previous pixel information gathered using other compression schemes. Statistics
5 are stored in a table indexed by the hybrid context, and entropy encoding is performed based on the probability information in the table indexed by the hybrid context for the current pixel. The combination of different probability determining schemes results in greater probability values for each pixel, and correspondingly higher compression ratios.

10 In a further specific embodiment, a thirteen bit hybrid context is formed using seven bits to store a quantized gray value difference determined using the inverse scheme, and using six bits to store recently scanned pixel values in the neighborhood in accordance with the context (JBIG) scheme. A table, indexed by the hybrid context, is formed containing the probability for a given pixel, having the respective six-bit JBIG
15 context and the respective seven-bit gray value difference, to be "0" or "1". An entropy encoder encodes the pixel value based on the prediction and the probability determined from the hybrid context.

BRIEF DESCRIPTION OF THE DRAWINGS

20

Figure 1A is a sample table representing a frame buffer.

Figure 1B is a sample table representing a threshold matrix.

Figure 1C is a sample table representing the resulting pixel array from thresholding figure 1A with 1B.

25 Figure 2A is a sample table representing a frame buffer with empty cells to be estimated by the present invention.

Figure 2B is a sample table representing a threshold matrix which is utilized in predicting the empty cells of figure 2A.

Figure 2C is a sample table representing the resulting pixel array which is utilized in predicting the empty cells of figure 2A.

Figure 3 is a sample probability table which is utilized in the context scheme to maintain statistics.

5 Figure 4 is a flow diagram demonstrating the switching method of the present invention.

Figure 5 is a flow diagram demonstrating the inverse scheme.

Figure 6 is a flow diagram demonstrating the process of calculating the g-estimate of an interval.

10 Figure 7 is a flow diagram demonstrating the context method.

Figure 8 is an illustration of one embodiment of a hybrid context.

Figure 9 is a flow diagram of an encoding process implementing a hybrid context.

15 Figure 10 is a mapping diagram for difference values in one embodiment of the invention.

Figure 11 is a pixel diagram illustrating nearest neighbors for a modified JBIG scheme.

Figure 12 is a threshold diagram for a multi-bit embodiment.

Figure 13A is a block diagram of a switching embodiment of the invention.

20 Figure 13B is a block diagram of a hybrid context embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

A method and apparatus for compression of data is described. In the following
25 description numerous, specific details, such as the use of the context/JBIG and inverse schemes, are set forth in order to provide a more thorough understanding of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well

known features have not been described in detail so as not to unnecessarily obscure the present invention.

Current compression methods are comprised of a modeling and an encoding stage. During the modeling stage, the probability that a pixel value will be a binary one
5 (1) or zero (0) is estimated. During the encoding stage, this probability is utilized to compress or encode the pixels. A disadvantage associated with limiting the compression to one modeling scheme is that some schemes work better on data that was originally text and some work better on data that was originally an image.

The present invention overcomes this disadvantage by applying a plurality of
10 different modeling schemes to each pixel. In one embodiment, the least costly scheme in terms of bits is used to compress the data. In a second embodiment, each scheme is reduced and combined into a single hybrid compression scheme having the advantages associated with each of its constituent schemes.

One embodiment of the invention, as illustrated in Figure 4, provides a method
15 of selectively switching between the most effective compression schemes on a pixel by pixel basis. The first step of the switching method is that of choosing a multiplicity of compression modeling schemes which are going to be used 400. In one embodiment of the invention, a scheme similar to that of JBIG, referred to as the context scheme, is utilized as one of the schemes in the modeling stage, and the "inverse" scheme is
20 utilized as another scheme in the modeling stage.

Prior to selecting a modeling scheme, a cost value is calculated for each scheme
404. Thus, the cost for the context scheme and the cost for the inverse scheme are calculated (see discussion below). The scheme with the lowest summed cost in a pixel neighborhood around the current pixel is selected 406, and the current pixel is encoded
25 using the selected scheme 408. Subsequent to the encoding, the next pixel to be encoded is selected, and the process is repeated. Such a calculation of costs and encoding for each pixel permits the most effective or lowest cost scheme to be utilized to encode each pixel independently.

In the above embodiment, the inverse scheme predicts the gray value of a frame buffer pixel. The inverse scheme, which is discussed more fully below, computes a probability that the current pixel is a 0 and a probability that it is a 1, the two probabilities summing to one. If either probability is high, i.e. close to one, then much less than one bit is required to encode the pixel bit, resulting in compression.

To establish the probability, the inverse scheme predicts the value of the current pixel by examining/analyzing its neighboring pixels. The threshold ranges of the neighboring pixels are intersected until either all of the ranges have been intersected or the ranges cease to intersect. If a neighboring pixel range does not intersect with the previously intersected pixel ranges, the previously intersected pixel ranges are used as the resulting range. A value within the resulting range is selected and differenced with the threshold value for the pixel to be estimated. A table is then created which records the probabilities that the true binary value will be a one (1) or a zero (0) based on the difference value. The probability that a pixel will be a zero (0) or a (1) may be utilized in the encoding stage.

Additionally, in the above embodiment, a scheme similar to that of JBIG, referred to as the "context" scheme, is utilized. A set of pixels in a frame buffer are selected and thresholded. The thresholded values are then concatenated to provide a binary number. A frequency table is compiled which records the number of times a binary one (1) or a binary (0) occurs for each sequence of binary numbers. The concatenated binary number is used as an index into the frequency table. As a result, the frequency table maintains the probabilities which may be utilized in the encoding stage. The values in the frequency table may be scaled down at intervals to reduce the storage requirements of the table.

Though various embodiments are described with reference to a "frequency" table, other means for estimating the probability of the current context may be used to implement an embodiment of the invention. Typically, a numerical value representative of a statistical history of a context is stored, though the extent of the history may vary

for different embodiments. The probability value and prediction are determinable from the numerical value representation, but need not be explicitly present in the table. Further, the stored numerical value may be updated (e.g., as a byproduct of an entropy encoding engine) with each context occurrence by increasing or decreasing the
5 respective numerical value based on the current pixel prediction and current actual pixel value.

INVERSE COMPRESSION SCHEME

10 The inverse scheme is a scheme which estimates the probability during the modeling stage for later use in the encoding stage of the present invention. Referring to Figure 4, the modeling stage includes steps 400, 402, and 404. The encoding stage includes steps 406 and 408. More specifically, the compression schemes are utilized after the pixel (p) has been selected 402 and before or as part of the calculation of costs
15 404. The inverse scheme is unique and novel in its approach to estimating the pixel value. The inverse scheme scans the pixel array ($P[i][j]$) row by row using information from previously scanned pixels and the threshold array to predict the current pixel.

 The estimate of the gray value G24 (Figure 2A) and the threshold value are then used to make a prediction, with a probability, as to the value B24. This is derived from
20 the previous occurrences of the arithmetic difference of the gray estimate and the threshold value.

 To perform the estimate, a small set of pixels (S) is utilized, near to $P[i][j]$ that precedes $P[i][j]$ in scan order. Thus, referring to Figure 5, the first step is the selection of a pixel set 502. In the present example, the scan order is from left to right. Thus,
25 the pixels of Figure 2C would be scanned in the following order: P_{11} , P_{12} , P_{13} , P_{14} , P_{15} , P_{21} , P_{22} , ..., P_{45} , where P_{ij} represents i rows and j columns. Referring again to figure 2C, the S-set would likely consist of the values scanned prior to the relevant pixel and nearby. The relevant pixel is P_{24} due to the fact that it corresponds to the

location in the G-array being estimated (200). Thus, the S-set would contain: P₁₁ (0); P₁₂ (0); P₁₃ (1); P₁₄ (0); P₁₅ (0); P₂₁ (0); P₂₂ (0); and P₂₃ (1).

For each pixel P[i][j] in the S-set, an estimate of G[i][j] (the value of the pixel prior to thresholding) may be obtained by knowing what value t was used to threshold G[i][j]. In other words, an estimate of G₁₁ (of Figure 2A) may be obtained by knowing that the value 128 (the value at T₁₁ of Figure 2B) was used as the threshold to obtain a binary 0 result (the value at P₁₁ of Figure 2C). Similarly, an estimate of G₂₂ (of Figure 2A) may be obtained by knowing that the value 219 (the value at T₂₂ of figure 2B) was used as the threshold to obtain a binary 0 result (the value at P₂₂ of Figure 2C).

Referring to Figure 5, after selecting a set of neighboring pixels 502, a pixel that has not been previously selected (NP) is selected from the pixel set (S) 504. Referring to Figure 2A, a pixel such as G₂₃ would be selected to begin the estimation. If this is the first pixel to be selected 506, then a first interval level must be set. Setting the first interval level consists of first determining whether or not the binary thresholded value (BTV) of the pixel (NP) is equal to zero (0) or one(1) 508. If the binary thresholded value (BTV) is equal to zero (0), then the value of the pixel (NP) may be assumed to be lower than the threshold level 510. Conversely, if the binary thresholded value (BTV) is equal to one (1), then the value of the pixel (NP) may be assumed to be higher than or equal to the threshold level 512. Thus, assuming, as in the preferred embodiment, that the pixel values are made up of 8 bits (256 different values), the minimum value is that of zero (0) and the maximum value is that of 255.

The following equations set forth what the range of G[i][j] (NP) is depending on the resulting binary thresholded value (BTV):

25

$$\begin{aligned} t \leq G[i][j] \leq 255, & \quad \text{if the BTV of } P[i][j] = 1 \\ 0 \leq G[i][j] < t, & \quad \text{if the BTV of } P[i][j] = 0 \end{aligned}$$

If the binary thresholded value (BTV) at $P[i][j]$ is equal to 1, then it may be inferred that the value at $G[i][j]$ is greater than or equal to the t-value which was used in $G[i][j]$'s (NP's) thresholding. Similarly, if the binary thresholded value (BTV) at $P[i][j]$ is equal to 0, then it may be inferred that the value at $G[i][j]$ is less than the t-value which was
5 used in $G[i][j]$'s thresholding.

For example, if the binary thresholded value (BTV) at $P[i][j]$ was equal to a binary 1 and the threshold level ($T[i][j]$) is equal to 200, then it may be inferred that the value at $G[i][j]$ is equal to or greater than 200. Similarly, if $P[i][j]$ is equal to a binary 0 and the threshold value ($T[i][j]$) is equal to 200, then it may be inferred that
10 the value at $G[i][j]$ is less than 200. Referring to Figure 2C, the binary thresholded value (BTV) at P_{13} is 1. The value which was used as a threshold was 101 (the value at T_{13} of Figure 2B). As a result, it can be estimated that the value of G_{13} (NP) of Figure 2A is greater than or equal to 101. Similarly, the binary thresholded value of P_{22} (of Figure 2C) is 0. The value which was used as a threshold for P_{22} was 219
15 (the value at T_{22} of Figure 2B). Consequently, it may be inferred that the value at G_{22} (NP) is less than 219. Although the above examples of Figure 2 demonstrate that the values of G_{13} and G_{22} are already known (147 and 205 respectively), to utilize the inverse scheme of the present invention, estimates of all pixels in an S-set must be performed.

20 Based on the estimates of the surrounding pixels, the inverse scheme estimates the desired value at $G[i][j]$ (200). To make the estimation of X at $G[i][j]$ (200), the intervals for each of the pixels in S are intersected. Hence, referring again to Figure 5, subsequent to the first pixel 506, the intervals associated with the remaining pixels (NP) in the pixels set (S) are intersected. The intersections are performed sequentially in
25 order of proximity to pixel p. Thus, the intervals associated with the pixels from the S set are intersected one at a time beginning with the pixel closest to pixel p. Referring to Figure 2A, if pixel X 200 is the pixel to be estimated and the first pixel to be analyzed is that of G_{23} (with a value of 234), the next pixel that would have been scanned is that of

G22 (with a value of 205). The next pixel that would have been scanned is that of G21, then G15, etc.

Referring again to Figure 5, the first step of intersecting is that of determining the binary thresholded value of the pixel which is currently being analyzed/intersected

5 514. If the binary thresholded value is zero (0), then once again, it is known that the value prior to thresholding was less than the threshold value (T). Hence, the maximum value of the pixel (NP) would be that of one less than the threshold value. For example, the binary thresholded value of P22 (NP) (of Figure 2C) is 0. The value which was used as a threshold for P22 was 219 (the value at T22 of Figure 2B).

10 Consequently, it may be inferred that the value at G22 (NP) is less than 219 or the maximum value of NP is one less than 219 or 218. Hence, NP has a range of from 0 to 218.

Similarly, if the binary thresholded value is equal to one (1), then once again, it is known that the value prior to thresholding was greater than or equal to the threshold

15 value (T). Hence, the minimum value of the pixel (NP) would be that of the threshold value. For example, referring to Figure 2C, the binary thresholded value (BTV) at P13 is 1. The value which was used as a threshold was 101 (the value at T13 of Figure 2B). As a result, it can be estimated that the value of G13 (NP) of Figure 2A is greater than or equal to 101. Hence, the minimum value of the pixel (NP) would be that of

20 101.

If the intersection of the intervals as already determined does not intersect with the current NP's range, the process stops. As a result, it must be determined if an intersection will occur at all (steps 516 and 518 of Figure 5). If the existing interval level is greater than the threshold level, and the binary thresholded level was a zero (0),

25 then the ranges do not intersect 516. For example, if the BTV is zero (0), the existing interval level is from 210 to 255, and the threshold level is a 200, then the ranges 0-199 and 210-255 do not intersect. As such, the process stops without including the current NP's range in the intersection.

Similarly, if the existing interval is less than the threshold level, and the binary thresholded level was a one (1), then the ranges do not intersect 518. For example, if the BTV is one (1), the existing interval level is from 0 to 150, and the threshold level is a 200, then the ranges 0-150 and 200-255 do not intersect. As such the process stops.

- 5 On the other hand, assuming that an intersection will occur, then the existing interval level is updated, restricting the interval to the range of the intersection (Steps 520, 522, and 524 of Figure 5). For example, if the range for some pixel A (a pixel in closest proximity to pixel p) was greater than or equal to 200 (200-255), and the range for some other pixel B (the pixel in next closest proximity to pixel p) was less than 210 (0-10 210), the intersection of the two intervals is from 200 to 210.

- The intersection of the ranges in close proximity to pixel p then continues. Continuing with the above example, the intersection 200-210 would then be intersected with the next previous pixel in the scan. Referring to Figure 2, assuming that the S-set consists of the nearest 7 pixels as discussed supra, to estimate the value of X 200, the 15 intervals for the pixels in the S-set (those in closest proximity to X) are estimated and intersected. Thus, the first pixel to be estimated is that of G₂₃. To estimate what G₂₃ would be, its binary thresholded value (BTV) and its threshold value are examined. The binary thresholded value (BTV) of P₂₃ is 1 and the threshold value is 197. Thus, it may be estimated that G₂₃ must be greater than or equal to 197. Due to the fact that 20 the pixels in Figure 2 are made up of 8 bits, the maximum value is 255. Hence, the interval for G₂₃ is from 197-255.

- The next pixel in closest proximity is that of G₂₂. The binary thresholded value (BTV) of P₂₂ is 0 and the corresponding threshold level at T₂₂ is 219. Therefore, the estimated value for G₂₂ must be lower than 219. The intersection of the intervals from 25 G₂₃ and G₂₂ is then calculated. G₂₂ ranges from 0 to 218 and G₂₃ ranges from 197-255. Therefore, the intersection must be lower than 218 but higher than 197. Such an intersection results in an interval from 197-218. This process is then repeated with the next pixel in the closest proximity to X 200.

The intersection of the intervals continues until one of two conditions has been met. The first condition occurs when all of the intersects of pixels in the S-set have been intersected ("Condition A") (526). The second condition occurs when the intersection results in an empty set (i.e. the threshold intervals did not intersect) ("Condition B") (516 and 518), in which case the previous non-empty intersection is examined.

Continuing with the example above, the next closest pixel is that of G₂₁. The binary thresholded value (BTV) at P₂₁ is 0 and the threshold level at T₂₁ is 111. Thus, the value of G₂₁ must be less than 111. The intersection of this interval with the above interval (197-218) results in a set which is between 0 and 111 and also between 197-218. Due to the fact that no such value exists, the two intervals do not intersect resulting in a "contradiction". As a result, "Condition B" has been met (the threshold intervals did not intersect). The intersection of the intervals process comes to a halt. The last intersection which maintained a value is preserved (197-218).

From the resulting intersection interval, a value is estimated (g-estimate) and selected for G[i][j] (pixel p). Referring to Figure 6, if "Condition A" has been met 602A, then the midpoint of the resulting intersection is used as the g-estimate 604. Thus, if all pixels are intersected, condition A is met and the midpoint is used. In the above example, if no contradiction had resulted and all of the pixel intervals had been intersected, the midpoint between 197 and 218 (the value 208) would be used as the g-estimate. If "Condition B" has been met 602B, then the value closest to the contradiction is utilized (Steps 606, 608, and 610). Thus, if pixel intervals do not intersect, condition B is met and an endpoint is used.

In the above example, the resulting intersection interval is from 197-218. Further, the contradiction and halting arose with a value of 111 (Step 606A of Figure 6). Therefore, the value closest to the contradiction, yet still within the interval is that of 197, which is then used as the g-estimate 608. If the contradiction had occurred with a non intersecting interval of 230-255 (Step 606B of Figure 6), then the endpoint of

218 would be used as the g-estimate because it is the closest value to the contradiction (230) (Step 610 of Figure 6).

The first step in estimating the current pixel probability is obtaining/computing a difference value. The difference value (diff) is the difference between the g-estimate for
 5 pixel p and the threshold value (T) for pixel p. Thus, diff equals the g-estimate minus $T[i][j]$ where $T[i][j]$ is the threshold value for the pixel at $P[i][j]$:

$$\text{diff} = \text{g-estimate} - T[i][j]$$

10 Once again continuing with the above example, the value 197 is the g-estimate. Diff is therefore equal to $197 - T[i][j]$. $T[i][j]$ in the present example is equal to 214 (the value at T24 220). Thus, diff is equal to $197 - 214$ or -17.

The value of diff is used as the context for obtaining an estimated probability from a table comprising statistical history information for each context. For eight-bit g-
 15 estimates and threshold values, diff may retain values in the range: -255, -254, ..., 253, 254.

(g values in 0, 1, .. 255

t values in 1, ..., 255)

20

The current pixel bit value and the associated diff value context may be used to update the statistical information in the table.

In one embodiment, statistics are maintained on the approximate number of times previously for each diff value that $P[i][j]$ has resulted in a binary 0 after
 25 thresholding and the number of times $P[i][j]$ has resulted in a binary 1 after thresholding. Thus, a table indexed by the diff value context may be created which maintains approximate statistical information gathered over a number of past pixels.

For example, if for seven out of 10 times when the diff value was a -4, P_{12} was a 0, and for three out of those ten times P_{12} was a 1, such frequencies/statistics or approximations thereof would be recorded. Based on this frequency, $prob_1$ may be estimated. $Prob_1$ is equal to the estimated probability that the pixel is a binary 1.

- 5 Therefore, as in the above example, if for the last three out of ten times when the diff value was a -4, $prob_1$ would retain the value of approximately 3/10 or 0.3. $Prob_0$ is equal to the estimated probability that the pixel is a binary 0. Similarly, if for the seven of the last ten times when the diff value was a -4, P_{12} was a 0, $prob_0$ would retain the value of approximately 7/10 or 0.7. It should be noted that the total of the probabilities
- 10 is 1:

$$prob_0 + prob_1 = 1$$

$$7/10 + 3/10 = 1.$$

- 15 Thus, both probabilities need not be estimated in all embodiments, as one probability may be determined from the other.

As previously stated, the "frequency" table implementation for estimating probabilities is one example of how probability information may be determined for each context. Other methods for estimating the probability associated with each diff value

20 context may also be used within the scope of the invention.

CONTEXT COMPRESSION SCHEME

- A second scheme used in some embodiments of the invention is a variation of
- 25 the JBIG scheme described above. Referring to Figure 7, the pixels in S or a similar set are utilized (802). The binary thresholded values (BTV) are concatenated to provide a binary number which is then used as an index into a table of frequency counts (804). The estimated probability table is compiled from each previous occurrence of the index

(context) (806). This method is referred to as the context method. Hence, if the context is from 4 pixels, a table such as in Figure 3 is maintained and utilized to retain statistics of prob₁ or prob₀ in the context/JBIG scheme.

Referring to the table, the probability that the next bit after 1011 would be a 1 is examined by looking to the appropriate location of the table at 1011 signified by 300. Thus, the table indicates that 4/10 or 4 out of the last 10 times, the next bit was a 1. Similarly, the probability that the next bit after 1101 would be a 1 is examined by looking to the appropriate location of the table at 1101. Thus, the table indicates that 8/13 or 8 out of the last 13 times, the next bit was a 1. The table may also contain approximations rather than exact values. After each pixel is encoded, the table is updated to record the appropriate statistics.

SWITCHING METHOD

The generation of a context (e.g., a diff value or a JBIG-type context) and assigning of a probability to the context is the modeling portion of the compression. Thus, the inverse scheme and context scheme comprise the modeling portion of an embodiment of the invention. If the probabilities are accurate and close to 1, high compression is achieved. In other words, the more accurate the probabilities are and closer to 1 (or 10 out of 10 times; or 7/7 times...), the higher the compression.

In some embodiments, an entropy encoder may update the prediction and probability data in a feedback arrangement. Modeling may thus be performed during the encoding process by an adaptive encoder. For example, if a certain context and bit value are occurring with greater frequency, the estimated probability for that context may be adjusted upwards if the prediction matches the actual bit value, or downwards if the prediction does not match the actual bit value.

The portion of compression which involves conversion of the prediction, the estimated probability, and the actual bit value into encoded bits is called entropy

encoding. Entropy is a term of art which means or is a measure of how much information is encoded in a symbol. Thus, the higher the entropy of a message, the more information it contains. In this respect, if one uses entropy encoding, bits of data may be encoded using a number of bits corresponding to their information content.

- 5 Very predictable bits have a lower entropy, so require fewer bits to encode.

As the probability of a certain bit increases, the compression obtained using entropy encoding increases. Thus, the objective is to use a scheme which results in the highest probability for each pixel. Therefore, the best possible entropy encoding would use the highest probability as obtained from various different schemes (e.g., the inverse
10 scheme, the JBIG scheme, etc.). The switching method of the present invention allows for bit-level selective switching from scheme to scheme to utilize the highest probability.

The information regarding the pixels prior to $P[i][j]$ (and the threshold array) is referred to as the context:

15

$$P[i][j] = \begin{cases} \text{prediction} = 1, \text{ with estimated probability} = \text{prob}_1 \\ \text{prediction} = 0, \text{ with estimated probability} = 1 - \text{prob}_1 = \text{prob}_0 \end{cases}$$

In this respect, the probability that $P[i][j]$ is equal to 1 is prob_1 and the probability that $P[i][j]$ is equal to 0 is prob_0 . Thus, using the context scheme and referring to Figure 3,
20 the estimated probability, prob_1 , that $P[i][j]$ will have a binary value of 1 (300) is approximately 4/10 or 0.4 and the estimated probability, prob_0 , that $P[i][j]$ will have a binary value of 0 is approximately
1- prob_1 or $1 - 0.4 = 0.6$ or 6/10.

Using arithmetic encoding (a type of entropy encoding), the higher the
25 probability, the higher the compression capability. In order to determine the number of bits which are required to encode a set of binary numbers, the probability is used in a

logarithmic calculation. More specifically, the number of bits required to encode a binary number with a known probability is:

$$\begin{aligned} & \log_2 \frac{1}{\text{prob}_1} \text{ if the outcome } P[i][j] = 1 \text{ and} \\ 5 \quad & \log_2 \frac{1}{1 - \text{prob}_1} \text{ if the outcome } P[i][j] = 0. \end{aligned}$$

Thus, compression will take place according to the above statistics using arithmetic encoding.

If prob_1 is an accurate reflection of the true probability, the expected cost in
10 encoding bits to encode $P[i][j]$ is given by:

$$\begin{aligned} & - (\text{prob}_0) \log_2 \text{prob}_0 - \text{prob}_1 \log_2 \text{prob}_1 \\ = & - (1 - \text{prob}_1) \log_2 (1 - \text{prob}_1) - \text{prob}_1 \log_2 \text{prob}_1 \end{aligned}$$

15 To keep this cost low, the estimate of prob_1 needs to be accurate and as close to 0 or 1 as possible. As the estimate of prob_1 moves away from 0.5, the number of bits required to encode the respective pixel bit decreases. This results due to the fact that if prob_1 is near 0 (i.e., prob_0 is near 1), then the number of bits required to encode that bit is close to 0. Similarly, if prob_1 is near 1, then the number of bits required to encode
20 that bit is also close to 0. Thus, the closer the probability is to 0 or 1, the lower the cost of encoding is.

The probability prob_1 is a conditional probability, conditioned on the context used. In the embodiment described herein, two different contexts (schemes) are utilized and whichever context (scheme) results in the lowest total cost over a set of previously
25 encoded pixels is the scheme used for encoding the current bit. In other embodiments, more than two schemes may be cost-evaluated for encoding each pixel bit.

At each previously scanned pixel, the cost to encode that pixel is computed as:

- $\log_2 \text{prob}_1^{\text{inverse}}$ for the inverse scheme
and - $\log_2 \text{prob}_1^{\text{context}}$ for the context scheme

- 5 if the pixel was a one. Similarly, - $\log_2 \text{prob}_0^{\text{inverse}}$ and - $\log_2 \text{prob}_0^{\text{context}}$ if the pixel was a zero, when $\text{prob}_1^{\text{inverse}}$ and $\text{prob}_1^{\text{context}}$ are the probability estimates from the inverse and context schemes.

The next step is to add up the cost for the set of previously encoded pixels, such as S, for each scheme. The scheme with the lower sum is used to predict $P[i][j]$. In
10 other words, the scheme with the lower summed cost is utilized for the encoding. Thus, if the summed cost of the inverse scheme over the pixel set S is 140 bits and the summed cost of the context scheme over the pixel set S is 165 bits, then the inverse scheme would be utilized for the encoding. The summed cost is recalculated for each new pixel scanned in the $P[i][j]$ array. Hence, the scheme used may be switched from
15 pixel to pixel depending on the cost associated with each scheme.

Figure 13A is a functional block diagram illustrating a switching embodiment of the invention. In Figure 13A, scanned data is stored in registers (e.g., a shift register) or memory cells 1500 for use in encoding the data. The neighboring pixel bit values 1505 of the set S from registers 1500 are provided to the multi-scheme context/index
20 generator 1501. Contexts or indices 1506 are generated for each scheme using methods such as those described for the inverse and JBIG schemes.

Contexts and indices (such as diff and the JBIG context) 1506 are presented to the statistics/probability tables 1502 to access the respective stored predictions and probabilities 1507 for each scheme. The prediction/probability values 1507 are
25 provided to the cost analysis switcher/selector 1503 for determination of the current encoding scheme, based on the cost analysis described above. The prediction/probability 1508 for the selected scheme is then provided to entropy encoding engine 1504. The current pixel bit value 1509 from registers 1500 is also

provided to entropy encoding engine 1504, and based on the value of prediction/probability 1508, entropy encoding engine 1504 produces encoded data 1510.

The functional blocks of Figure 13A may be embodied in dedicated electronic hardware, or they may be embodied in software functions implemented in general electronic hardware, or they may be embodied in a combination of dedicated hardware and software functions.

HYBRID CONTEXT METHOD

10

A second embodiment of the invention utilizes multiple compression schemes to form a hybrid context for each pixel. By using a hybrid context, the embodiment is able to take advantage of the unique statistical aspects of each compression scheme in determining the probability for each possible hybrid context. The resulting hybrid context can provide improved prediction ability (i.e., higher probabilities) with correspondingly higher compression ratios.

The hybrid context is formed by assigning a portion of bits in the context to each compression scheme. The size of the total context and the apportionment of the bits between the respective schemes may vary for different embodiments. For the purpose of clarity, a specific embodiment having a two-scheme thirteen-bit context is described.

In this embodiment, the thirteen-bit context is separated into a seven-bit portion assigned to the inverse scheme, and a six-bit portion assigned to an abbreviated JBIG scheme. This apportionment is illustrated in Figure 8. As shown, the six least significant bits (B6-B1) of the thirteen-bit context store the JBIG context information, and the seven most significant bits (D6-D0) store the inverse scheme information. Similar embodiments may have a different ordering of this context information.

The set of previous pixels S' is a reduced set from S described earlier. The set S' consists of six previous pixel values as shown in Figure 11. B0 represents the current pixel, whereas B1-B6 represent the six nearest neighbors, with B1 being the nearest neighbor, B2 being the second nearest neighbor, etc. The values of B1-B6 provide the abbreviated JBIG context portion for the hybrid context.

With respect to the inverse scheme context portion, the "gray value" estimation is performed as described previously for the inverse scheme by progressing through the respective ranges of the previous pixel values, preferably in the order of B1, B2, B3, etc., to maintain the priority of the nearest pixels since the nearest pixels contain the highest predictive information. In other embodiments, to reduce the amount of calculation per pixel, the "gray value" range for several pixels may be precalculated by column and stored to render repeated calculation of the intersections unnecessary. The calculation reduction is achieved with some compromise in priority of the pixel information.

The resulting difference value (diff) between the gray value estimate and the current pixel threshold is typically within the range of -255 to 254, e.g., for eight-bit pixel data. In order to reduce this difference value into the prescribed seven bits, the difference value is remapped into the range of 0 to 127. One such mapping is illustrated in Figure 10.

In Figure 10, the difference value (diff) is mapped into the seven bit difference value (diff') for generation of the inverse scheme portion of the hybrid context as follows:

$$\text{diff}' = \begin{cases} 0, & \text{for diff} < -64 \\ \text{diff} + 64, & \text{for } -64 \leq \text{diff} \leq 63 \\ 127, & \text{for diff} > 63 \end{cases}$$

The above distribution for diff assures better probability resolution for gray value estimates near the threshold value at the cost of clamping diff for gray value estimates further away from the threshold, where the probability is less likely to vary substantially for different diff values.

5 With the hybrid context formed as described, statistical analysis of a set of past data may be performed to establish a probability table indexed by specific context values for use in the entropy encoding. As described previously, this may be done by tracking the number of occurrences of each particular context, and the number of times each particular context resulted in a "1" or a "0." The resulting probability for a particular
10 context to yield a "1," for example, would then be an approximate count of times that particular context produced a "1" divided by the overall approximate count for occurrences of that particular context. Also, as mentioned previously, an adaptive entropy encoder may update statistics during encoding to maintain probability data, or some combination of statistical analysis and adaptive encoding may be used.

15 Figure 9 is a flow diagram for the encoding process using the hybrid context. In step 1100, the six-bit JBIG context is determined from six previous pixels. In step 1101, the seven-bit inverse scheme context (diff) is determined. The relative order in which steps 1100 and 1101 are performed is not critical to the operation of the embodiment. In step 1102, the two context portions are combined into a single hybrid
20 context, and, in subsequent step 1103, the context is used as an index into the statistics table to obtain the prediction and estimated probability for that context. In step 1104, the prediction, estimated probability and actual current bit value are provided to an entropy encoder for encoding of the current bit.

 On the decoding side, an entropy decoder will receive the encoded bit. By
25 performing similar operations to steps 1100-1104, the hybrid context is formed and applied to the same or a similar statistics table to obtain the prediction and probability. The entropy decoder is thus able to decode the current bit value.

Figure 13B is a functional block diagram illustrating a hybrid context embodiment of the invention. In Figure 13B, scanned data is stored in registers (e.g., a shift register) or memory cells 1500 for use in encoding the data. The neighboring pixel bit values 1505 (e.g., from the set S or S') from registers 1500 are provided to the multi-scheme context/index generator 1501. Contexts or indices 1506 are generated for each scheme using methods such as those described for the inverse and JBIG schemes.

Contexts and indices (such as diff or diff' and the JBIG context) 1506 are provided to hybrid context generator 1511 to be combined into hybrid context 1512. Hybrid context 1512 is presented to the statistics/probability table 1502 to access the respective stored prediction and probability 1508 for the hybrid context scheme. The prediction/probability 1508 for the current hybrid context is then provided to entropy encoding engine 1504. The current pixel bit value 1509 from registers 1500 is also provided to engine 1504, and based on the value of prediction/probability 1508, engine 1504 produces encoded data 1510.

The functional blocks of Figure 13B may be embodied in dedicated electronic hardware, or they may be embodied in software functions implemented in general electronic hardware, or they may be embodied in a combination of dedicated hardware and software functions.

ENTROPY/ARITHMETIC ENCODING

The present invention permits the selective switching of compression schemes based on cost analysis, or the combining of compression schemes in a hybrid context. The objective is to represent the pixel array in fewer bits than n rows by m columns.

The encoding stage of the invention utilizes the probabilities obtained in the modeling stage to encode the bits such that they are represented in fewer bits. This is performed through entropy encoding. In an embodiment of the present invention, the method of arithmetic encoding is utilized to perform entropy encoding.

The above description has focused on the compression of image data compressed (thresholded) into one-bit data which is then further compressed by the methods of the invention. It will be obvious to those skilled in the art that the methods of the invention may be similarly applied to multi-bit data. For example, data may be
5 thresholded as shown in Figure 12 such that two-bit data is the result. Rather than one threshold value for each pixel, three threshold values are provided for each pixel to divide the initial pixel value range into four sections, each section represented by a two bit value (00, 01, 10, 11).

Multi-bit schemes may be encoded on a bit by bit basis. Bitwise encoding may
10 be performed by designating bit-planes for the data, such as a most significant bit-plane and least significant bit-plane. Each bit plane is then encoded using the bi-level data schemes described above, though lesser bit-planes may take advantage of the predictive information in the more significant bits by using nearest neighbor sets including bits in adjacent bit-planes.

15 Magnitude code representations may be used to permit encoding of subsequent bits based on the encoding of the previously encoded bits of a pixel. For example, if the three-bit magnitude code values 000, 001, 011 and 111 are assigned to the two-bit pixel values 00, 01, 10 and 11, then the second bit need be encoded only if the most significant bit is a zero, and likewise, the least significant bit need be encoded only if
20 the second bit is zero. Thus, multi-bit compression may be performed using the methods of the present invention.

Thus a method and apparatus for data compression has been described in conjunction with one or more specific embodiments. The invention is defined by the claims and their full scope of equivalents.

CLAIMS

1. A scheme for predicting a gray value of a frame buffer cell comprising the steps of:
- 5 selecting a pixel-set (S) from previously scanned pixels in a pixel array;
intersecting intervals from pixels in the pixel-set, said intersecting process comprising the steps of:
- determining an interval (I) from a first pixel in the pixel-set;
determining a pixel-interval (PI) from another pixel in the pixel-set;
- 10 intersecting said interval (I) with said pixel interval (PI) to create a new interval (NI)
if NI is not an empty set, adjusting I to equal NI; and
selecting a gray value from the interval (I).
- 15 2. The scheme of claim 1 wherein said pixel-set (S) is constructed from pixels closest in proximity to the pixel of the frame buffer cell being predicted.
3. The scheme of claim 1 wherein said step of intersecting intervals is conducted in order of proximity to the pixel of the frame buffer cell being predicted.
- 20 4. The scheme of claim 1 wherein said pixel interval (PI) is greater than or equal to a threshold value if a corresponding binary thresholded pixel value is 1.
5. The scheme of claim 1 wherein said pixel interval (PI) is less than a
25 threshold value if a corresponding binary thresholded pixel value is 0.
6. The scheme of claim 1 wherein said gray value is a midpoint of the interval (I) if NI does not contain the empty set.



7. The scheme of claim 1 wherein said gray value is an endpoint of the interval (I) if NI is an empty set.

5 8. The scheme of claim 7 wherein said endpoint is the endpoint closest to the pixel-interval (PI).

9. The scheme of claim 1 further comprising the step of:
comparing said gray value to a threshold matrix value.

10

10. The scheme of claim 9 further comprising the steps of:
forming a difference value from said comparing step; and
utilizing said difference value as a pixel context.

15 11. The scheme of claim 10 wherein said step of utilizing said difference value comprises the steps of:
accessing statistical information indexed by said difference value; and
performing entropy encoding using said statistical information.

20 12. A method for compressing data in a computer system comprising the steps of:

providing a plurality of modeling schemes;

calculating an estimated cost to encode a previous set of pixels for each modeling scheme;

25 selecting a lowest cost scheme from said plurality of modeling schemes; and
encoding the pixel based on said lowest cost scheme.

13. The method of claim 12 wherein two modeling schemes are provided.



14. The method of claim 12 wherein one of the modeling schemes is a context scheme.

5 15. The method of claim 12 wherein one of the modeling schemes is an inverse scheme.

16. The method of claim 15 wherein said calculating step comprises the steps of:

10 calculating a probability that a pixel bit value for each pixel in said previous set is 1; and

determining a number of bits based on said probability.

17. The method of claim 16 wherein for each pixel P_{ij} in said previous set, 15 the encoding cost for the inverse scheme is approximately equal to:

$-\log \text{prob}_{1^{ijj}, \text{inverse}}$ if said pixel in a 1

else $-\log \text{prob}_{0^{ijj}, \text{inverse}}$ if said pixel in a 0

where $\text{prob}_{1^{ij}, \text{inverse}}$ is the estimate derived from previous occurrences of P_{ij} 's context; where $\text{prob}_{1^{ij}, \text{inverse}} = 1 - \text{prob}_{0^{ij}, \text{inverse}}$.

20

18. The method of claim 14 wherein said calculating step comprises the steps of:

calculating a probability that a pixel bit value is 0; and

determining a number of bits based on said probability.

25



19. The method of claim 18 wherein for each pixel P_{ij} in said previous set, the encoding cost for the context scheme is approximately equal to:

$-\log \text{prob}_{1^{ij}, \text{context}}$ if said pixel in a 1

else $-\log \text{prob}_{0^{ij}, \text{context}}$ if said pixel in a 0

5 where $\text{prob}_{1^{i,j}, \text{inverse}}$ is the estimate derived from previous occurrences of P_{ij} 's context; where $\text{prob}_{1^{i,j}, \text{inverse}} = 1 - \text{prob}_{0^{i,j}, \text{inverse}}$.

20. The method of claim 12 wherein said encoding step comprises arithmetic encoding.

10

21. The method of claim 13, wherein said total cost for all pixels in said previous set is the sum of the cost for the individual pixels, where there is a cost for the context scheme and for the inverse scheme and where the scheme having the lower cost is used to encode a current pixel.

15

22. An apparatus for predicting a gray value of a frame buffer cell comprising:

means for intersecting intervals from pixels in a pixel-set (S) selected from previously scanned pixels in a pixel array, said intersecting means comprising:

20 means for determining a resulting interval (RI) from a first pixel in the pixel-set;

means for determining a pixel-interval (PI) from another pixel in the pixel-set;

25 means for intersecting said resulting interval (RI) with said pixel interval (PI) to create a new interval (NI)

means for adjusting RI to equal NI, if NI is not an empty set; and

means for selecting a gray value from the resulting interval (RI).



23. The apparatus of claim 22 wherein said pixel-set (S) comprises pixels closest in proximity to the pixel of the frame buffer cell being predicted.

24. The apparatus of claim 22 wherein said means for intersecting intervals
5 comprise means to intersect in order of proximity to the pixel of the frame buffer cell being predicted.

25. The apparatus of claim 22 wherein said pixel interval (PI) is greater than or equal to a threshold value if a corresponding bi-level pixel value is 1.
10

26. The apparatus of claim 22 wherein said pixel interval (PI) is less than a threshold value if a corresponding bi-level pixel value is 0.

27. The apparatus of claim 22 wherein said gray value is a midpoint of the
15 resulting interval (RI) if NI is not an empty set.

28. The apparatus of claim 22 wherein said gray value is an endpoint of the resulting interval (RI) if NI is an empty set.

20 29. The apparatus of claim 28 wherein said endpoint is the endpoint closest to the pixel-interval (PI).

30. The apparatus of claim 22 further comprising:
means for comparing said gray value to a threshold matrix value.

25

31. The apparatus of claim 30 further comprising:
means for applying a difference value generated from said comparing means as a pixel context.



32. The apparatus of claim 31 wherein said applying means comprises:
means for accessing statistical information indexed by said difference value; and
an entropy encoder receiving said statistical information.

5

33. An apparatus for compressing data in a computer system comprising:
means for calculating an estimated cost to compress a previous set of pixels for
each one of a plurality of modeling schemes;
means for selecting a lowest cost scheme; and
10 means for encoding the pixel based on a selected scheme.

34. The apparatus of claim 33 wherein said plurality of modeling schemes
comprise two modeling schemes.

15 35. The apparatus of claim 33 wherein one of the modeling schemes is a
context scheme.

36. The apparatus of claim 33 wherein one of the modeling schemes is an
inverse scheme.

20

37. The apparatus of claim 33 wherein said means for calculating further
comprise:
means for calculating a probability that a bit value is 1; and
means for determining a number of bits based on said probability.

25



38. The apparatus of claim 37 wherein for each pixel P_{ij} in said previous set, the encoding cost for the inverse scheme is approximately equal to:

-log prob $_{1^{i,j, inverse}}$ if said pixel in a 1

else -log prob $_{0^{i,j, inverse}}$ if said pixel in a 0

5 where prob $_{1^{i,j, inverse}}$ is the estimate derived from previous occurrences of P_{ij} 's context; where prob $_{1^{i,j, inverse}} = 1 - \text{prob}_{0^{i,j, inverse}}$.

39. The apparatus of claim 36 wherein said means for calculating further comprises:

10 means for calculating a probability that a bit value is 0; and
 means for determining a number of bits based on said probability.

40. The apparatus of claim 39 wherein said number of bits is approximately equal to:

15 $\log_2 \frac{1}{\text{prob}_0}$, wherein prob $_0$ is said probability.

41. The apparatus of claim 33 wherein said means for encoding comprise an arithmetic encoder.

20 42. A method for compressing data in a computer system, comprising the steps of:

generating a first context based on a first compression scheme;

generating a second context based on a second compression scheme;

forming a hybrid context comprising said first context and said second context;

25 accessing statistical information indexed by said hybrid context; and

performing entropy encoding using said statistical information.



43. The method of claim 42 wherein said step of generating a first context comprises the steps of:

estimating a gray value for a pixel; and

determining a difference between said gray value and a current pixel threshold

5 value.

44. The method of claim 43 wherein said step of estimating said gray value comprises the steps of:

determining a first value range for a first adjacent pixel;

10 determining a second value range for a second adjacent pixel; and

determining an intersection range of said first value range and said second value range.

45. The method of claim 43 wherein said difference comprises a first
15 number of bits, said method further comprising the step of:
mapping said difference into a second number of bits smaller than said first
number of bits.

46. The method of claim 42 wherein said step of generating a second context
20 comprises the step of:

determining a set of neighboring pixels; and

concatenating a plurality of bit values of said neighboring pixels to form said
second context.

25 47. The method of claim 42 wherein said statistical information comprises a
prediction and a probability.



48. The method of claim 47 further comprising the steps of:
counting occurrences of each hybrid context value in a portion of data; and
determining said prediction and said probability from said counting.

5 49. The method of claim 46 wherein said step of performing entropy
encoding comprises the step of:
performing arithmetic encoding based on said prediction and said probability.

10 50. An apparatus for compressing data in a computer system comprising:
a plurality of pixel values stored in memory;
a context generator receiving a set of pixel values from said memory, said set of
pixel values comprising pixel values adjacent a current pixel value, said generator
providing a plurality of contexts associated with a plurality of compression schemes;
tables indexed by said plurality of contexts, said tables having statistical
15 information;
cost-based selector receiving said statistical information from said tables and
providing statistical information associated with a lowest cost compression scheme; and
an encoding engine receiving said current pixel value from said memory and
said statistical information from said cost-based selector, said encoding engine
20 providing encoded output.

51. The apparatus of claim 50 wherein said context generator comprises:
means for determining an inverse scheme difference value; and
means for concatenating previous pixel values.

25 52. The apparatus of claim 50 wherein said encoding engine comprises:
an arithmetic encoder.



53. The apparatus of claim 50 wherein said statistical information comprises a prediction and a probability value associated with each context.

54. The apparatus of claim 50 wherein said cost-based selector comprises:
5 means for calculating a cost based on a probability value; and
means for selecting a lowest cost scheme.

55. An apparatus for compressing data in a computer system comprising:
a plurality of pixel values stored in memory;
10 a context generator receiving a set of pixel values from said memory, said set of pixel values comprising pixel values adjacent a current pixel value, said generator providing a plurality of contexts associated with a plurality of compression schemes;
a hybrid context generator receiving said plurality of contexts and generating a hybrid context;
15 a table indexed by said hybrid context, said table having statistical information;
and
an encoding engine receiving said current pixel value from said memory and said statistical information from said table, said encoding engine providing encoded output.

20

56. The apparatus of claim 55 wherein said context generator comprises:
means for determining an inverse scheme difference value; and
means for concatenating previous pixel values.

25 57. The apparatus of claim 55 wherein said encoding engine comprises:
an arithmetic encoder.



58. The apparatus of claim 55 wherein said statistical information comprises a prediction and a probability value associated with each context.

59. The apparatus of claim 55 wherein said hybrid context generator
5 comprises:
means for concatenating said plurality of contexts.

60. The apparatus of claim 59 wherein said hybrid context generator further
comprises:
10 means for mapping a context having a first number of bits into a second number
of bits, said second number being less than said first number.

61. A computer program product comprising:
a computer usable medium having computer readable program code embodied
15 therein for performing data compression operations in a computer system, said
computer program product comprising:

computer readable program code configured to cause a computer to generate a
first context based on a first compression scheme;

20 computer readable program code configured to cause a computer to generate a
second context based on a second compression scheme;

computer readable program code configured to cause a computer to form a
hybrid context comprising said first context and said second context;

computer readable program code configured to cause a computer to access
statistical information indexed by said hybrid context; and

25 computer readable program code configured to cause a computer to perform
entropy encoding using said statistical information.



62. The computer program product of claim 61 wherein said computer readable program code configured to cause a computer to generate a first context comprises:

- computer readable program code configured to cause a computer to estimate a
5 gray value for a pixel; and
- computer readable program code configured to cause a computer to determine a difference between said gray value and a current pixel threshold value.

63. The computer program product of claim 62 wherein said computer
10 readable program code configured to cause a computer to estimate a gray value comprises:

- computer readable program code configured to cause a computer to determine a first value range for a first adjacent pixel;
- computer readable program code configured to cause a computer to determine a
15 second value range for a second adjacent pixel; and
- computer readable program code configured to cause a computer to determine an intersection range of said first value range and said second value range.

64. The computer program product of claim 62 wherein said difference
20 comprises a first number of bits, said computer readable program code configured to cause a computer to generate a first context further comprising:

- computer readable program code configured to cause a computer to map said difference into a second number of bits smaller than said first number of bits.



65. The computer program product of claim 61 wherein said computer readable program code configured to cause a computer to generate a second context comprises:

computer readable program code configured to cause a computer to determine a
5 set of neighboring pixels; and

computer readable program code configured to cause a computer to concatenate a plurality of bit values of said neighboring pixels to form said second context.

66. The computer program product of claim 61 wherein said statistical
10 information comprises a prediction and a probability.

67. The computer program product of claim 66 wherein said computer usable medium further comprises:

computer readable program code configured to cause a computer to count
15 occurrences of each hybrid context value in a portion of data; and

computer readable program code configured to cause a computer to determine said prediction and said probability from said count of occurrences.

68. The computer program product of claim 67 wherein said computer
20 readable program code configured to cause a computer to perform entropy encoding comprises:

computer readable program code configured to cause a computer to perform arithmetic encoding based on said prediction and said probability.

