



US006941019B1

(12) **United States Patent**
Mitchell et al.

(10) **Patent No.:** **US 6,941,019 B1**

(45) **Date of Patent:** **Sep. 6, 2005**

(54) **REENTRY INTO COMPRESSED DATA**

(75) Inventors: **Joan LaVerne Mitchell**, Longmont, CO (US); **Nenad Rijavec**, Longmont, CO (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/569,117**

(22) Filed: **May 10, 2000**

(51) **Int. Cl.**⁷ **G06K 9/36**

(52) **U.S. Cl.** **382/232; 382/235**

(58) **Field of Search** **382/232-253, 382/297; 341/51, 107, 65; 348/390.1, 408.1; 375/240.01-240.26; 358/426.01-426.16**

(56) **References Cited**

U.S. PATENT DOCUMENTS

| | | |
|-------------|---------|-----------------|
| 4,596,039 A | 6/1986 | Mitchell et al. |
| 4,610,027 A | 9/1986 | Anderson et al. |
| 4,627,020 A | 12/1986 | Anderson et al. |
| 4,646,356 A | 2/1987 | Anderson et al. |
| 4,658,430 A | 4/1987 | Anderson et al. |
| 4,689,824 A | 8/1987 | Mitchell et al. |
| 4,725,815 A | 2/1988 | Mitchell et al. |
| 4,837,845 A | 6/1989 | Pruett et al. |
| 4,888,645 A | 12/1989 | Mitchell et al. |
| 4,891,643 A | 1/1990 | Mitchell et al. |

(Continued)

FOREIGN PATENT DOCUMENTS

| | | | | |
|----|--------------|---------|-------|-----------|
| EP | 1037165 A2 * | 9/2000 | | G06T 3/00 |
| JP | 02698037 | 12/1992 | | |
| WO | WO96/36941 | 11/1996 | | |

OTHER PUBLICATIONS

Glen G. Langdon, Jr., "An Introduction to Arithmetic Coding," *IBM J. Res. Develop.* 28, 135-149 (Mar. 1984).

(Continued)

Primary Examiner—Amelia M. Au

Assistant Examiner—Ishrat Sherali

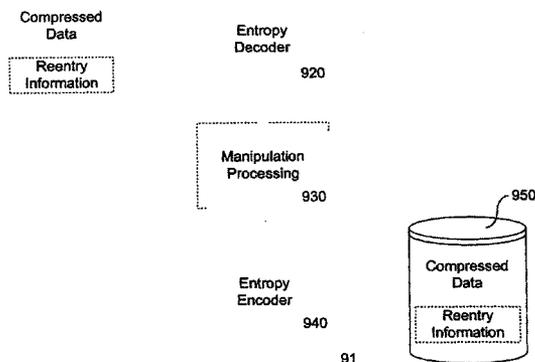
(74) *Attorney, Agent, or Firm*—Blakely, Sokoloff, Taylor & Zafman LLP

(57) **ABSTRACT**

Apparatus and methods are provided for entering compressed data streams at selected reentry points to initiate decoding thereby allowing efficient manipulation of the compressed data and minimizing storage requirements. The reentry information preferably includes bit-level pointers and sufficient state information to initialize the decoder properly. This enables decoding without having to resume at independently decodable points, such as JPEG restart markers. For example, in the context of a JPEG image, in addition to the typical information available to the decoder that has been passed in earlier markers, the reentry information for a given MCU boundary may include: a bit-level pointer to the first block's DC Huffman code, the position of the output, and a DC predictor for each component of the MCU. This allows decompression to be performed in the appropriate order to accomplish various data manipulation operations, such as rotation, thus significantly reducing buffering requirements. Reentry information into a compressed data stream can be generated during encoding, decoding, partial encoding, partial decoding, entropy encoding, and/or entropy decoding. In addition, a reentry decoder may quickly interpret the compressed data sufficiently to preserve desired reentry information and discard unneeded output of the decoding process and terminate immediately after the last desired reentry point. This enables buffering of pieces of compressed data with associated reentry information rather than buffering the entire decompressed data. Additionally, when a subset of the reconstructed data is needed the step of recompressing the individual pieces can be avoided by saving reentry information with the associated pieces of compressed data.

60 Claims, 21 Drawing Sheets

900



U.S. PATENT DOCUMENTS

| | | | | |
|-----------|------|---------|------------------------|---------|
| 4,905,297 | A | 2/1990 | Langdon, Jr. et al. | |
| 5,327,248 | A * | 7/1994 | Miller et al. | 382/233 |
| 5,341,440 | A * | 8/1994 | Earl et al. | 382/233 |
| 5,708,732 | A | 1/1998 | Merhav et al. | |
| 5,751,865 | A | 5/1998 | Micco et al. | |
| 5,867,598 | A | 2/1999 | de Queiroz | |
| 6,081,211 | A * | 6/2000 | de Queiroz et al. | 341/65 |
| 6,246,801 | B1 * | 6/2001 | Boyce et al. | 382/248 |
| 6,298,166 | B1 * | 10/2001 | Ratnakar et al. | 382/248 |
| 6,381,371 | B1 * | 4/2002 | Epstein et al. | 382/246 |
| 6,690,832 | B1 * | 2/2004 | Mitchell et al. | 382/233 |

OTHER PUBLICATIONS

Joan L. Mitchell, Ph.D., William Pennebaker, Ph.D., "JPEG", 1993 *AIMM Show and Conference—Imaging: Transformation at Work*, Chicago, Illinois (Apr. 5-8, 1993).
 K.L. Anderson, J.L. Mitchell, G. Goertzel, W.B. Pennebaker, "Compression of Bilevel Images."
 P. J. Sementilli, "Image Processing in the Quantized DCT Domain," *IBM Confidential* (Nov. 11, 1992).
 Joan L. Mitchell, "Facsimile Image Coding," *AFIPS—Conference Proceedings, National Computer Conference*, vol. 49, 423-426 (1980).
 W.B. Pennebaker, G. Goertzel, J.L. Mitchell, "Recent Advances in Data-Conversion Facsimile Techniques," *Journ. Appl. Photog. Eng.*, vol. 6, No. 4, 93-96 (Aug. 1980).
 S.T. Klein, Y. Wiseman, "Parallel Huffman Decoding," *IEEE Computer Society Proceedings, Data Compression Conference*, Snowbird, Utah (Mar. 28-30, 2000).
 W.B. Pennebaker, J.L. Mitchell, "JPEG Still Image Data Compression Standard," Chapman & Hall, International Thomson Publishing (1993), ISBN 0-442-01272-1, pp. 73-78.
 W.B. Pennebaker, J.L. Mitchell, "JPEG Still Image Data Compression Standard," Chapman & Hall, International Thomson Publishing (1993), ISBN 0-442-01272-1, Chapter 7, JPEG Syntax and Data Organization, pp. 97-120.
 W.B. Pennebaker, J.L. Mitchell, "JPEG Still Image Data Compression Standard," Chapman & Hall, International Thomson Publishing (1993), ISBN 0-442-01272-1, Chapter 8, Entropy Coding Concepts, pp. 135-147.
 W.B. Pennebaker, J.L. Mitchell, "JPEG Still Image Data Compression Standard," Chapman & Hall, International Thomson Publishing (1993), ISBN 0-442-01272-1, Chapter 9, JPEG Binary Arithmetic Coding, pp. 149-167.
 M. Rabbani, P.W. Jones, "Digital Image Compression

Techniques," *Tutorial Texts in Optical Engineering*, SPIE Optical Engineering Press (1991), ISBN 0-8194-0648-1, vol. TT7, Chapter 3, Variable-Length Codes, pp. 22-32.
 "Algorithm for Rotating an Image in Run End Form," *IBM Tech. Disclosure Bulletin* (Nov. 1989), pp. 299-302.
 K.L. Anderson, F.C. Mintzer, G. Goertzel, J.L. Mitchell, K.S. Pennington, W.B. Pennebaker, "Binary-Image-Manipulation Algorithms in the Image View Facility," *IBM J. Res. Develop.* (Jan. 1987), vol. 31, No. 1, pp. 16-31.
 J.L. Mitchell, W.B. Pennebaker, "Software Implementations of the Q-Coder," *IBM J. Res. Develop.* (Nov. 1988), vol. 32, No. 6, pp. 753-774.
 W.B. Pennebaker, J.L. Mitchell, "Probability Estimation for the Q-Coder," *IBM J. Res. Develop.* (Nov. 1988), vol. 32, No. 6, pp. 737-752.
 J.L. Mitchell, W.B. Pennebaker, "Optimal Hardware and Software Arithmetic Coding Procedures for the Q-Coder," *IBM J. Res. Develop.* (Nov. 1988), vol. 32, No. 6, pp. 727-736.
 W.B. Pennebaker, J.L. Mitchell, G.G. Langdon, Jr., R.B. Arps, "An Overview of the Basic Principles of the Q-Coder Adaptive Binary Arithmetic Coder," *IBM J. Res. Develop.* (Nov. 1988), vol. 32, No. 6, pp. 717-726.
 R.B. Arps, T.K. Truong, D.J. Lu, R.C. Pasco, T.D. Friedman, "A Multi-Purpose VLSI Chip for Adaptive Data Compression of Bilevel Images," *IBM J. Res. Develop.* (Nov. 1988), vol. 32, No. 6, pp. 775-795.
 M.J. Slattery, J.L. Mitchell, "The Qx-coder," *IBM J. Res. Develop.* (Nov. 1998), vol. 42, No. 6, pp. 767-784.
 T.C. Bell, J.G. Cleary, I.H. Witten, "Text Compression," Prentice Hall PTR (1990), ISBN 0-13-911991-4, Chapter 8, Dictionary Techniques, pp. 206-243.
 J. Banks, "Integrated Circuit Cards (ICCs) and Optical Memory Cards (OMCs) Hold a Lot in Store," *ISO Bulletin* (Nov. 1993), vol. 24, No. 11, pp. 9-11.
 R. Hunter, A.H. Robinson, "International Digital Facsimile Coding Standards," (Invited Paper) *Proceedings of the IEEE, Special Issue on digital encoding of graphics* (Jul. 1980), vol. 68, No. 7, pp. 854-867.
 D.A. Huffman, "A Method for the Construction of Minimum-Redundancy Codes," *Proceedings of the I.R.E.* (Sep. 1952), vol. 40, No. 9, pp. 1098-1101.
 W.B. Pennebaker, J.L. Mitchell, "A Fast-Moving International Standard for Still Image Compression," *ISO Bulletin* (Nov. 1993), vol. 24, No. 11, pp. 2-8.

* cited by examiner

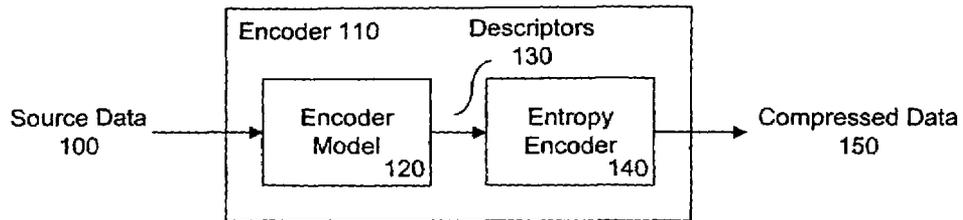


Figure 1 (Prior Art)

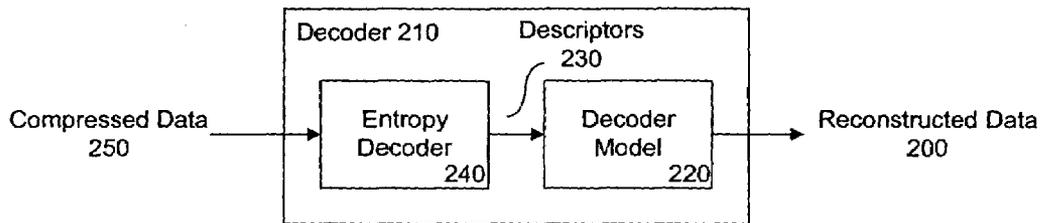


Figure 2 (Prior Art)

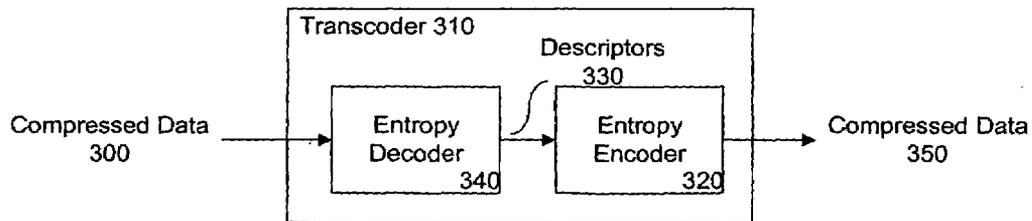


Figure 3 (Prior Art)

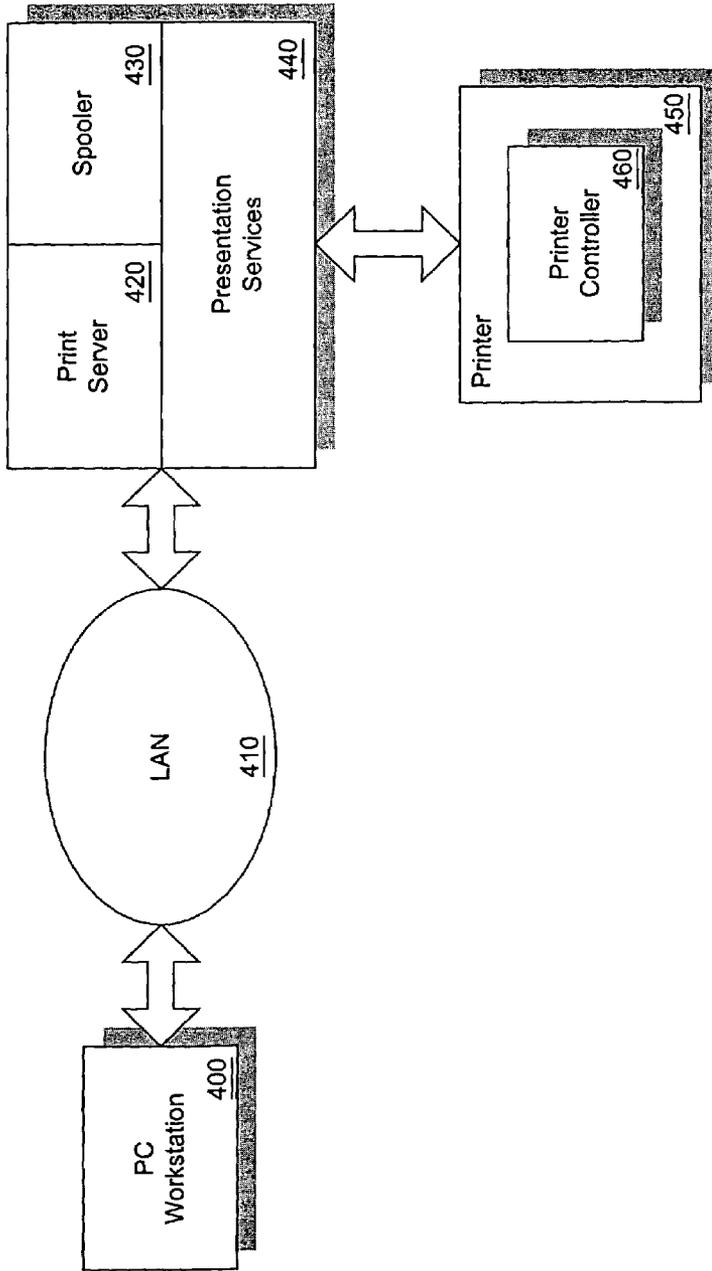


Figure 4

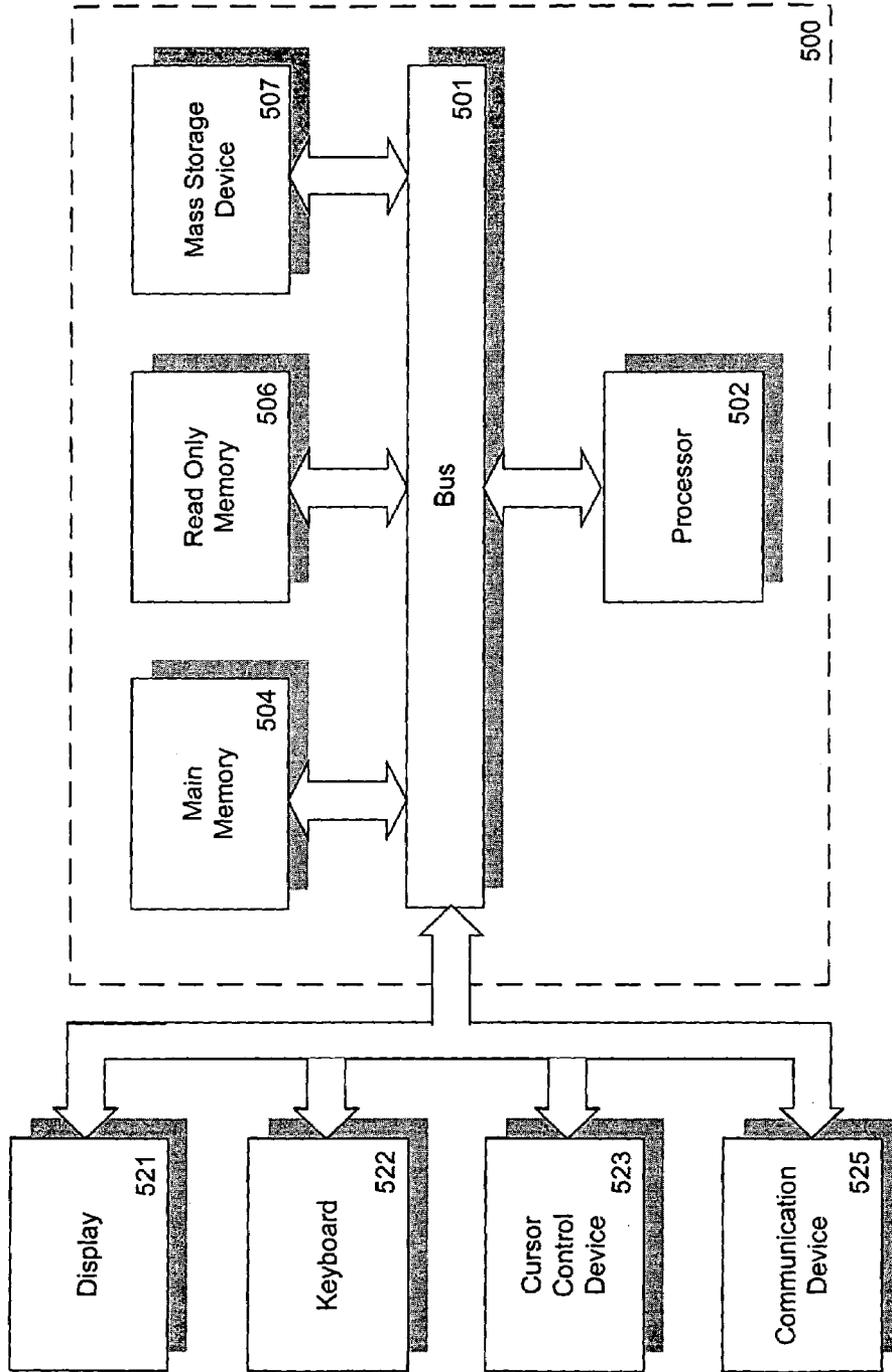


Figure 5

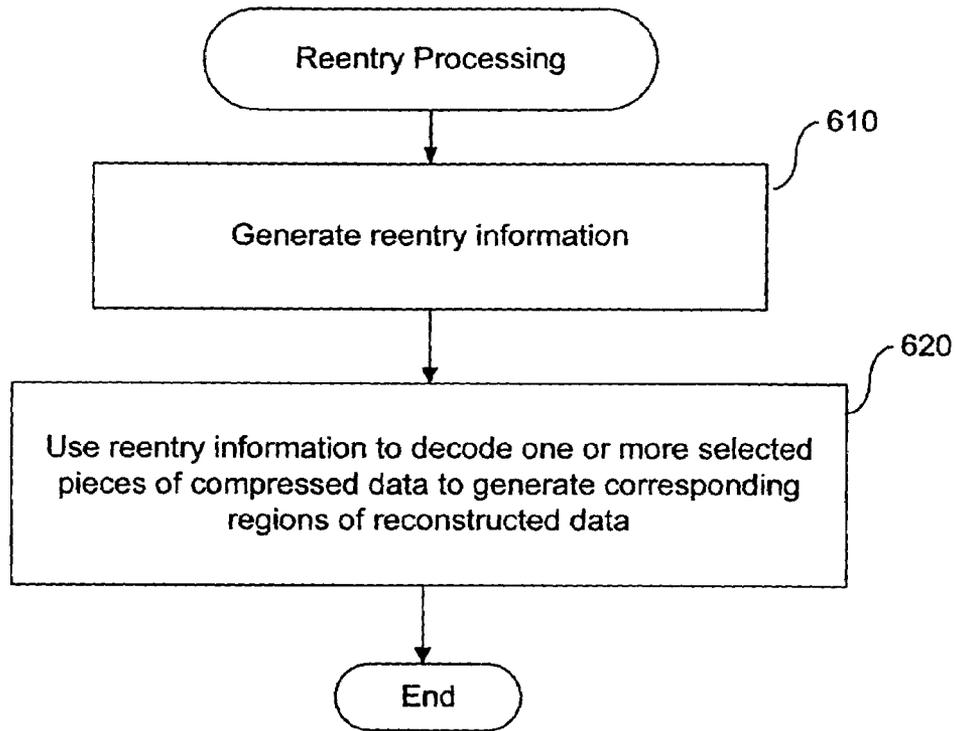


Figure 6

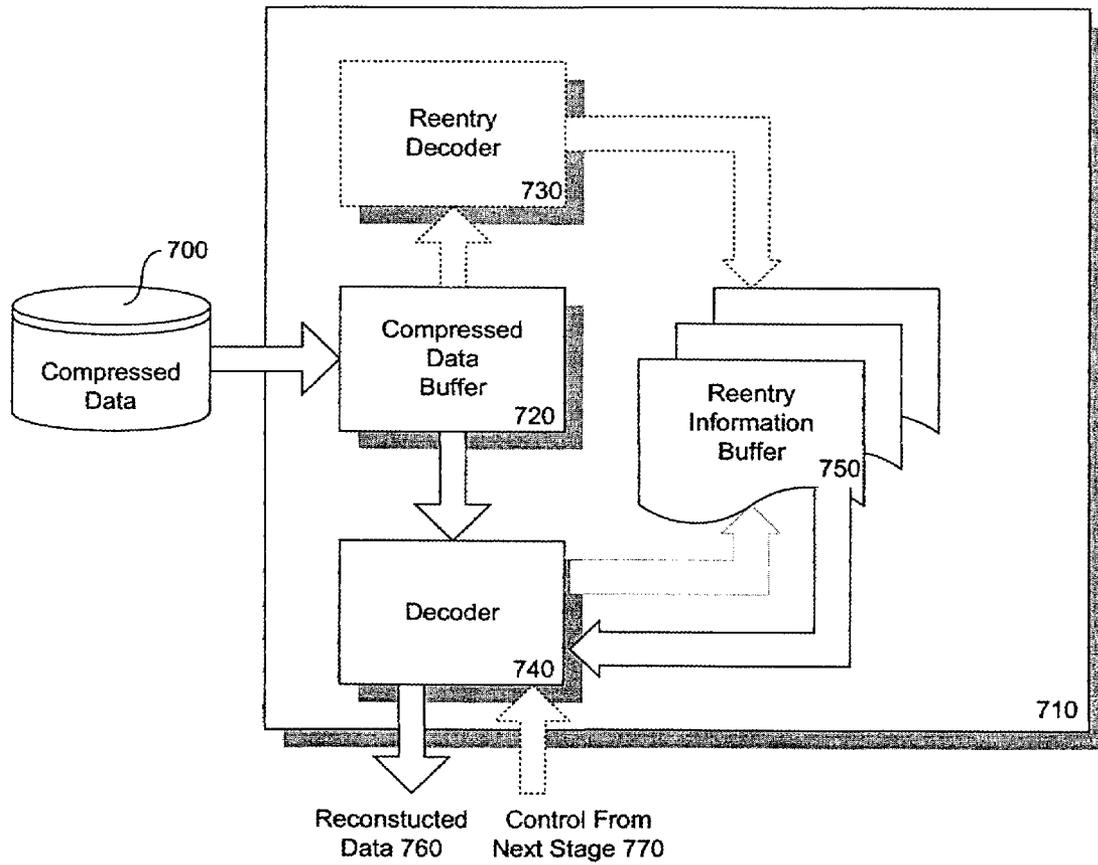


Figure 7

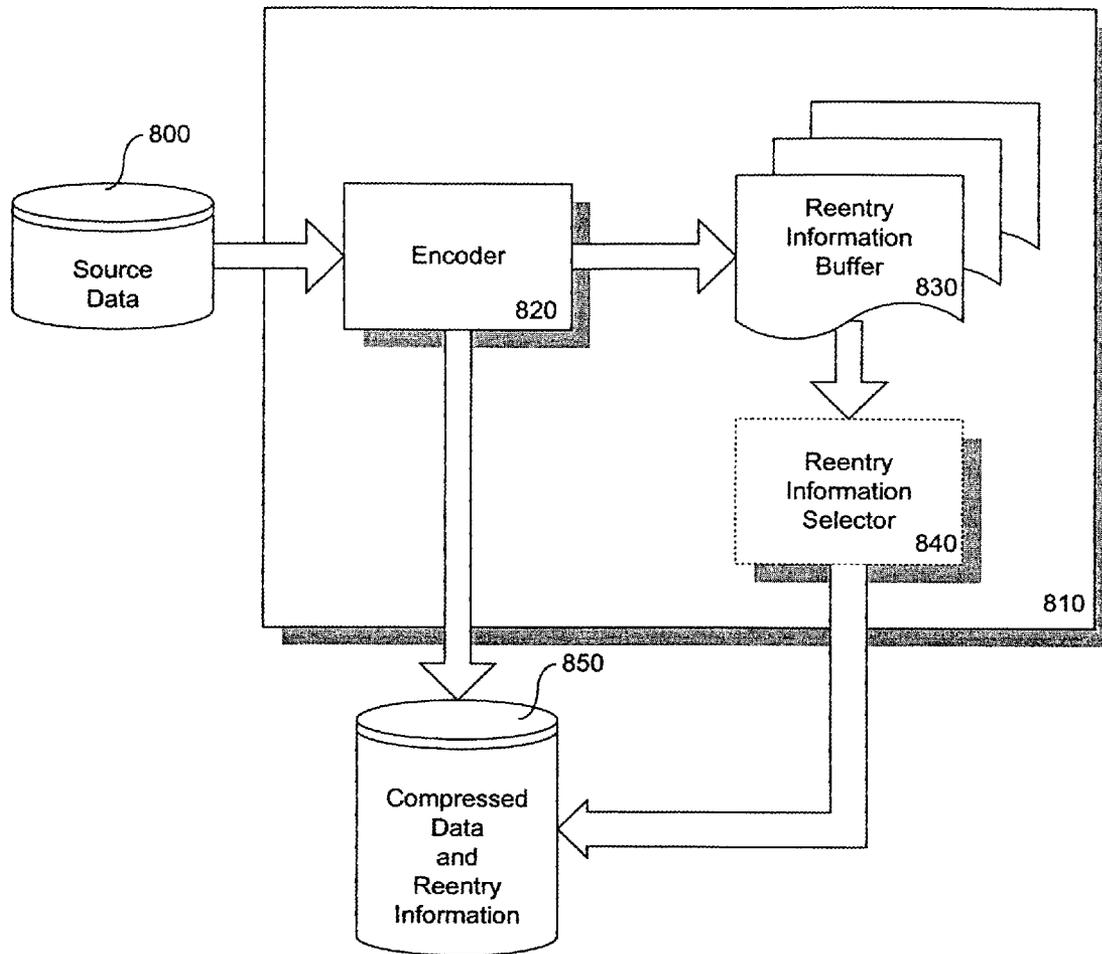


Figure 8

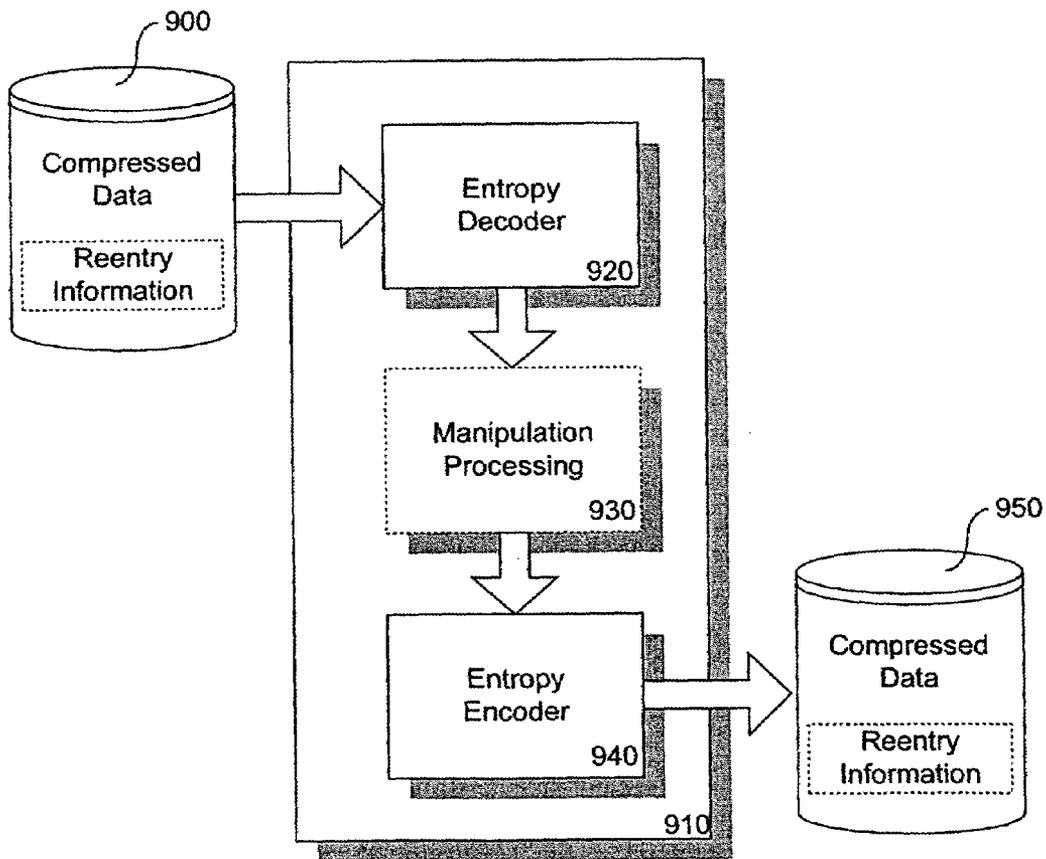


Figure 9

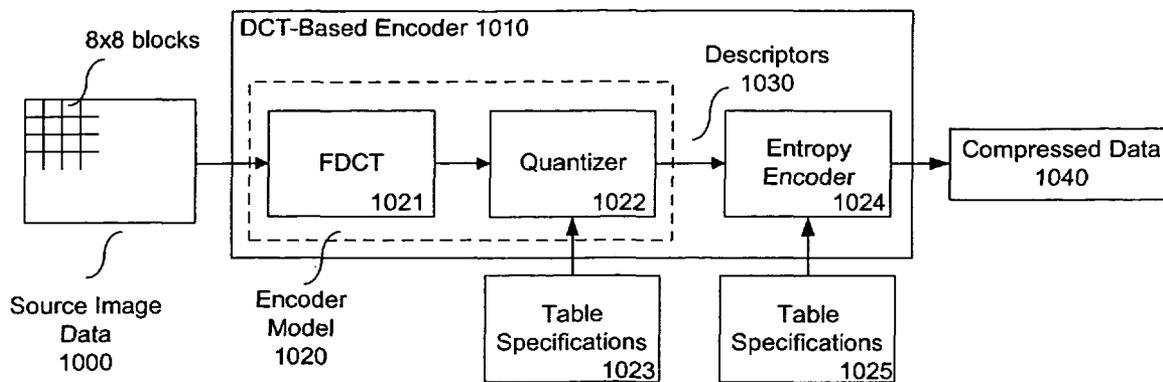


Figure 10 (Prior Art)

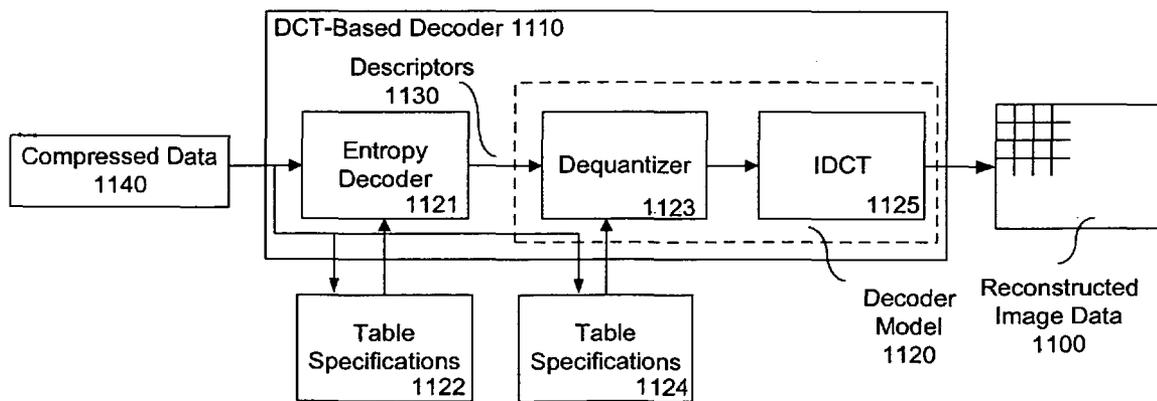
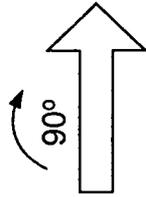


Figure 11 (Prior Art)

| | | |
|------|------|------|
| 1201 | 1202 | 1203 |
| 1204 | 1205 | 1206 |
| 1207 | 1208 | 1209 |

Output Image 1210



| | | |
|------|------|------|
| 1203 | 1206 | 1209 |
| 1202 | 1205 | 1208 |
| 1201 | 1204 | 1207 |

Input Image 1200

Figure 12 (Prior Art)

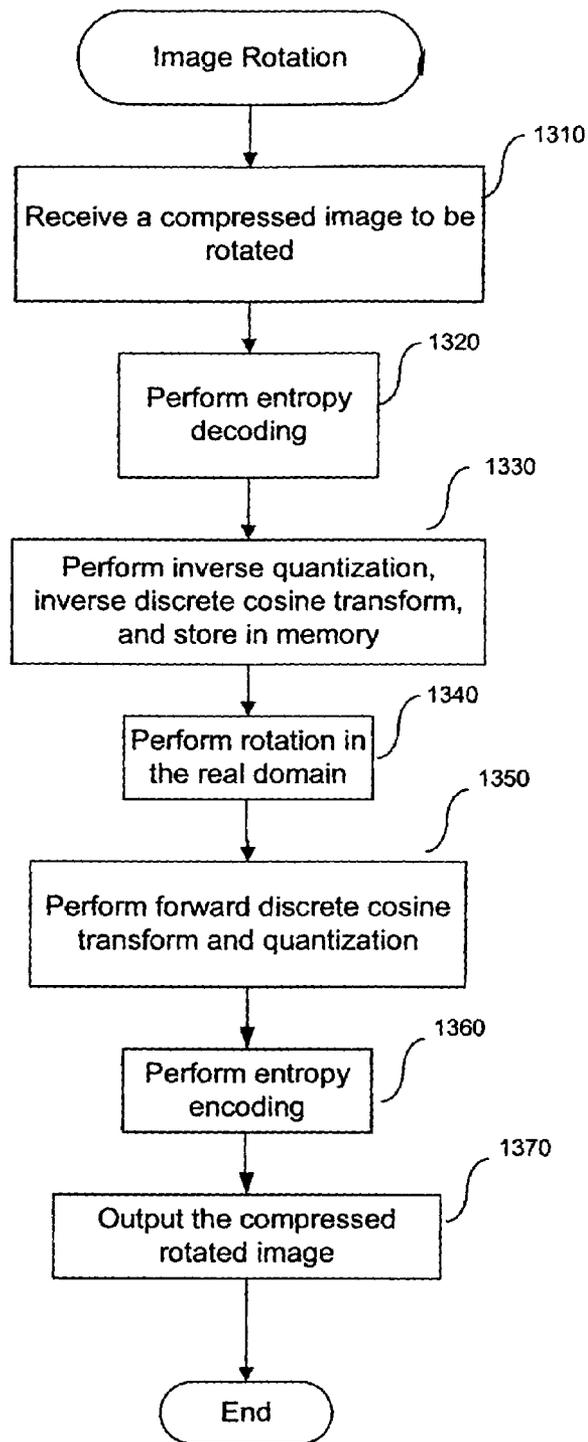


Figure 13 (Prior Art)

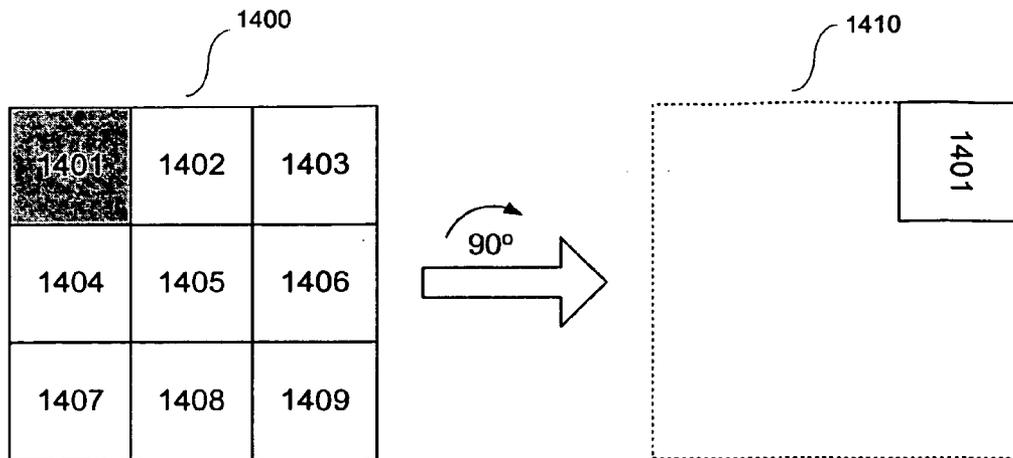


Figure 14A (Prior Art)

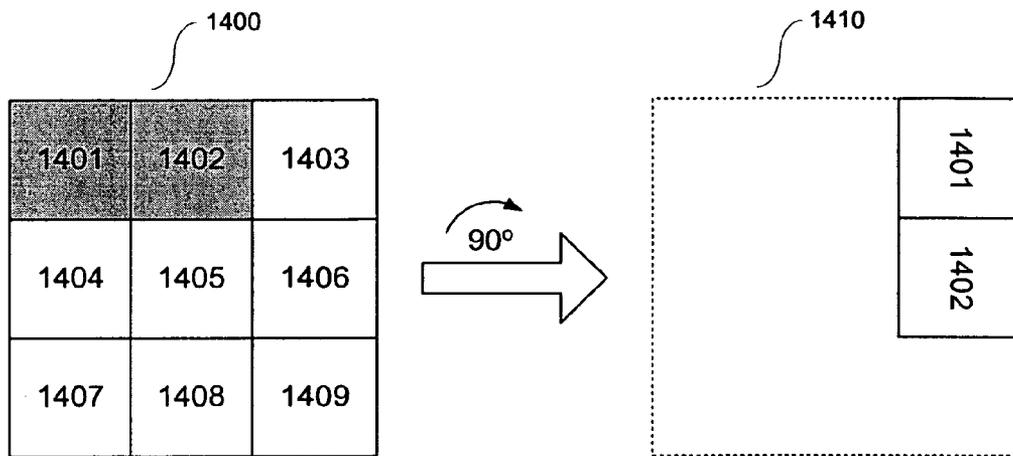


Figure 14B (Prior Art)

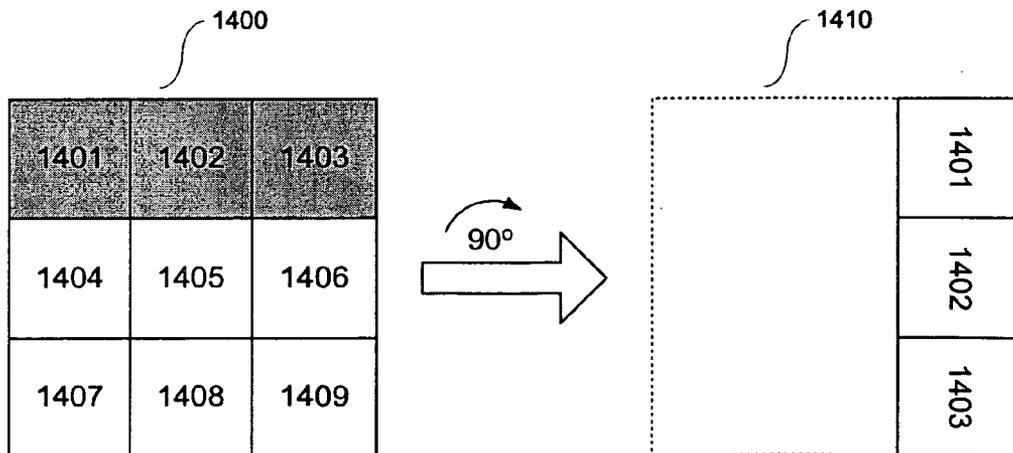


Figure 14C (Prior Art)

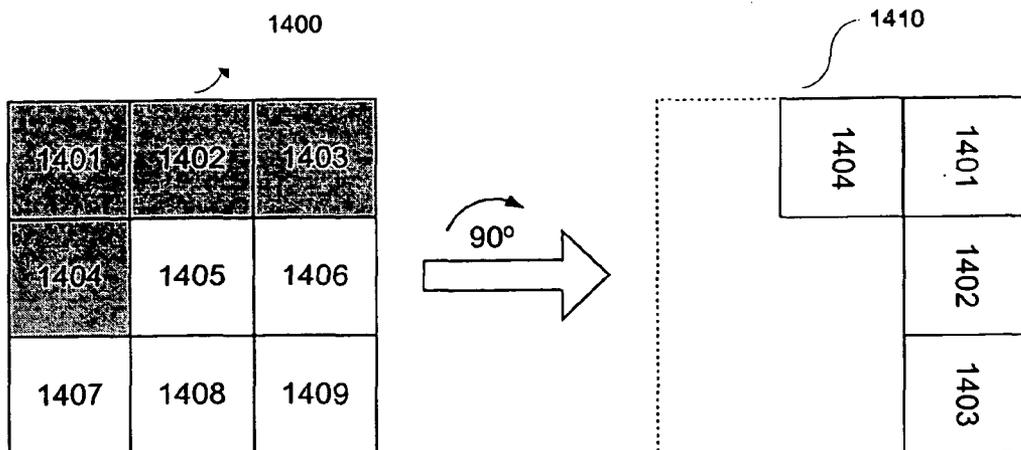


Figure 14D (Prior Art)

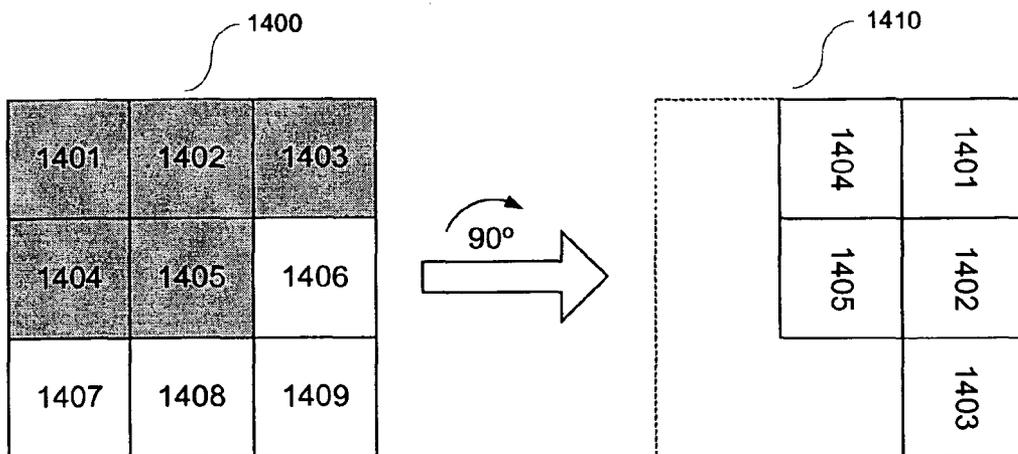


Figure 14E (Prior Art)

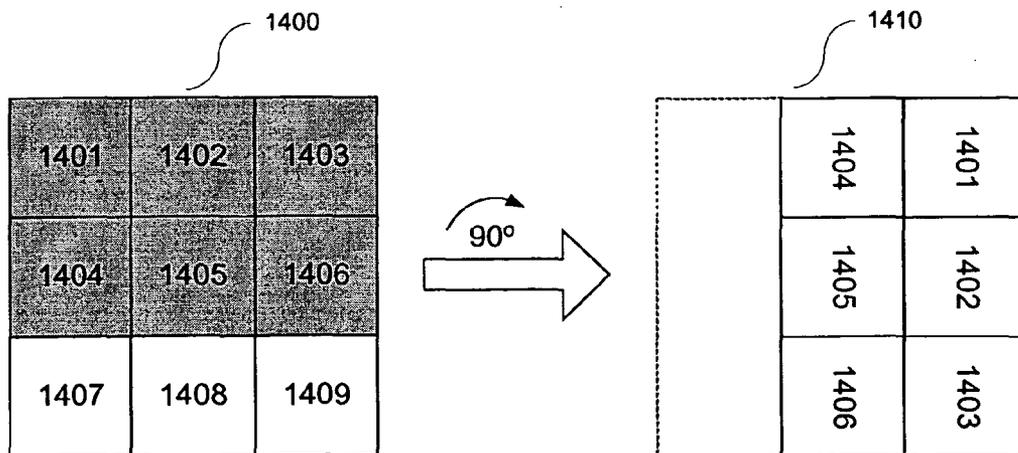


Figure 14F (Prior Art)

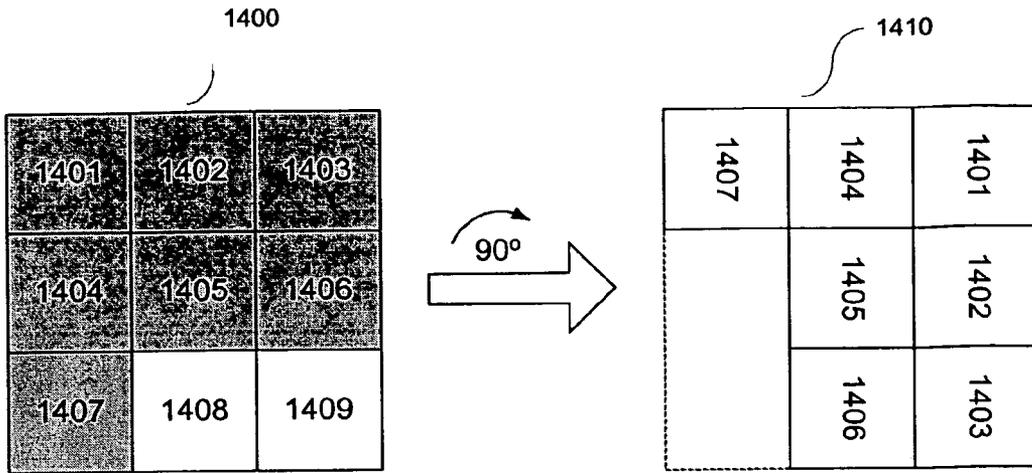


Figure 14G (Prior Art)

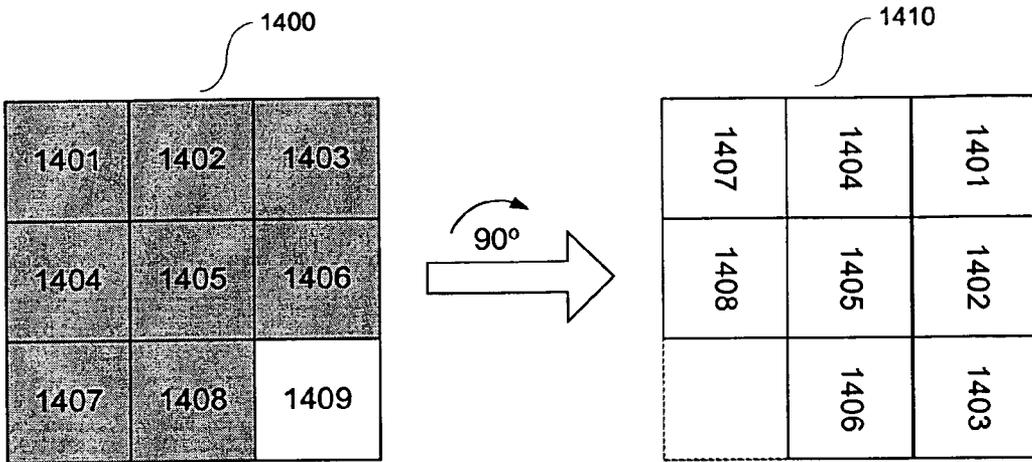


Figure 14H (Prior Art)

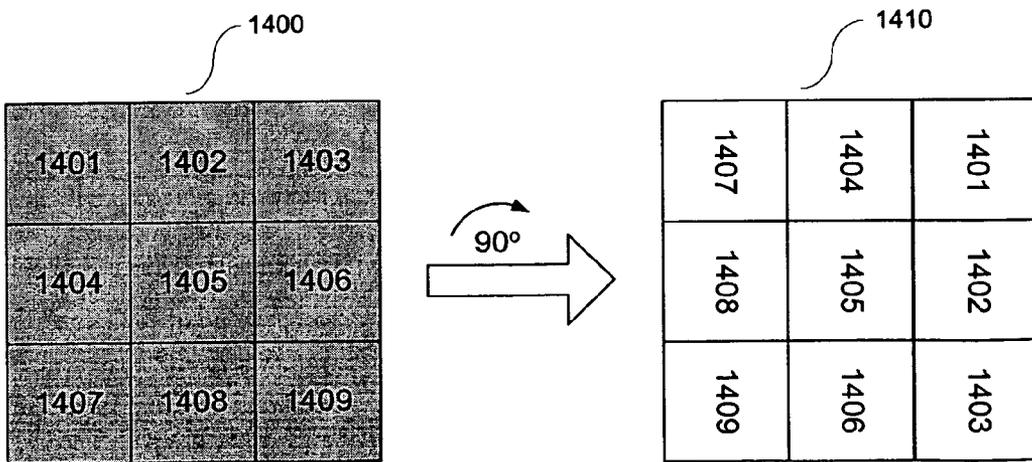


Figure 14I (Prior Art)

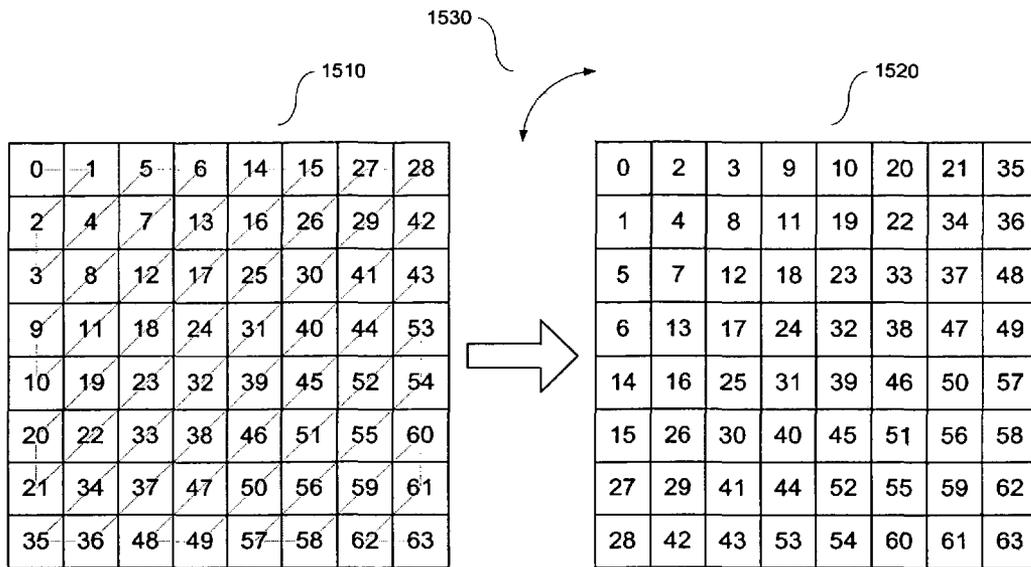


Figure 15 (Prior Art)

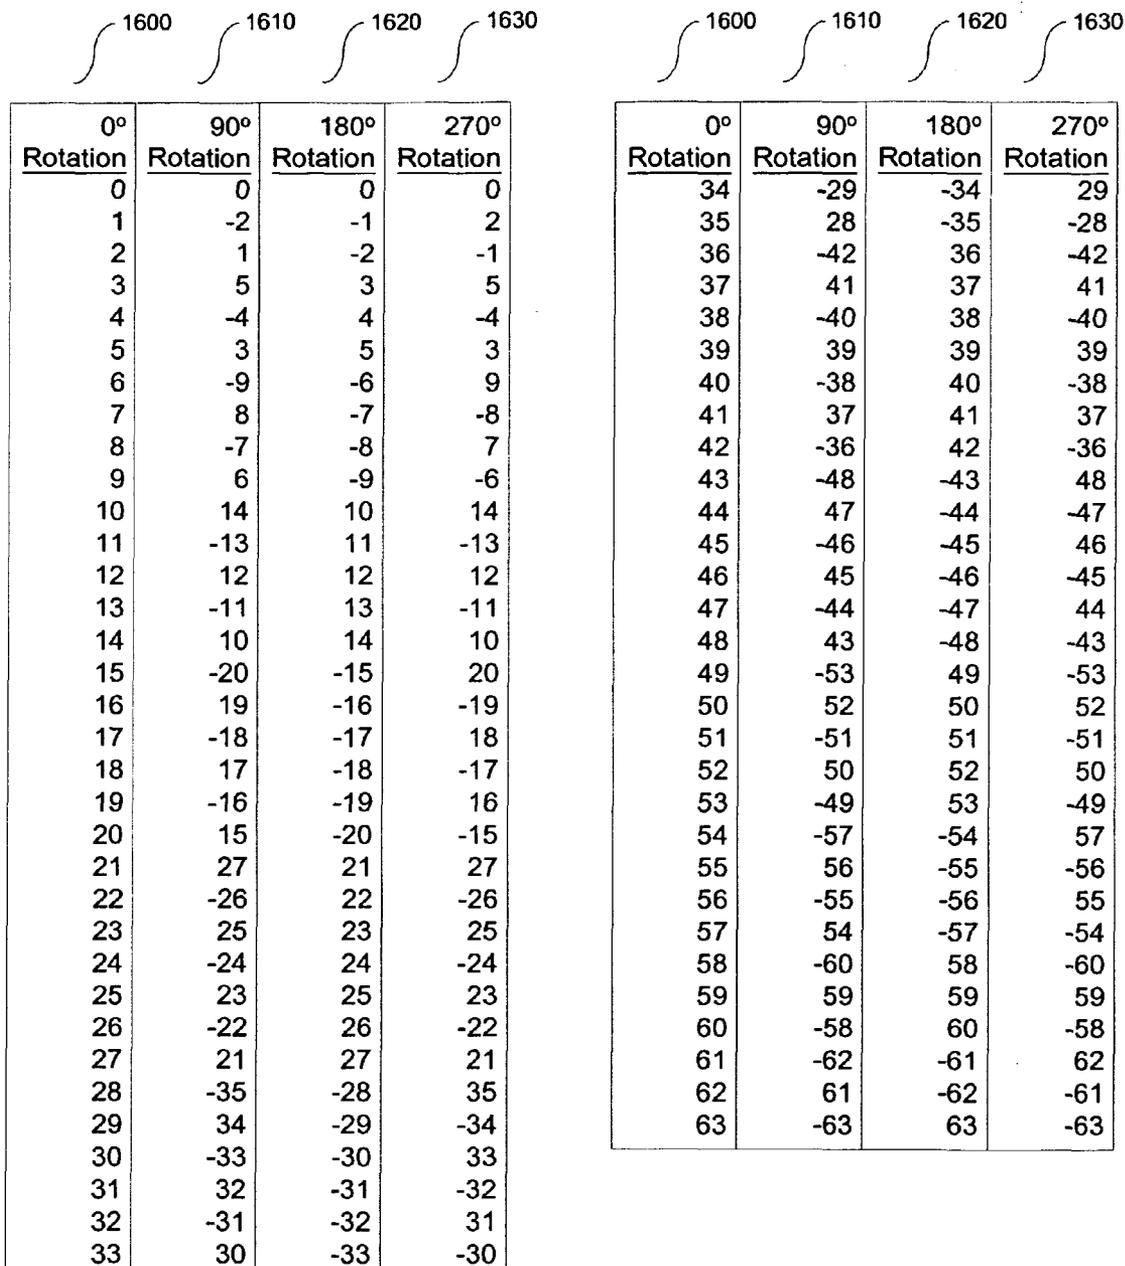


Figure 16 (Prior Art)

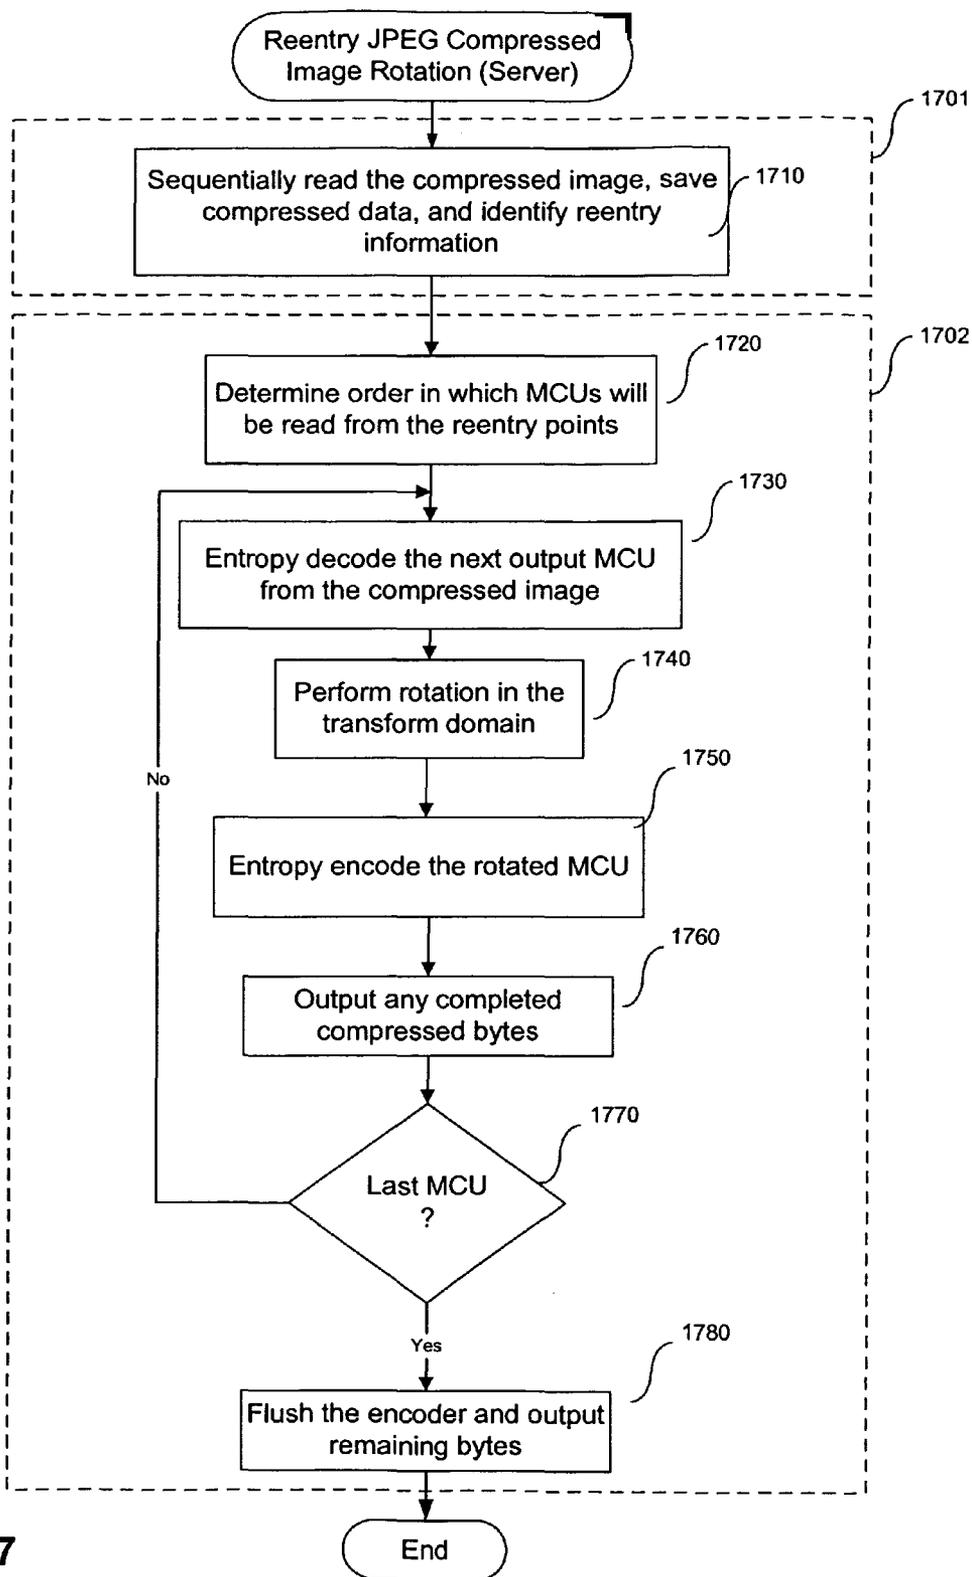


Figure 17

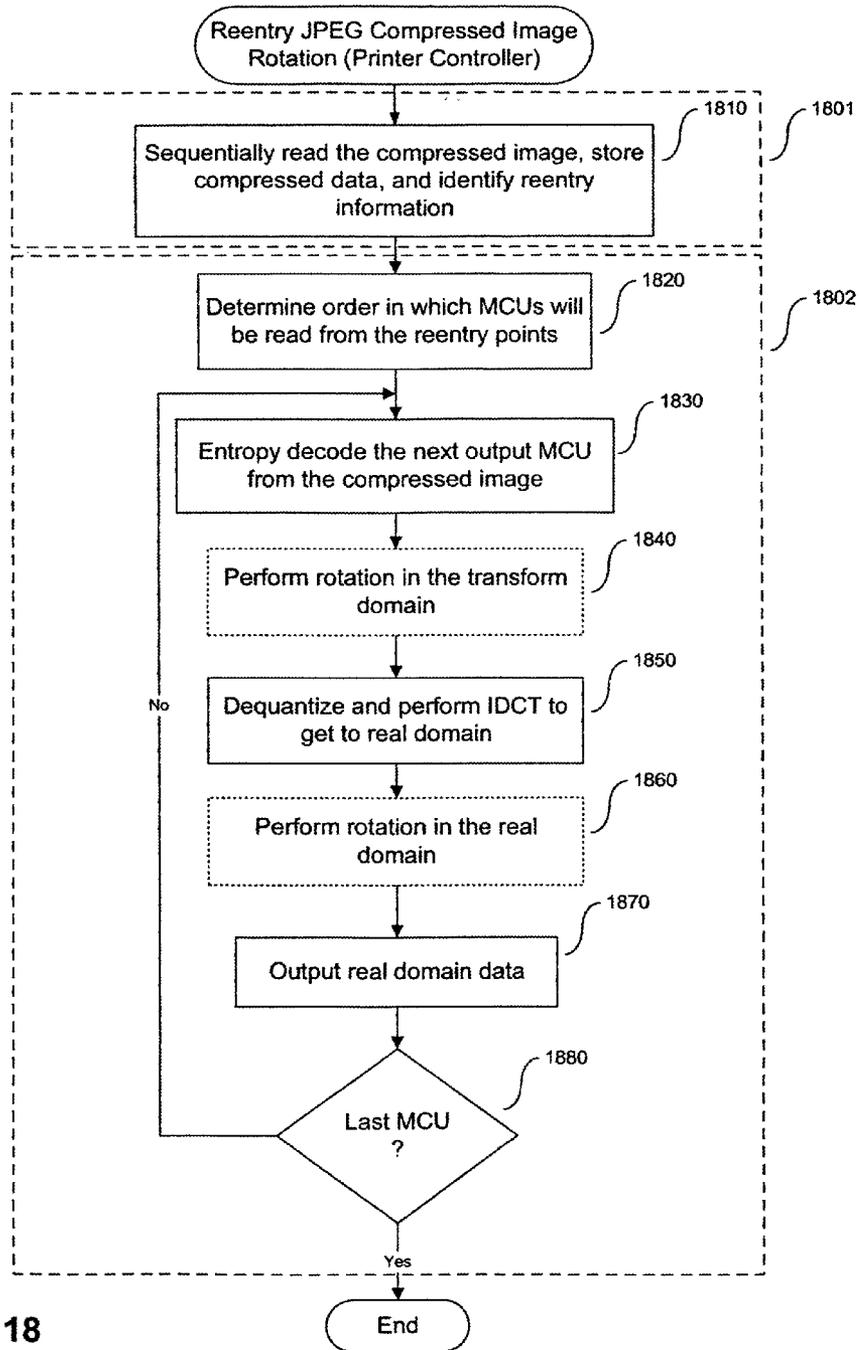


Figure 18

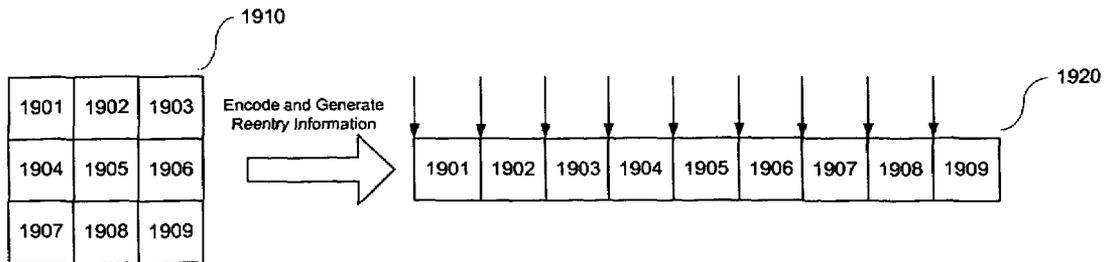


Figure 19A

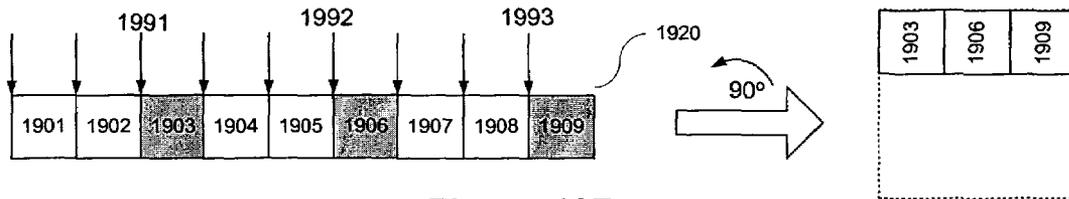


Figure 19B

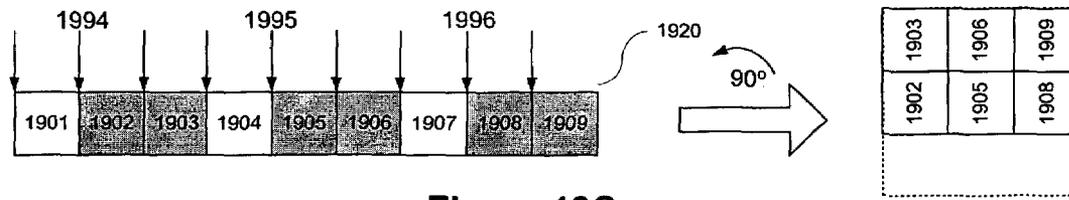


Figure 19C

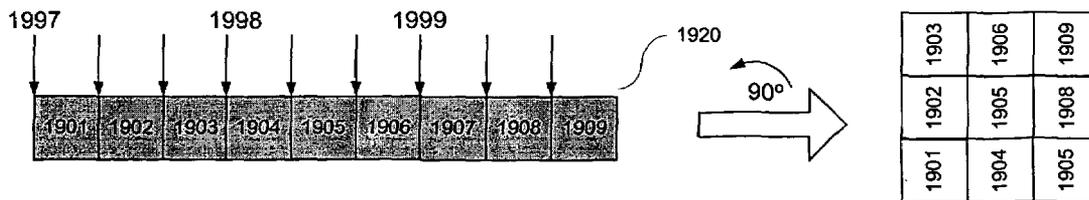


Figure 19D

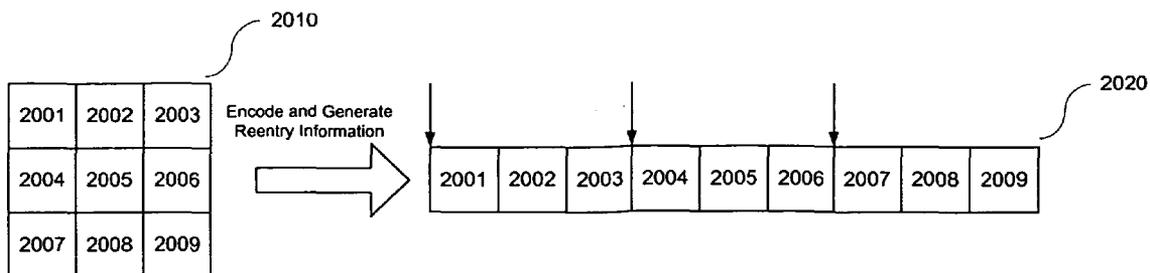


Figure 20A

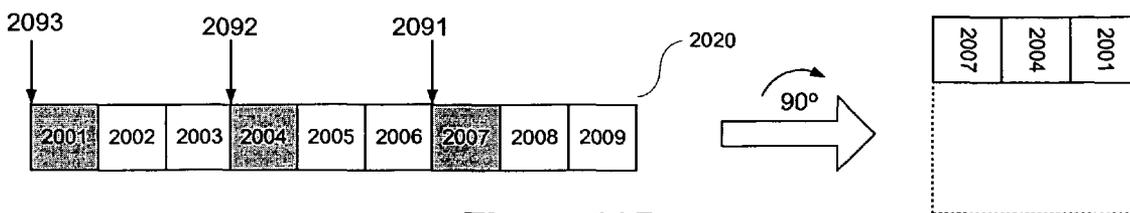


Figure 20B

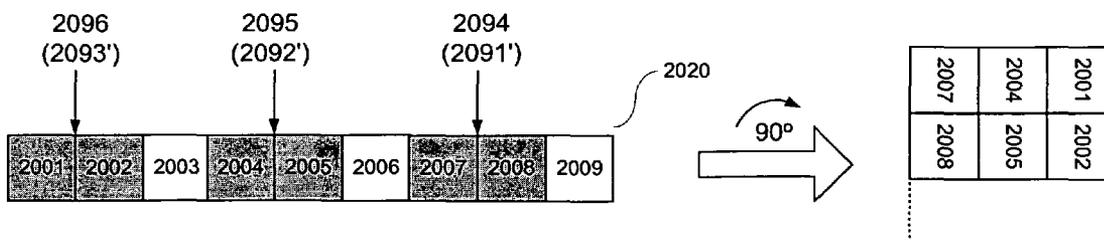


Figure 20C

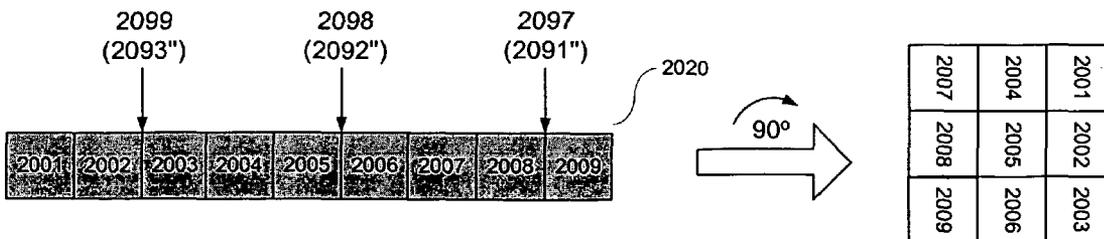


Figure 20D

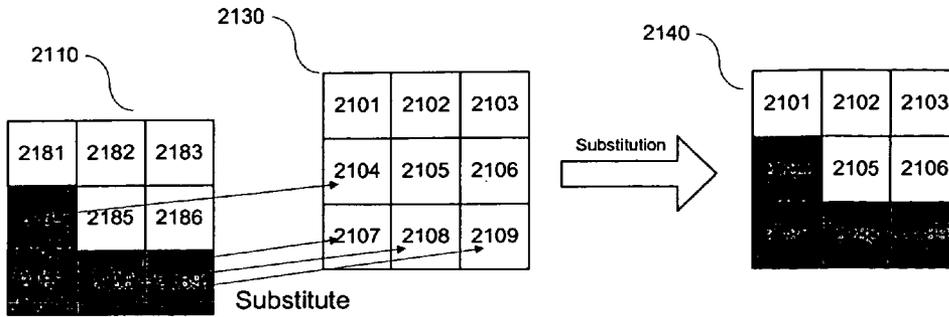


Figure 21A

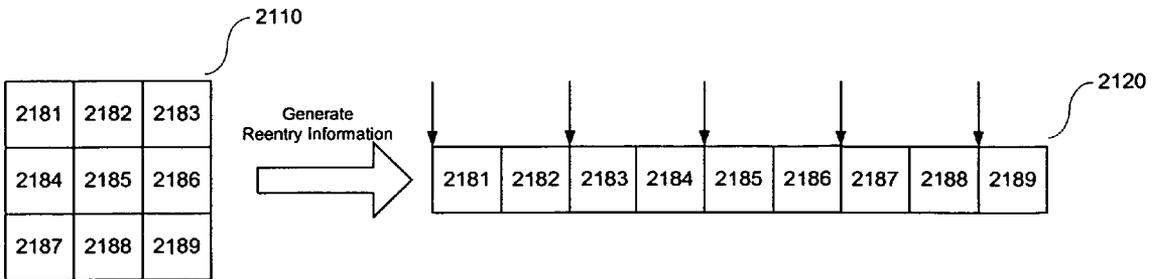


Figure 21B

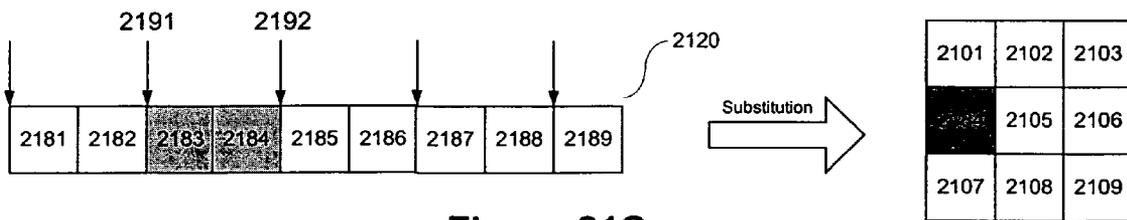


Figure 21C

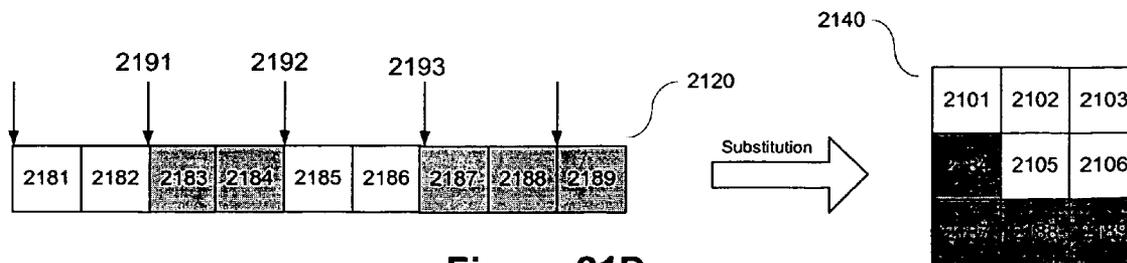


Figure 21D

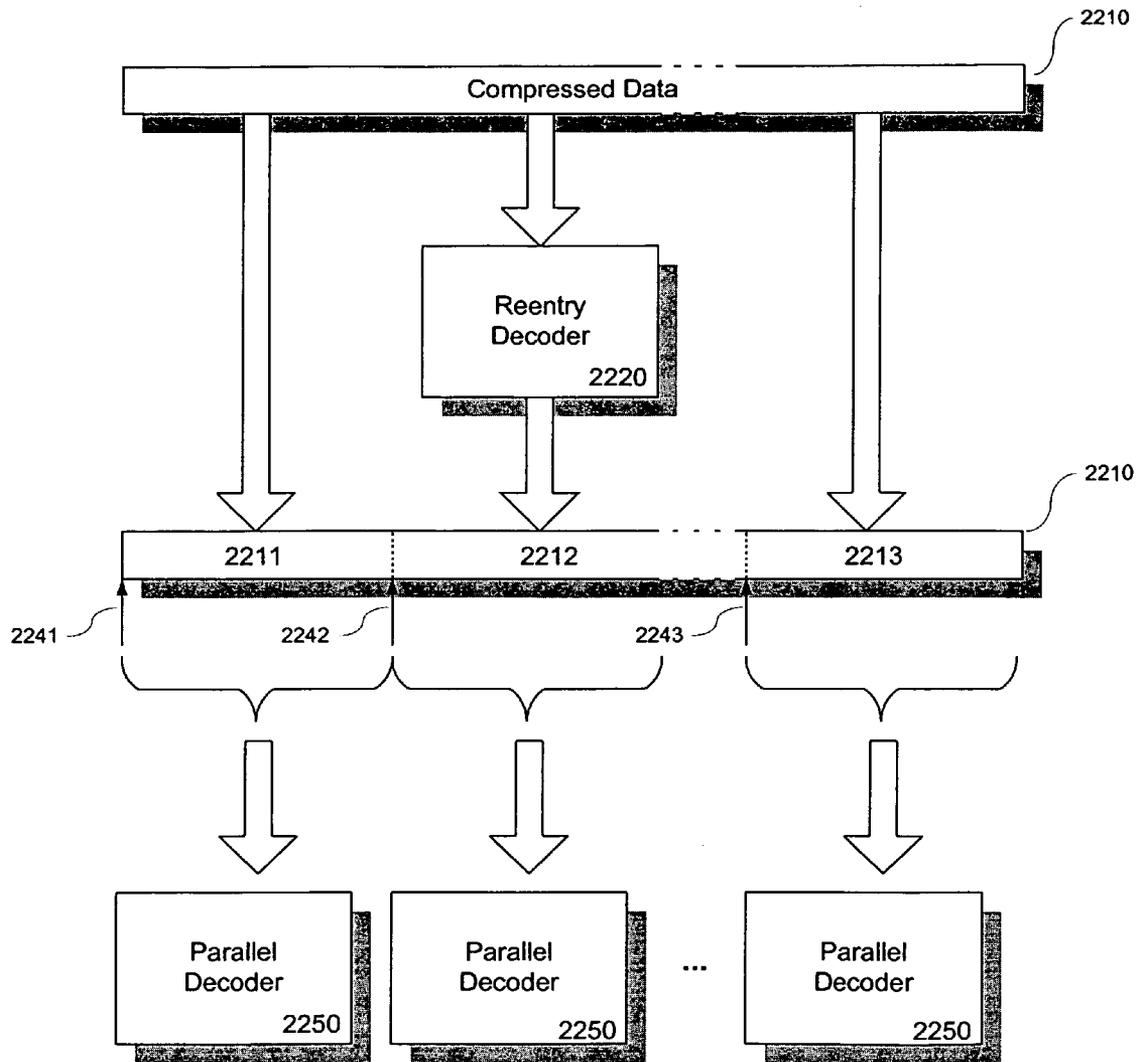


Figure 22

REENTRY INTO COMPRESSED DATA**CROSS REFERENCE TO RELATED APPLICATIONS**

The present application is related to application Ser. No. 09/569,573 filed on May 10, 2000, by Nenad Rijavec and Joan L. Mitchell for "Reordering of Compressed Data".

COPYRIGHT NOTICE

Contained herein is material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction of the patent disclosure by any person as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all rights to the copyright whatsoever.

BACKGROUND OF THE INVENTION**1. Field of the Invention**

The invention relates generally to the generation and use of reentry information into a compressed data stream. More specifically, the invention relates to apparatus and methods for entering compressed data streams at selected points to initiate decoding. These reentry points may be chosen to minimize storage and may employ bit-level resolution pointers. The reentry information may be generated during encoding, decoding, entropy encoding, entropy decoding, partial encoding, partial decoding, or reentry decoding. Embodiments of the present invention are thought to be particularly advantageous for high speed printing applications where the compressed data stream represents images and there are requirements to minimize storage, allow parallel decoding, and achieve high throughput.

2. Description of the Related Art

The purpose of data compression is to represent source data with less data in order to save storage costs or transmission time and costs. Data compression is regarded as "lossless" if the reconstructed data matches the source data. However, in applications where some loss is acceptable, such as pictures, approximating the source data with the reconstructed data, rather than reproducing it exactly achieves a more effective compression. This is called "lossy" compression.

The two basic components of data compression systems are the encoder and the decoder. The encoder compresses the source data (the original digital information) and generates a compressed data stream. The compressed data stream may be either stored or transmitted, but at some point is fed to the decoder. The decoder recreates or reconstructs the data from the compressed data stream. In general, a data compression encoding system can be broken into two basic parts: an encoder model and an entropy encoder. (Some like to have a third part, e.g., a statistical model, which, for simplicity, is currently included in the entropy encoder.) The encoder model generates a sequence of descriptors that is an abstract representation of the data. The entropy encoder takes these descriptors, converts them into symbols, and compresses the symbols taking advantage of the statistics to form compressed data. Similarly to the encoder, the decoder can be broken into basic parts that have an inverse function relative to the parts of the encoder.

Generally, for lossy compression, it is the encoder model and the decoder model that introduce the loss. Generally, the entropy encoding and decoding are lossless. In such cases, lossless transcoding becomes possible. The compressed data

is entropy decoded to intermediate data, which may or may not be identical to the descriptors, and then the intermediate data can be fed into a different entropy encoder to create a different compressed data stream. Examples of entropy encoders are Huffman encoders and arithmetic encoders. Converting between these methods is an example of transcoding. For a more detailed discussion of Huffman variable-length coding see: M. Rabbani and P. W. Jones, *Digital Image Compression Techniques*, Tutorial Texts in Optical Engineering, Vol. TT7, SPIE Optical Engineering Press, Bellingham, Wash. 1991; and W. B. Pennebaker and J. L. Mitchell, *JPEG: Still Image Data Compression Standard*, Van Nostrand Reinhold, N.Y. (1993).

Without loss of generality, most examples herein will be presented with respect to images as the source data. In the context of high speed printing, data compression is key to lessening the amount of data transmission. As a practical matter, large-scale digital color printing has been unaffordable until recently for most applications. Consequently, the field of image processing has not yet had to address many of the problems associated with efficiently processing the absolute torrent of data required by continuous-tone color images.

An image, in the context of this application, is an electronic representation of a picture as an array of raster data. Image data can be generated by a computer program, or formed by electronically scanning such items as illustrations, drawings, photographs, and signatures. For monochrome images each sample in the array of raster data has an intensity value. For traditional facsimile images only two values are allowed (black or white) and so 1-bit per sample is sufficient. For continuous-tone pictures 8-bits per sample is more common. Some applications such as medical images require higher precision and use 12 bits per sample. Color images use three or four values for each position in the raster array. Typically scanners and displays use red, green, and blue (RGB) as primary colors or components. Printers often use cyan, magenta, and yellow (CMY) inks or toners. Sometimes a fourth black ink or toner is added (CMYK). The components may be interleaved (RGB,RGB,RGB) or separated into multiple arrays (RRR,BBB,GGG).

The size of digital images is rapidly increasing. In the late 1970s, 8.5 inch×11-inch facsimile images were standardized at 1728 picture elements (pels) per line by about 1100 lines (nominally 200 pels/inch by 100 lines/inch). A finer resolution of 200 lines/inch was also allowed. Now, many digital printers print binary images at 600 pels/inch by 600 lines/inch, an order of magnitude increase in data. Some of today's printers are supporting 1200 pels/inch by 1200 lines/inch.

Data compression was essential to making digital facsimile practical. Two international facsimile data compression standards were developed. The CCITT Group 3 (G3) facsimile machines were made for the public phone network and originally expected errors during transmission. Group 4 (G4) facsimile machines were intended for data networks and assumed that transmissions were error-free.

The G3 standard Modified Huffman (MH) algorithm coded each binary line one-dimensionally (i.e., independently) with unique end-of-line (EOL) codes separating each line. Transmission errors were detected when the decoded lines did not match the expected line length. Since the compression process could be restarted on the next line by searching for the EOL, only one line was corrupted. The G3 Modified READ (MR) algorithm encoded some lines two-dimensionally (i.e. with reference to the previous line) and could not be restarted except at the one-dimensional lines. A

tag bit after the EOL code indicated whether the next line was coded with one-dimensionally or two-dimensionally. The G3 standard required that at least every other line be coded with MH (and thus restartable after errors) for the standard resolution and at least every fourth line be coded with MH for the finer resolution. In the absence of transmission errors, the G3 algorithms are lossless (i.e., the decoder's output image exactly matches the encoder's input image).

G4 machines used the Modified Modified READ (MMR) data compression algorithm. MMR assumed an all white history line before the top image line and encoded all lines two-dimensionally (i.e., with reference to the history line) using the same two-dimensional codes as MR. The EOL codes were not used because transmission was assumed to be error-free. The G4 algorithm was soon extensively used in the error-free computer environments too.

IBM MMR was derived from the G3 MR algorithm before the G4 MMR standard was established. It is defined in K. L. Anderson, F. C. Mintzer, G. Goertzel, J. L. Mitchell, K. S. Pennington, and W. B. Pennebaker, "Binary Image Manipulation Algorithms in the Image View Facility," *IBM J. Res. Develop.*, vol. 31, 16-31 (1987). The first line is coded exactly like the first line of G3 MR starting the compressed data with an EOL with a 1-D tag. The first line is 1-D coded without the need for a history line. An EOL with a tag follows the first compressed line of data. If the tag indicates that 2-D coding follows, an arbitrary number of lines are encoded the same as G4 MMR. EOLs with 1-D or 2-D tags are allowed to be encoded at the start of any line. After the 1-D tag, the line is compressed with MH and must be followed by another EOL. The IBM MMR compressed data terminates with six EOLs with 1-D tags just like G3 MR.

Detailed examples of such data compression and decompression algorithms are given in: J. L. Mitchell, K. L. Anderson, and G. Goertzel, "Method for Encoding and Decoding a Digital Image," U.S. Pat. No. 4,725,815 issued Feb. 16, 1988; and J. L. Mitchell, K. L. Anderson, and G. Goertzel, "Method for Encoding and Decoding a Digital Image," U.S. Pat. No. 4,888,645 issued Dec. 19, 1989. The use of intermediate data called "run ends" is disclosed in the above. Rather than the individual bits representing the source image, the encoder model converts each raster line into a sequence of numbers that represent the distance from the left edge of the last pel in each run. An example of raster to run end conversion is given in: K. L. Anderson, G. Goertzel, and J. L. Mitchell, "A Method for Converting a Bit Map of an Image to a Run Length or Run End Representation," U.S. Pat. No. 4,610,027 issued Sep. 2, 1986. The entropy encoder takes the run ends and converts them into Huffman codes according to the appropriate standard.

Adaptive Bilevel Image Compression (ABIC) is a lossless binary arithmetic coding algorithm. For more details, refer to R. B. Arps, T. K. Truong, D. J. Lu, R. C. Pasco, and T. D. Friedman, "A multi-purpose VLSI chip for adaptive data compression of bilevel images," *IBM J. Res. Develop.*, Vol. 32, 775-795 (1988) and to G. G. Langdon Jr., J. L. Mitchell, W. B. Pennebaker, and J. J. Rissanen, entitled "Arithmetic Coding Encoder and Decoder System", U.S. Pat. No. 4,905,297 issued Feb. 27, 1990.

ABIC encodes an image as one compressed data stream. Unlike the G3 MH, G3 MR, or the Joint Bi-Level Image Experts Group (JBIG) encoding algorithm, an ABIC compressed image has no markers or EOL codes. Decoding must be sequentially from the upper left pel to the bottom right pel in raster scan order.

The Joint Bi-Level Image Experts Group (JBIG) standardized the next generation facsimile data compression technique based on arithmetic coding. JBIG-1 allows for byte-aligned markers to separate the compressed data and identify strips of predetermined number of lines. The SDNRM marker (0xFF02) restarts the arithmetic coder A-register and C-register for the next stripe but keeps the probability estimates and uses the previous line(s) as history. The compressed data also starts byte-aligned. The SDRST marker (0xFF03) starts coding as if this next line were a new image. Thus, the SDRST identifies an independently decodable piece of compressed data that follows. The SDRST is required at the end of bit planes. This allows the compressed data to be shuffled without decoding to convert between stripes organized by bit plane and stripes organized by full width rectangles of multiple-bit image. Also the markers allow shuffling of the progressively encoded JBIG-1 image to shuffle the multiple resolutions of compressed data organized by resolution layer to compressed data organized by full-width regions of the image. Stripes must occur at a fixed number of complete lines except at the bottom of the image. Consequently, there is no ability to reenter the compressed data stream within a line. Stripe boundaries are selected at the encoder and are maintained for all resolution levels.

Compression of continuous-tone color images can be lossless, but the 24-(three 8-bit component colors) to 32-(four 8-bit component colors) fold increase in the data makes lossy compression often more practical. The JPEG (Joint Photographic Experts Group) and MPEG (Moving Picture Experts Group) standards are examples of lossy data compression standards. See generally: W. B. Pennebaker et al.; J. L. Mitchell, W. B. Pennebaker, C. E. Fogg, and D. J. LeGall, *MPEG Video Compression Standard*, Chapman & Hall, NY (1997); B. G. Haskell, A. Puri, and A. Netravali, *Digital Video Compression Standard, An Introduction to MPEG-2*, Mitchell & Pennebaker, Editors, Chapman & Hall, NY (1997); and K. R. Rao and J. J. Hwang, *Techniques and Standards for Image, Video, and Audio Coding*, Prentice Hall PTR, Upper Saddle River, N.J. (1996).

Both the MPEG and JPEG standards employ transform coding. Each color component is divided up into 8x8 blocks. The forward Discrete Cosine Transform (FDCT) of each block is performed. The transform coefficients are then quantized. This step introduces the largest loss as the quantized coefficients are rounded to integers. Then a lossless entropy coding technique (Huffman coding for MPEG and baseline JPEG or also arithmetic coding for JPEG) encodes the quantized integers. The decoder performs the entropy decoding to recover the quantized coefficients. Then an inverse quantization (or dequantization) step multiplies each coefficient by its quantization. The dequantized coefficients are fed to an inverse Discrete Cosine Transform (IDCT) to reconstruct the 8x8 samples in the block.

JPEG employs the concept of a "minimum coded unit" (MCU), which refers to a group of one or more DCT blocks in lossy coding and samples in lossless coding. In JPEG, entropy coding is always performed on a complete MCU. MPEG has a similar concept and calls it a macroblock.

JPEG images contain byte-aligned "markers." A marker is a unique code that can be located by scanning the compressed data stream in which it is embedded. In JPEG, markers are unique byte-aligned codes that can be used to identify the location and purpose of header and entropy-coded segments within an image. A marker comprises a marker code that identifies the function of the particular segment it precedes and a prefix, i.e., 0xFF. Most JPEG markers are followed by length fields and communicate

header information. A few markers such as the Restart Markers indicate where the image data can be restarted and thus independently encoded or decoded. The Restart Markers include in their three least significant bits a modulo eight counter. This enables detection of any corruption of a previous Restart Marker as part of an error recovery mechanism.

The MPEG video frames often depend upon previous or future frames. At the start of independent frames (I-frames) the decoding is restartable. The MPEG syntax has unique 32-bit byte-aligned start codes. The sequence header, group-of-pictures header, picture header, and slice header all have their own start code. MPEG I-frames are also restartable at the slice headers.

An overview of audio coding with some details regarding MPEG audio is given in Section 10.3 entitled "Audio coding" of Rao et. al.

Various text compression algorithm are described in T. C. Bell, J. G. Cleary, and I. H. Witten, *Text Compression*, Prentice Hall PTR, Englewood Cliffs, N.J., (1990). Lempel-Ziv (LZ) compression techniques tend to split into two types LZ1 and LZ2. The LZ1 keeps a history of the last N bytes and tries to code the data as a pointer into the history buffer followed by the number of matching bytes. Otherwise, it signals that it is going to send the raw data. LZ2 type compression, e.g., LZW, constructs a dictionary of patterns already encountered in the data. After each repeat of a pattern, a new pattern consisting of the old pattern extended by one character is added to the dictionary.

Images need to be transferred to a print server in compressed form; and after manipulation by the print server, the print server often needs to recompress the manipulated image. This is particularly important if the printer controller has only hardware decoding and the bandwidth between the print server and the printer controller is inadequate.

While JPEG has a structure that supports independently decodable pieces of compressed data between Restart Markers, going all the way to the real domain is time consuming and wasteful. The markers are very helpful to enable parallel encoding and/or decoding starting at marker boundaries. It should be noted that the encoder can only place Restart Markers at fixed intervals in the source data and consequently the Restart Markers may not be optimally placed for the decoder. To change the placement of the Restart Markers would require at least a transcoding of the compressed data stream.

An example of arbitrary parallel decoding is given in S. T. Klien et al., S. T. Klien and Y. Wiseman, *Parallel Huffman Decoding*, Proceedings: Data Compression Conference, pp. 383-392, Snowbird, Utah, (2000), (hereinafter "Klien et al."). Klien et al. describe a parallel algorithm for decoding a Huffman encoded image that exploits the tendency of Huffman codes to resynchronize quickly. When more than one processor is available, the compressed data stream can be split into pieces, and each processor can be assigned one piece of the compressed data stream for decompression. Klien et al. suggest letting each processor overflow beyond its assigned piece into the next piece until its results synchronize with the processor that has been assigned to the next piece of compressed data. Once synchronization has been detected, the processor can stop or be assigned another piece. Synchronization is detected because each processor has saved the index of the last bit in each codeword. The processor of the previous piece can examine this list. Importantly, the assigned pieces do not necessarily begin at codeword boundaries. Therefore, the first codes are expected to be erroneously decoded until the Huffman property of

self-synchronization occurs. When applied to JPEG compressed images, the position and the DC predictors of the partially decoded image are guessed and subsequently corrected after synchronization has been established. Further, the approach suggested by Klien et al. has no additional information available as they try to reenter the compressed data stream at arbitrary boundaries. Consequently, correct display of the image must wait until all synchronization points have been established.

One very common image manipulation operation is rotation by a multiple of 90°. For example, images are often rotated to accommodate a particular page orientation, e.g., landscape or portrait, and/or to accommodate particular user-specified job attributes, such as impositioning. While current print servers typically support various forms of image manipulation, these manipulations may introduce multi-generation losses and make inefficient use of processing resources. As discussed further below, image rotation processes must typically accumulate the whole image in memory or on disk before the rotated image can be output. When rotating images, often the first pixels read by the rotation process contain some of the last pixels to be output.

Examples of prior art rotation of binary images in the real domain are: K. L. Anderson, F. C. Mintzer, G. Goertzel, and J. L. Mitchell, entitled "Method for Rotating a Binary Image", U.S. Pat. No. 4,627,020 issued Dec. 2, 1986; and D. R. Pruett, G. Goertzel, and G. R. Tompson (sic), entitled "Method for Rotating a Binary Image," U.S. Pat. No. 4,837,845 issued Jun. 6, 1989. A disadvantage of working in the real domain is that the entire image has to be decompressed, temporarily buffered, rotated, and then recompressed. This is costly in terms of both storage and time.

The above patents disclose storing the full image plus a much smaller temporary buffer. When the full image cannot be in contiguous storage, the method disclosed in K. L. Anderson and J. L. Mitchell, "System for Rotating Binary Images," U.S. Pat. No. 4,658,430 issued Apr. 14, 1987 may be used. However, it still needs sufficient storage for the entire source image plus an additional buffer.

To avoid the time to decode all the way to the raster image, rotation on run end data is disclosed in K. L. Anderson, "Fast Algorithm for Rotating an Image in Run End Form," *IBM Technical Disclosure Bulletin*, Vol. 32 no.6B pp. 299-302 (1989). For typewritten text documents this run end data is significantly smaller than the source data. It has the disadvantage, however, that in the worst case there is a 16 to 1 expansion for alternating single pel runs since each run end is saved in 16 bits.

Rotation of continuous-tone images is less complicated because the pixels are generally on byte-boundaries, on the other hand they are likely to be 8 to 24 times larger. In addition, decoding to the real domain, and after rotation, reencoding is CPU intensive. An additional complexity is that lossy decoding and then reencoding has a multi-generation problem, namely the recompressed data doesn't match the previously compressed data.

The aforementioned approach is limited in that it converts between the transform domain and the real domain to prepare for the rotation process. W. B. Pennebaker, I. R. Finlay, J. L. Mitchell, K. L. Anderson, P. J. Sementilli, Jr. entitled "Intermediate Format for Representing Transform Data" Japanese patent JA02698034 issued Sep. 19, 1997 discloses lossless rotation of JPEG images in the transform domain. The entropy-decoded DCT transform coefficients are transposed within a block, some signs are changed, and the blocks must still be reordered. While CPU cycles have been drastically cut by avoiding the dequantization, IDCT,

FDCT, and requantization, as above, the intermediate data could potentially expand (e.g., 2 to 1) and the buffering requirements are still large.

BRIEF SUMMARY OF THE INVENTION

Apparatus and methods are described for entering compressed data streams at selected reentry points thereby allowing efficient manipulation of the compressed data and minimizing storage requirements. According to one embodiment of the present invention, reentry information into a compressed data stream is generated during the process of decoding to source data, during the process of encoding from source data, during a partial decoding of the compressed data stream, during entropy encoding of the compressed data stream, or during entropy decoding of the compressed data stream. Then, the reentry information may be used to decode one or more selected pieces of compressed data to generate corresponding regions of reconstructed data. Advantageously, in this manner, pieces of compressed data need not be sequentially processed according to the order in which they appear in the compressed data stream. Rather, pieces of compressed data may be selectively extracted from different parts of the compressed data stream from locations previously identified, which may represent convenient places at which to reenter the compressed data or may identify places at which data manipulation is anticipated, for example.

According to another embodiment, efficient rotation of Joint Photographic Experts Group (JPEG) compressed data is provided. Reentry information into the JPEG compressed data is identified. An order in which Minimum Coded Units (MCUs) of the JPEG compressed data will be used is determined. Then, the reentry information is used to decode output MCUs from the JPEG compressed data. Finally, rotation is performed in the transform domain by transposing coefficients and appropriately negating certain coefficients' signs.

According to yet another embodiment, compressed image data is printed in a novel way. The compressed image data is received in the form of a compressed data stream. Reentry information into the compressed data stream is generated. Then, the reentry information is used to decode one or more selected pieces of compressed image data to generate corresponding regions of reconstructed data. Finally, the regions of reconstructed data are printed.

According to another embodiment, a compressed data stream may be decoded by multiple parallel decoders by initiating decoding of the compressed data stream by the multiple parallel decoders based upon reentry information into the compressed data stream.

Other features of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 is a simplified block diagram of a typical (prior art) encoder.

FIG. 2 is a simplified block diagram of a typical (prior art) decoder.

FIG. 3 is a simplified block diagram of a typical (prior art) transcoder.

FIG. 4 is a simplified block diagram depicting a local area network (LAN) printing environment.

FIG. 5 is an example of a typical computer system upon which one embodiment of the present invention or components thereof may be implemented.

FIG. 6 is a flow diagram illustrating high-level reentry processing of a compressed data stream according to one embodiment of the present invention.

FIG. 7 is a block diagram illustrating an exemplary decoder architecture for generating reentry information according to one embodiment of the present invention.

FIG. 8 is a block diagram illustrating an exemplary encoder architecture for generating reentry information according to another embodiment of the present invention.

FIG. 9 is a block diagram illustrating an exemplary architecture for manipulating and/or transcoding a compressed data stream according to one embodiment of the present invention.

FIG. 10 is a simplified block diagram of a prior art DCT-based encoder.

FIG. 11 is a simplified block diagram of a prior art DCT-based decoder.

FIG. 12 is a simplified block diagram that conceptually illustrates the result of rotating an image by 90°.

FIG. 13 is a flow diagram illustrating a prior art approach for rotation of compressed images.

FIGS. 14A–14I illustrate the order in which portions of the image are processed according to the prior art rotation approach of FIG. 13.

FIG. 15 illustrates the transposition of coefficients within each block for transform domain rotations.

FIG. 16 is a table of the zig-zag scan ordered coefficients and their respective signs to accomplish 0°, 90°, 180°, and 270° rotation.

FIG. 17 is a flow diagram illustrating a JPEG compressed image rotation processing technique that employs reentry according to one embodiment of the present invention.

FIG. 18 is a flow diagram illustrating a JPEG compressed image rotation processing technique that employs reentry according to another embodiment of the present invention.

FIGS. 19A–19D conceptually illustrate the processing of FIGS. 17 and 18 for a 270° rotation.

FIGS. 20A–20D conceptually illustrate the processing of FIGS. 17 and 18 using an optimization for 90° rotation.

FIGS. 21A–21D conceptually illustrate substitution processing according to one embodiment of the present invention.

FIG. 22 conceptually illustrates parallel decoding processing according to one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Apparatus and methods are described for entering compressed data streams at selected reentry points thereby allowing efficient manipulation of the compressed data and minimizing storage requirements. Because of the absolute torrent of data that requires processing in a very short computational window, efficient manipulation of compressed images, such as rotation, is important in the context of large-scale digital color printing, for example.

In brief, according to one embodiment of the present invention a two stage process may be employed in processing a compressed data stream. Rather than simply sequen-

tially processing a compressed data stream according to the order in which it appears, a two stage technique is employed which allows portions of the reconstructed data to be extracted from different parts of the compressed data stream. During a reentry information generation stage, a scan is made through the compressed data stream to identify locations in the compressed data and preserve associated state information to allow subsequent decoding to resume at those locations. The locations identified may represent convenient places at which to reenter the compressed data or they may identify places at which data manipulation is anticipated. For example, reentry information may be provided on a periodic basis, at particular intervals, or reentry information may simply be generated when conditions are considered favorable. During a second stage, the reentry information generated by the first stage is used. Since there are multiple reentry points into the compressed data stream, the entire reconstructed data need not be stored since the reentry points may be used to selectively access desired reconstructed data. Additionally, access to the compressed data stream is no longer restricted to sequential, rather the compressed data stream may be accessed and decoded in an order that is appropriate for a particular data manipulation operation.

In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a 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 some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

The present invention includes various steps, which will be described below. The steps of the present invention may be performed by hardware components or may be embodied in machine-executable instructions, which may be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware and software.

The present invention may be provided as a computer program product that may include a machine-readable medium having stored thereon instructions that may be used to program a computer (or other electronic devices) to perform a process according to the present invention. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet or optical cards, flash memory, or other type of media/machine-readable medium suitable for storing electronic instructions. Moreover, the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

For convenience, various embodiments of the present invention will be described with reference to the JPEG data stream. However, the present invention is not limited to any particular compressed data format. In alternative embodiments, various other compressed data streams may be employed, such as satellite data transmissions, audio, e.g., MPEG audio, video, e.g., MPEG-1 and MPEG-2 video, speech, text, e.g., LZW, and image compressed formats, such as that produced by ITU-T T.4 Group 3 two dimensional coding standard for facsimile (also referred to as G3), and ITU-T T.6 Group 4 two dimensional coding standard for facsimile (also referred to as G4).

Also, while embodiments of the present invention will be described with reference to two particular entropy coding techniques, e.g., Huffman coding and Arithmetic coding, other entropy coding techniques may be employed. Finally, while only a limited number of presentation devices are specifically referred to herein, such as printers, the method and apparatus described herein are equally applicable to other presentation devices that are capable of producing images on a physical medium, such as displays, facsimile devices, copiers, scanners, and the like. Additionally, the method and apparatus described herein are applicable to other intermediate devices that are capable of retransmission, such as a server, router, bridge, print server, or a base station for satellite or cellular communications, for example.

While embodiments of two-level processing are described in terms of entropy decoders and encoders and models as the two levels, it is contemplated that other one, two or more level processing may be employed. Less than the full number of levels may need to be undone before the manipulation process in which case the encoder would be replaced with a process that redoes as much as the decoder had to undo. An example of a two-level data compression coding scheme without entropy coding could be based upon a hierarchical scheme. For a graphics source image (i.e., computer generated without noise), a 2:1 along each axis reduced version of the graphics source image is first sent in the data stream. The full original-sized image is predicted from a bilinear interpolation 2:1 scale-up in both directions from the reduced image. Each prediction difference can be stored in a byte (i.e., 8-bit modulo arithmetic on the 9-bit prediction difference). Then, along the horizontal axis, traditional run-coding of constant values works well since constant errors in the original image will give zero differences after interpolation of the reduced image. This run-coding does not have to be performed by an entropy coder. Rather, the coding scheme may use fixed-length codes for each run. Alternatively, it might replace repeats of four or more differences that are the same with three differences and the fourth byte interpreted as how many more (from 0-255) will be the same. In this example, the partial decoder removes the run-length codes. After rotation, the partial encoder run-length codes on the new axis, but all of the prediction, interpolation, scale-up, etc. is not affected by the rotation.

Finally, for convenience, embodiments of the present invention will be described with reference to rotation of compressed images. However, such examples are merely illustrative of the power and efficiencies that can be achieved by employing the techniques of the present invention. The applicability of the present invention is not limited to any particular type of data manipulation. In alternative embodiments, various other data manipulations may take advantage of the techniques described herein, such as clipping, substitution, merging, editing, dubbing, etc.

The verbs "print" and "output" in the context of this application, refer broadly to the act of communicating a data stream containing data that is destined for a presentation device or something capable of retransmitting the data stream to a presentation device. These definitions include, but are not limited to, the traditional meaning of print, e.g., sending data to a printer in order to produce character shapes, graphics pictures, images, or other symbols with ink, toner, or the like on a physical medium, such as paper. According to these definitions, data may be said to be "printed" or "output" to various types of presentation devices, such as printers, facsimile devices, fax servers, email servers, pagers, televisions, file viewers, copiers, and

other devices that are capable producing images on a physical medium. As used herein, the phrase “compressed data” generally refers to any type of encapsulated data that cannot normally be reentered without additional side information or pointers at a granularity finer than a byte.

Encoding, Decoding, and Transcoding Overview

FIG. 1 is a simplified block diagram of a typical (prior art) encoder 110. Source data 100 is input to the encoder 110. The encoder model 120 converts the source data 100 into descriptors 130. The descriptors 130 are input into an entropy encoder 140. Examples of entropy encoders include Huffman encoders and arithmetic coding encoders. The output of the entropy encoder 140 is compressed data 150.

FIG. 2 is a simplified block diagram of a typical (prior art) decoder 210. Each step shown performs essentially the inverse of its corresponding main procedure within the encoder shown in FIG. 1. The compressed data 250 is input to the decoder 210 where it is first processed by an entropy decoder 240 which recovers the descriptors 230 from the compressed data 250. This part of the process is typically lossless. The descriptors 230 are input into a decoder model 220 that outputs reconstructed data 200 which may exactly match the source data 100 (lossless) or only approximate the source data 100 (lossy).

FIG. 3 is a simplified block diagram of a typical (prior art) transcoder 310. The compressed data 300 is input to the transcoder 310 where it is first processed by an entropy decoder 340 (or other partial decoder) that recovers the descriptors 330 from the compressed data 300. The descriptors 330 are then input into an entropy encoder 320 (or other partial reencoder). The output of the entropy encoder 320 is compressed data 350. Examples of transcoding are converting between Huffman and arithmetic coding for DCT-based JPEG.

Exemplar Presentation Environment

A simplified printing environment will briefly be described with reference to FIG. 4. In this example, a personal computer workstation 400 is coupled to a print server 420 via a LAN 410. The print server 420 includes a spooler 430 for controlling the spooling of data files and presentation services 440 for generating appropriate commands to drive an attached printer 450. The print server 420 may also include other components that are not shown for performing basic tasks, such as monitoring and configuring attached printers, and providing print job management. At any rate, when the PC workstation 420 has data to print, it sends print data to the print server 420. Among the functions typically provided by a print server is the conversion of the data stream containing the print data to a data stream supported by the printer to which the print data is destined. For instance, the printer 450 may accept the Intelligent Printer Data Stream (IPDS), PostScript, or some other printer data stream. Therefore, in this example, the print server 420 also includes a means for converting between the various input data streams that may be received and the data streams accepted by the printers 450. The print server 420 may also be configured to perform manipulation of compressed images, such as rotation, to allow more efficient processing by the printer 450. Alternatively, the compressed image manipulation processing described herein may be performed local to the printer 450, e.g., by the printer controller 460.

An Exemplary Computer Architecture

Having briefly described an exemplary environment in which the present invention may be employed, an exemplary

machine in the form of a computer system 500 in which features of the present invention may be implemented will now be described with reference to FIG. 5. Computer system 500 may represent a workstation, host, server, print server, or printer controller. Computer system 500 comprises a bus 501 for communicating information, and a processing means such as processor 502 coupled with bus 501 for processing information. Computer system 500 further comprises a random access memory (RAM) or other dynamic storage device 504 (referred to as main memory), coupled to bus 501 for storing information and instructions to be executed by processor 502. Main memory 504 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor 502. Computer system 500 also comprises a read only memory (ROM) and/or other static storage device 506 coupled to bus 501 for storing static information and instructions for processor 502.

A data storage device 507 such as a magnetic disk or optical disc and its corresponding drive may also be coupled to bus 501 for storing information and instructions. Computer system 500 can also be coupled via bus 501 to a display device 521, such as a cathode ray tube (CRT) or Liquid Crystal Display (LCD), for displaying information to an end user. Typically, an alphanumeric input device 522, including alphanumeric and other keys, may be coupled to bus 501 for communicating information and/or command selections to processor 502. Another type of user input device is cursor control 523, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 502 and for controlling cursor movement on display 521.

A communication device 525 is also coupled to bus 501. Depending upon the particular presentation environment implementation, the communication device 525 may include a modem, a network interface card, or other well known interface devices, such as those used for coupling to Ethernet, token ring, or other types of physical attachment for purposes of providing a communication link to support a local or wide area network, for example. In any event, in this manner, the computer system 500 may be coupled to a number of clients and/or servers via a conventional network infrastructure, such as a company's Intranet and/or the Internet, for example.

The present invention is related to the use of computer system 500 to direct the execution of one or more software and/or firmware routines to manipulate compressed images as discussed herein. As computer system 500 executes the one or more routines, the processor 502 may access compressed image data stored within main memory 504, ROM 506, or another storage device to manipulate the compressed image in accordance with desired presentation attributes. Importantly, the present invention is not limited to having all of the routines located on the same computer system. Rather, individual objects, program elements, or portions thereof may be spread over a distributed network of computer systems. Additionally, it is appreciated that a lesser or more equipped computer system than the example described above may be desirable for certain implementations. Therefore, the configuration of computer system 500 will vary from implementation to implementation depending upon numerous factors, such as price constraints, performance requirements, and/or other circumstances. For example, according to one embodiment of the present invention, an embedded printer controller may comprise only a processor and a memory for storing static or dynamically loaded instructions and/or data.

Reentry Processing Overview

FIG. 6 is a flow diagram illustrating high-level reentry processing of a compressed data stream according to one embodiment of the present invention. In one embodiment, the processing described below may be performed under the control of a programmed processor, such as processor 502. However, in alternative embodiments, the processing may be fully or partially implemented by any programmable or hardcoded logic, such as Field Programmable Gate Arrays (FPGAs), TTL logic, or Application Specific Integrated Circuits (ASICs), for example.

At processing block 610, reentry information into a compressed data stream is generated. As discussed below, the reentry information may be generated by an encoder, a decoder, a reentry decoder, an entropy decoder, or an entropy encoder. The purpose of the reentry information is to allow a decoder to initiate decoding at specified locations in the compressed data stream. This is very useful, for example, in connection with parallel decoding and also for minimizing local storage by being able to extract regions of reconstructed data from the compressed data instead of having to store the entire reconstructed data.

Preferably, the reentry information includes optional state information to allow the decoder to resume decoding and a pointer into the compressed data stream that identifies the location at which decoding is to resume. This pointer often has granularity finer than byte boundaries. This reentry information is particularly useful when the compressed data does not contain restart points, such as JPEG restart markers and G3 end-of-line codes. This also avoids extra overhead in the compressed data as the reentry information can be generated and used locally in the decoder. Importantly, in a closed environment, codes to separate independently decodable pieces of compressed data may be inserted arbitrarily (i.e., violate the rules of the data compression standard, such as JPEG restart markers on other than fixed restart intervals). State information may be needed to correctly interpret the results. State information in the context of markers may be as simple as the location of the marker within the compressed data stream. For illegally placed markers, state information may also include the location in the reconstructed image of independently decodable pieces of compressed data.

At processing block 620, the reentry information is used to decode one or more selected pieces of compressed data to generate corresponding regions of reconstructed data. The encoder may have saved pointers and associated state information into the compressed data stream periodically in the compressed data or periodically in the source data. The decoder, however, knows the relevant pieces of information and preserves only what it needs. For example, the decoder may selectively discard certain pieces of the compressed data stream and save reentry information to initiate decoding into each of the remaining pieces of the compressed data stream. This has the advantage that the decoder does not have to serially decompress the compressed data from the start of the compressed data stream. Rather, the decoder may skip to the closest pointer preceding the relevant piece of compressed data and start decoding at the point. The decoder is not restricted to accessing the information from the pointers in sequential order. This provides limited random access into the compressed data stream.

Decoder Architecture

FIG. 7 is a block diagram illustrating an exemplary decoder architecture for generating reentry information according to one embodiment of the present invention. In

one embodiment, the components described below may be implemented as one or more software or firmware routines. However, in alternative embodiments, the components may be fully or partially implemented by any programmable or hardcoded logic, such as Field Programmable Gate Arrays (FPGAs), TTL logic, Application Specific Integrated Circuits (ASICs) or other special purpose hardware, for example.

In this example depicted, the decoder architecture 710 includes an optional reentry decoder 730, a compressed data buffer 720, a decoder 740, and a reentry information buffer 750. Although in the embodiment depicted the reentry decoder 730, the compressed data buffer 720, the decoder 740, and the reentry information buffer 750 are shown as separate elements, various other implementations may combine one or more of these elements.

According to one embodiment, the decoder 740 may generate reentry information into the compressed data stream as it decodes and reconstructs data 760. Alternatively, the decoder 740 may use reentry information 750 previously generated by the optional reentry decoder 730 to identify appropriate pieces of compressed data in the compressed data buffer 720 and reconstruct the desired regions of reconstructed data 760.

When the decoder 740 reconstructs the data the first time through the compressed data in the compressed data buffer 720, it may store reentry information at predetermined or advantageous locations, such as buffer boundaries, in the compressed data stream. If the processing stage following the decoder 740 needs access to one or more regions of the reconstructed data 760, the decoder 740 may regenerate the desired reconstructed data by using the reentry information in the reentry information buffer 750 to access the corresponding compressed data from the compressed data buffer 720. Advantageously, this means the processing stage following the decoder 740 may be designed with less than worst case buffering because it can request desired regions of reconstructed data at less than the cost of a full decode from the start of the compressed data stream to the end of the desired region. For example, an image manipulation stage following the decoder 740 may provide optional control information 770 to the decoder 740 requesting reconstructed data 760 that the image manipulation stage discarded because its buffers were too small. Then, when the image manipulation stage is ready for the discarded information and has some free buffers, the decoder 740 may use the reentry information to generate the requested reconstructed data 760. One particularly useful environment for such regeneration would be data transmission between two devices. In this case, if the reconstructed data 760 is received by the destination in error, e.g., a packet is corrupted or lost during transmission, the decoder 740 may resend the reconstructed data 760 by using the reentry information to decode the relevant pieces of compressed data without having to store the reconstructed data and without having to restart from the beginning of the compressed data 700.

As will be described further below, one use of the reentry decoder 730 may be to facilitate parallel decoding by multiple decoders. The reentry decoder 730 may be a special purpose decoder that only entropy decodes the compressed data stream and stores reentry information at predetermined or advantageous locations, such as buffer boundaries, in the compressed data stream. Alternatively, the reentry decoder 730 may be a fully functional decoder and after it completes its reentry decoding processing it switches to become one of the parallel decoders. Importantly, the reentry decoder 730 does not need to process the entire contents of the com-

pressed data buffer **720**. It can stop once it has completed the desired reentry information and stored it in the reentry information buffer **750**.

According to one embodiment, the reentry decoder **730** may generate a coarse grid of reentry information. The decoder **740** may take the closest preceding coarse reentry point and while decoding save a finer grid of reentry information. The decoder **740** may also at any time supplement the current reentry information. For example, JBIG has SDRST markers that allow independent decoding from these byte-aligned markers. While decoding, either or both of the reentry decoder **730** and the decoder **740** may save additional reentry information which, for example, could allow decoding between these markers.

Encoder Architecture

FIG. **8** is a block diagram illustrating an exemplary encoder architecture for generating reentry information according to another embodiment of the present invention. In this example, the encoder architecture **810** includes an encoder **820**, a reentry information buffer **830**, and an optional reentry information selector **840**. Although in the embodiment depicted the encoder **820** and the reentry information buffer **830** are shown as separate elements, various other implementations may combine these elements.

According to one embodiment, the encoder **820** may generate compressed data and reentry information **850** into the compressed data stream as it encodes source data **800**. Alternatively, the optional reentry information selector **840** may select a subset of the reentry information generated by the encoder **820** for storage. An example of this may be the encoder **820** generating reentry information for every line and the reentry information selector **840** choosing reentry information for G4 MMR compressed data for which the history line is all white and the current line is not all white thereby minimizing cost of storing the reentry information. A reentry decoder may start decoding at the closest reentry information saved by the encoder and save additional reentry information such as the location of the first data that the printer needs to process. Thus, the full decode can be avoided until the relevant region of source data has been found.

Architecture for Manipulation and/or Transcoding

FIG. **9** is a block diagram illustrating an exemplary architecture **910** for manipulating and/or transcoding a compressed data stream according to one embodiment of the present invention. The embodiment shows compressed data **900** with optional reentry information input into an entropy decoder **920** (or other partial decoder). After optional manipulation by the manipulation processing **930**, the entropy encoder **940** (or other partial reencoder) outputs compressed data **950** with optional reentry information. If the reentry information is present in compressed data **900**, the entropy decoder does not have to serially decode the compressed data **900**. In an example of straight transcoding, the entropy encoder **940** may chose to refer to the same regions of reconstructed data with its reentry information. Whereas, if the reentry information is related to compressed buffer boundaries, the entropy encoder **940** will typically refer to different regions.

DCT-based Encoder/Decoder

FIG. **10** is a simplified block diagram of a prior art DCT-based encoder. Source image data **1000** sampled data in 8x8 blocks is input to a DCT-based encoder **1010**. Each 8x8 block is transformed by the FDCT **1021** into a set of 64 values, referred to as the DCT coefficients. One of these

values is referred to as the DC coefficient, and the other 63 values are referred to as AC coefficients. Each of the 64 coefficients is then quantized by quantizer **1022** using one of 64 corresponding values from a quantization table **1023**. After quantization, the DC coefficient and the 63 AC coefficients are output as descriptors **1030** from the encoder model **1020** and input into the entropy encoder **1024**. Inside the entropy encoder **1024**, the previous quantized DC coefficient is used to predict the current DC coefficient and the difference is encoded. The 63 AC coefficients, however, are not differentially encoded but, rather, are converted into a zig-zag sequence. One of two entropy encoding procedures may be used by the entropy encoder **1024**, Huffman encoding or arithmetic encoding. If Huffman encoding is used, then Huffman table specifications are provided by the table specifications **1025**, but if arithmetic encoding is used, then arithmetic coding conditioning table specifications are provided by table specifications **1025**. The output of the entropy encoder **1024** is the compressed data **1040**.

FIG. **11** is a simplified block diagram of a prior art DCT-based decoder. Each processing block shown performs essentially the inverse of its corresponding main procedure within the encoder shown in FIG. **10**. The compressed image data **1140** is input to the DCT-based decoder **1110** where it is first processed by an entropy decoder **1121** which decodes the zig-zag sequence of the quantized DCT coefficients. This is done using either Huffman table specifications or arithmetic coding conditioning table specifications **1122** which were extracted from the compressed data **1140**, depending on the coding used in the encoder. The entropy decoder **1121** also recovers the DC value from the difference and the DC prediction. Then, the descriptors **1130** (i.e., the quantized DCT coefficients output from the entropy decoder **1121**) are input to the dequantizer **1123** which, using table specifications **1124** (the quantization table extracted from the compressed data **1140**), outputs dequantized DCT coefficients to IDCT **1125**. The output of the decoder model **1120** is reconstructed image data **1100**.

In FIGS. **10** and **11**, the entropy encoder **1024** and the entropy decoder **1121** were defined to make transcoding at the descriptor-level **1030** and **1130** convenient. This same definition may be used for reentering the data stream if the reentry information saves the DC predictor per component and has a bit-level pointer into the compressed data for an MCU boundary for JPEG and a bit-level pointer to a macroblock for MPEG. In some embodiments, it may also be convenient to store the location of the MCU or macroblock. Saving reentry information for JPEG MCU boundaries or MPEG macroblock boundaries minimizes storage for the reentry information while allowing reentry into the compressed data at points that are not required to be byte-aligned. For the cost of additional overhead in saving the reentry information, this approach allows arbitrary reentry into the compressed data.

Exemplary Manipulation Processing

FIG. **12** is a simplified block diagram that conceptually illustrates the result of rotating an image by 90°. In FIG. **12**, an input image **1200** and an output image **1210** are depicted. The output image **1210** corresponds to a 90° (clockwise) rotation of the input image **1200**. The input image **1200** is shown divided into nine portions **1201–1209**. While portions **1201**, **1202** and **1203** are the first encountered in a sequential scan of the input image **1200**, portion **1203** is the last portion of the output image **1210** to be output. Consequently, as illustrated by FIGS. **14A–14I**, image rotation

processes must accumulate the whole image in memory or on disk before the rotated image can be output.

FIG. 13 is a flow diagram illustrating a prior art approach for rotation of compressed images. At step 1310, a compressed image to be rotated is received. At step 1320, the compressed image is entropy decoded. Then, at step 1330, inverse quantization and inverse DCT are performed and the output is stored. The rotation is performed by reordering the image samples in the real domain at step 1340. Then, in step 1350, the FDCT and requantization are performed. In step 1360, the entropy encoding is performed. Finally, in step 1370, the compressed, rotated image is output.

Notably, many disadvantages arise from this prior art approaches conversion between the transform domain and the real domain. First, it requires both IDCT and FDCT and the accompanying multi-generation problems due to the integer representation in the real domain. Additionally, this approach may require the entire image to be buffered since, as illustrated by FIGS. 14A–14I, the first pixels read by the rotation process contain some of the last pixels to be output. Finally, such an approach is also inefficient in terms of its use of CPU cycles.

FIGS. 14A–14I illustrate the order in which portions of the image are processed according to the prior art rotation approach of FIG. 13. In FIGS. 14A–14I, input image 1400 is sequentially processed to produce a rotated image 1410. In this example, the input image 1400 is rotated by 90 degrees (clockwise) to produce the rotated image 1410. If output of the rotated image 1410 is desired, the first blocks read and processed (e.g., blocks 1401, 1402, and 1403) by the rotation process cannot be immediately output and therefore need to be buffered. The order of processing of input image 1400 is block 1401, block 1402, block 1403, block 1404, block 1405, block 1406, block 1407, block 1408, and block 1409. In contrast, the output order, as shown in FIG. 14I, is block 1407, block 1404, block 1401, block 1408, block 1405, block 1402, block 1409, block 1406, and block 1403. Consequently, as mentioned above, prior art rotation techniques may require worst case buffering of the entire image.

In view of the foregoing, it is desirable to reduce the buffering requirements for such compressed data manipulations. Additionally, it would be advantageous to be able to selectively extract desired pieces of compressed data from a compressed data stream thereby allowing immediate output of processed data. For example, referring back to FIGS. 14A–14I, if blocks 1407, 1404, and 1401 could be accessed and manipulated first by the rotation process, then upon completion of this row, the blocks could be immediately output.

DCT Coefficient Manipulation for Transform Domain Rotations

FIG. 15 illustrates the transposition of coefficients within each block for transform domain rotations. Working with the zig-zag ordering of the DCT coefficients specified by JPEG, rotations by integer multiples of 90 degrees (clockwise) are accomplished as shown in the columns of FIG. 16. The entries in 8×8 block 1510 give the positions of the DCT coefficients in zig-zag scan order (also shown by the zig-zag line in block 1510). The entries in 8×8 block 1520 show the new positions of the DCT coefficients after the transpose. In this example, the transposition essentially comprises moving a DCT coefficient from its present position in block 1510 to a corresponding position in block 1520 on the other side of diagonal 1530. For example, before transposition of the DCT coefficients of block 1510, DCT coefficient 1 is located

in the first row, second column of the block 1510. After transposition, coefficient 1 is located in the second row, first column of block 1520.

Importantly, the transpose operation depicted in FIG. 15 requires less CPU cycles than would be required if the image data was rotated in the real domain. Recall, in order to rotate in the real domain and return to the compressed format, an inverse DCT and forward DCT are required, resulting in image degradation and additional processing. Therefore, one advantage of rotating in the transform domain, e.g., the quantized DCT coefficient domain, is that the quality of the image will be unaffected since it is a lossless process.

FIG. 16 is a table of the zig-zag scan ordered coefficients and their respective signs to accomplish 0°, 90°, 180°, and 270° rotation. In this example, the number in column 1600 represents the DCT coefficient position prior to rotation; and the values in columns 1610, 1620, and 1630 represent where that DCT coefficient that will be in zig-zag order after 90 degree, 180 degree, and 270 degree rotations, respectively, and whether or not a negation operation is needed. According to the table of FIG. 16, rotations by integer multiples of 90 degrees (clockwise) are accomplished by transposing the DCT coefficients of block 1510 and negating zero or more of the transposed DCT coefficients. For example, the DCT coefficient in position 1 of the zig-zag order prior to a 90 degree rotation will move to position 2 and be negated.

Reentry JPEG Compressed Image Processing Overview

A technique for processing a compressed image will now be briefly summarized with reference to rotation of JPEG compressed images. According to one embodiment, rather than simply sequentially processing MCUs according to the order in which they appear in the input data stream, MCUs are extracted from different parts of the compressed image in accordance with the order in which the MCUs will appear in the output data stream. During a reentry information generation stage, the compressed image is partially entropy decoded to identify pointers, e.g., offsets, to each desired reentry point, e.g., at MCU boundaries. Importantly, since MCUs are not necessarily byte aligned in the coded data stream, it should be understood that these pointers are bit-level pointers into entropy encoded image data. That is, the pointers contain enough resolution to identify the byte and bit position within the byte of the start of the code for the desired MCUs. Other state information may also be recorded as necessary during this stage. For example, in the case of a JPEG image, history information, such as DC predictors, are needed because DC values are differentially encoded. Consequently, determining the DC value for a current block depends upon the preceding DC value within the scan that is coded for the same component. Therefore, in the case of a JPEG image, a predictor per component is also recorded during this stage. During a second stage through the compressed image, e.g., a manipulation stage, such as a rotation stage, decoding may be restarted at each saved pointer. The compressed image is processed in such an order that the resulting MCUs, when rotated, can be assembled into a proper rotated image. Preferably, the order of processing also allows rendering of the output image to begin immediately, for example, by assembling the output image in a top-down order. Advantageously, in this manner, there is no longer a need to accumulate the whole input image in memory before beginning to output the rotated image.

Server-side, Reentry JPEG Compressed Image Rotation Processing

FIG. 17 is a flow diagram illustrating a JPEG compressed image rotation processing technique that employs reentry for

use on a print server according to one embodiment of the present invention. In this example, the rotation process generally breaks down into a reentry information identification stage **1701** and a reentry information usage stage **1702**. The identification stage **1701** is represented by processing block **1710** and the usage stage **1702** includes processing blocks **1720–1780**.

The rotation method begins at processing block **1710**, where a compressed image is sequentially read and reentry information into the compressed image is identified by the print server. This reentry information may be generated by a reentry decoder that generates the exact list of reentry information. Alternatively, the reentry information may have been saved with the compressed data by the encoder. Also it is possible to provide the reentry information concurrently with the compressed data. In another embodiment, the reentry information may have been previously generated during a prior rotation and saved locally in anticipation of subsequent reuse. As shown in FIGS. **19** and **20**, 90 degree counterclockwise rotation requires more reentry information than 90 degree clockwise rotation. In such a situation, the identification stage **1701** may selectively discard unnecessary reentry information points thereby retaining in local storage only those that will be needed in the usage stage **1702**.

At processing block **1720**, a determination is made with regard to the order in which MCUs will be read from the compressed image based upon the reentry information identified and the particular rotation desired.

At processing block **1730**, the next MCU to be placed into the output image is entropy decoded. Note, as explained with reference to FIG. **11**, the output of the entropy decoder **1121** is the complete DC coefficient (prediction has been removed).

At processing block **1740**, the MCU is rotated in the transform domain. It is then re-entropy encoded at processing block **1750**. Any completed compressed bytes are output at processing block **1760**. Preferably, the output is in MCU row order thereby allowing immediate transfer to the printer controller.

At processing block **1770**, a determination is made whether or not this is the last MCU of the output image. If so, encoding is terminated by flushing the encoder and outputting the final bytes of compressed data at processing block **1780**. Otherwise, processing returns to processing block **1730**.

Advantageously, in this manner, the printer controller is able begin decoding in MCU row order as soon as it receives the first compressed byte of the rotated image from the server. Notably, in the prior art, effectively the whole rotated transform coefficients need to be buffered before recompression for output to the printer controller can begin.

Note that in this example it is assumed that the rotation will not cause unwanted data to remain in the picture (e.g., the printer can discard the rotated padding if it is no longer on the right or the bottom of the image).

Printer Controller Reentry JPEG Compressed Image Rotation Processing

FIG. **18** is a flow diagram illustrating a JPEG compressed image rotation processing technique that employs reentry for use in a printer controller according to another embodiment of the present invention. As above, in this example, the rotation processing process generally breaks down into a reentry information identification stage **1801** and a reentry information usage stage **1802**. The identification stage **1801**

is represented by processing block **1810** and the usage stage **1802** includes processing blocks **1820–1880**.

The rotation method begins at processing block **1810**, where a compressed image is sequentially read and reentry information into the compressed image is identified by the printer controller. This reentry information may be generated by a reentry decoder that generates the exact list of reentry information. Alternatively, the reentry information may have been saved with the compressed data by the encoder. Also it is possible to provide the reentry information concurrently with the compressed data. In another embodiment, the reentry information may have been previously generated during a prior rotation and saved locally in anticipation of subsequent reuse.

At processing block **1820**, a determination is made with regard to the order in which MCUs will be read from the compressed image based upon the reentry information identified and the particular rotation desired.

At processing block **1830**, the next MCU to be placed into the output image is entropy decoded. Note, as explained with reference to FIG. **11**, the output of the entropy decoder **1121** is the complete DC coefficient (prediction has been removed).

At processing block **1840**, the MCU is optionally rotated in the transform domain. Dequantization and IDCT is then performed to return to the real domain at processing block **1850**. At processing block **1860**, rotation is performed in the real domain if rotation was not performed in processing block **1840**. While blocks **1850** and **1860** are shown separately in this example, it is contemplated that these two processing blocks could be combined thereby allowing rotation and IDCT to be performed simultaneously. At processing block **1870**, the real domain data is output. At processing block **1880**, a determination is made whether or not this is the last MCU to be processed. If so, rotation is complete. Otherwise, processing returns to processing block **1830**.

Reentry JPEG Compressed Image Manipulation Processing Illustrations

FIGS. **19A–19D** conceptually illustrate the processing of FIGS. **17** and **18** for a 270° rotation. In FIG. **19A**, a source image **1910** is illustrated having blocks **1901** to **1909**. The source image **1910** is encoded and reentry information is generated by either an encoder or by a reentry decoder. The saved reentry information (represented by arrows) includes a bit-level pointer into the compressed data **1920** and the DC prediction for the block. The encoding and/or reentry information generation may be performed at different points in time. For example, instead of encoding source image **1910** and generating reentry information in the encoder, a reentry decoder could operate on the compressed image data to generate the reentry information at a later time. For convenience, the compressed data for each block is shown as the same size as the source blocks **1901** to **1909** and identified with the same number.

Assuming the compressed data **1920** and the reentry information are contained in storage and rotated output is desired and following the method described with regard to FIG. **18**, e.g., printer controller processing, the decoder selectively extracts pieces of compressed data according to the desired output order, entropy decodes the next piece, performs rotation in the transform domain (following FIGS. **15** and **16**), dequantizes and performs the IDCT to output the rotated image data. Ultimately, the rotation in the transform domain could be skipped and the rotation could be performed in the real domain as described above. In this

illustration, the order in which the pieces of compressed data are selected is as shown by the reentry points labeled **1991** to **1999**. Therefore, the decoder starts decoding at reentry point **1991** and decodes block **1903** outputting it in rotated order. At the end of block **1903**, the decoder skips to reentry point **1992** and proceeds to decode and output block **1906**. This process is repeated until the rotated image is output.

Alternatively, if the desired output is a rotated image in compressed form, e.g., server processing, as described with respect to FIG. **17**, then after identifying the reentry points **1991** to **1999** in the desired order, the decoder selectively extracts pieces of compressed data according to the desired output order, entropy decodes the next piece, performs rotation in the transform domain (following FIGS. **15** and **16**), entropy encodes the rotated, quantized, transform, data, i.e., MCU, and terminates the encoding after all the MCUs have been rotated.

FIGS. **20A–20D** conceptually illustrate the processing of FIGS. **17** and **18** using an optimization for 90° rotation. In this example, the overhead for storing reentry information may be reduced by taking advantage of the fact that the end of decoding a particular piece of compressed data is the start of a piece of compressed data that will be accessed later. Consequently, less initial reentry pointers may be stored. For example, after competing entropy decoding of block **2007**, the reentry information to later entropy decode block **2008** is available and may replace the reentry information including reentry point **2091**, which is why reentry point in FIG. **20C** is labeled both **2094** and **2091**. Similarly, after competing entropy decoding of blocks **2004** and **2001**, the reentry information to later entropy decode blocks **2005** and **2002**, respectively, is available and may replace the reentry information including reentry points **2092** and **2093**. The remainder of the process is as described above with reference to FIGS. **19A–19D**.

FIGS. **21A–21D** conceptually illustrate substitution processing according to one embodiment of the present invention. In this example, reentry information into a first image **2110** in compressed form is used to take independent blocks of image data from the first image **2110** and to merge them into a second image **2130** on block boundaries. The concept works for more complex merging where overlapped partial blocks require more complex processing.

In a scenario where the second image **2130** is already decompressed and in storage, there is no need to decode the unused pieces of reconstructed data from the first image **2110**. FIG. **21B** illustrates that the first image **2110** exists in compressed form **2120** with periodic reentry information. FIG. **21C** shows that the closest reentry information **2191** to the desired block **2184** enables decoding of blocks **2181** and **2182** to be skipped. Either the decoder may initiate decoding at reentry information **2191** or a reentry decoder may find the precise point to start decoding at **2184**. The output replaces block **2104**. In FIG. **21D**, the reentry information **2193** is already correctly located and the rest of the compressed data may be decompressed and the output replaced.

When substitution of pieces of the first image **2110** are to be substituted into a compressed version of the second image **2130**, then it is helpful to have reentry information into both compressed data streams. Additionally, the entropy code of the first substituted DC term needs to be corrected to reflect the predictor differences between the images at that point.

In an alternative embodiment, such a substitution can be made more efficient by generating reentry information precisely where needed.

Importantly, while this example has been illustrated with two-dimensional blocks, a similar substitution process may be used for dubbing one-dimensional audio data.

Parallel Decoding

FIG. **22** conceptually illustrates parallel decoding processing according to one embodiment of the present invention. In this example, compressed data **2210** is input into a reentry decoder **2220**. The reentry decoder **2220** generates reentry information **2241**, **2242**, and **2243** into the compressed data **2210** that allows pieces of the compressed data **2211**, **2212**, and **2213** to be independently passed to multiple parallel decoders **2250**. While in this example, a one-to-one correspondence exists between the number of pieces of compressed data **2211**, **2212**, and **2213** and the number of parallel decoders **2250**, in a preferred embodiment, fewer parallel decoders **2250** may be employed and each may simply request another piece of compressed data upon completion of decoding of its piece. Importantly, the encoder and the compressed data are not affected. Rather, the reentry decoder **2220** may locally decide how many and what size to make the pieces of compressed data.

When there are fewer parallel decoders **2250** than pieces of compressed data **2211**, **2212**, and **2213**, the pieces may be treated as a pool from which each parallel decoder **2250** pulls. That is, the parallel decoders **2250** decode from their current reentry starting point to the next reentry starting point and stop. Then, they request the next available reentry point from the pool.

Selecting Reentry Information to Minimize Storage

For the LZW text compression algorithm, the maximum size of the dictionary may be set in advance. New leaves are added until the maximum size is reached. The already generated leaves may be continued to be used until the content changes sufficiently so that the compression degrades. Then a “clear” code can tell the decoder that it should clear out the dictionary and reset to the minimum bits per codeword. The encoder is allowed to reset the dictionary by sending a “clear” code at anytime. Reentry information may include a bit pointer to the leading one bit contained in the clear code (decimal value 256 with a number of leading zeros dependent upon the size of the current codes). The encoder may also record information regarding the count of characters already compressed. Here is an example of where the encoder might choose to output the reentry information every time it encoded a “clear” code. The decoder may search for the closest preceding pointer to the desired data.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A method comprising generating reentry information into a compressed data stream, wherein source data corresponding to the compressed data stream comprises image data that is compressed according to a transform based data compression technique, and wherein the reentry information for an MCU boundary of the compressed data stream comprises a bit-level pointer to the start of a code, the position of the output, and a DC predictor for each component of the MCU.

2. The method of claim **1**, wherein said generating reentry information is performed during the process of decoding to reconstructed data.

23

3. The method of claim 1, wherein said generating reentry information is performed during entropy encoding or decoding of the compressed data stream.

4. The method of claim 1, wherein said generating reentry information is performed during partial decoding of the compressed data stream.

5. The method of claim 1, comprising using the reentry information to decode one or more selected pieces of compressed data to generate corresponding regions of reconstructed data.

6. The method of claim 1, further comprising selecting a subset of the reentry information.

7. The method of claim 1, further comprising using the reentry information to partially decode one or more selected pieces of compressed data to generate intermediate data.

8. The method of claim 7, further comprising creating one or more new pieces of compressed data based upon the intermediate data.

9. The method of claim 7, further comprising creating one or more regions of reconstructed data based upon the intermediate data.

10. The method of claim 1, further comprising using the reentry information to partially decode one or more selected pieces of compressed data to transcode one or more pieces of compressed data into new pieces of compressed data.

11. The method of claim 7, further comprising manipulating the intermediate data to generate manipulated intermediate data.

12. The method of claim 11, further comprising creating one or more new pieces of compressed data based upon the manipulated intermediate data.

13. The method of claim 11, further comprising creating one or more regions of reconstructed data based upon the manipulated intermediate data.

14. The method of claim 1, wherein the transform based data compression technique employs the discrete cosine transform (DCT).

15. The method of claim 1, wherein the image data is compressed according to one of the Joint Photographic Experts Group (JPEG) standards.

16. The method of claim 1, wherein the start of a code is the start of a Huffman code.

17. The method of claim 1, wherein the reentry information for other than an MCU boundary comprises a bit-level pointer to the start of a Huffman code, an indication of an appropriate Huffman table associated with the Huffman code, the position of the output, an indication of the block of the MCU, an indication of which coefficient of the block, and the DC predictors for each component at the current decoding step.

18. The method of claim 1, wherein the reentry information for an MCU boundary comprises a byte-level pointer to the last byte input to the arithmetic decoder, probability index and MPS sense for the relevant contexts, the contents of the A register and the C register, the position of the output, sufficient history, and a DC predictor for each component of the MCU.

19. The method of claim 1, wherein the reentry information for other than an MCU boundary comprises a byte-level pointer to the last byte input to the arithmetic decoder, probability index and MPS sense for the relevant contexts, the contents of the A register and the C register, the position of the output, all indication of the block of the MCU, an indication of which coefficient of the block, sufficient history, and the DC predictors for each component at the current decoding step.

20. The method of claim 1, wherein source data corresponding to the compressed data stream comprises audio data.

24

21. The method of claim 20, wherein the audio data is compressed according to one of the Moving Pictures Expert Group (MPEG) standards.

22. The method of claim 1, wherein source data corresponding to the compressed data stream comprises video data.

23. The method of claim 22, wherein the video data is compressed according to one of the Moving Pictures Expert Group (MPEG) standards.

24. The method of claim 1, wherein source data corresponding to the compressed data stream comprises bi-level image data.

25. The method of claim 24, wherein the bi-level image data is compressed according to the G3 MH standard.

26. The method of claim 25, wherein the reentry information comprises bit-level pointers to the start of a Huffman code, a color of a Huffman table associated with the Huffman code, and an indication of the position of the output.

27. The method of claim 24, wherein the bi-level image data is compressed according to the G3 MR standard.

28. The method of claim 27, wherein the reentry information comprises bit-level pointers to the start of a Huffman code, a color of the output, an indication of the appropriate Huffman table, an indication of whether the line identified by the bit-level pointer was encoded one or two dimensionally, an indication of the position of the output, and sufficient previous line history.

29. The method of claim 24, wherein the bi-level image data is compressed according to the G4 MMR standard.

30. The method of claim 29, wherein the reentry information comprises bit-level pointers to the start of a Huffman code, a color of the output, an indication of the appropriate Huffman table, the position of the output, and sufficient previous line history.

31. The method of claim 24, wherein the bi-level image data is compressed according to Adaptive Bilevel Image Compression (ABIC).

32. The method of claim 31, wherein the reentry information comprises a byte-level pointer to the last byte input to the decoder, probability index and MPS sense for the 128 contexts, the contents of the A register and the C register, the position of the next output, and sufficient previous bits.

33. The method of claim 1, wherein source data corresponding to the compressed data stream comprises scientific data.

34. The method of claim 1, wherein the reentry information includes one or more pointers into the compressed data stream, each of the one or more pointers capable of identifying a point in the compressed data stream with a level of granularity that is finer than a byte boundary.

35. The method of claim 1, wherein the reentry information includes state information associated with each of one or more pointers into the compressed data stream, the state information enabling extraction of information from the compressed data stream at a point identified by the associated pointer.

36. The method of claim 35, wherein the state information enables decoding to be restarted at the point identified by the associated pointer.

37. The method of claim 1, wherein the reentry information corresponds to predetermined locations in source data represented by the compressed data stream.

38. The method of claim 1, wherein the reentry information is chosen to minimize storage required for the reentry information.

39. A method of selectively extracting information from a compressed data stream by using reentry information, source data corresponding to the compressed data stream comprising image data that is compressed according to a transform based compression technique, the reentry infor-

mation including state information associated with each of one or more pointers, the pointers including a bit-level pointer to the start of a code, the position of the output, and a DC predictor for each component of the MCU, and where the reentry information is for an MCU boundary of the compressed data stream, the state information enabling extraction of information from the compressed data stream at a point identified by the associated pointer.

40. A method of selectively extracting information from a compressed data stream by using reentry information, the reentry information for an MCU boundary of the compressed data stream, data corresponding to the compressed data stream comprising image data that is compressed according to a transform based compression technique, including one or more pointers into the compressed data stream, the pointers including a bit level pointer to the start of a code, the position of the output, and a DC predictor for each component of the MCU, and each of the one or more pointers capable of identifying a point in the compressed data stream with a level of granularity that is finer than a byte boundary.

41. The method of claim **40**, wherein the reentry information is available with the compressed data stream.

42. The method of claim **41**, wherein the reentry information is generated during the process of decoding to reconstructed data.

43. The method of claim **41**, wherein the reentry information is generating during entropy encoding or decoding of the compressed data stream.

44. The method of claim **41**, wherein the reentry information is generating during partial decoding of the compressed data stream.

45. A method comprising:

preserving reentry information into a compressed data stream, source data corresponding to the compressed data stream comprising image data that is compressed according to a transform based data compression technique;

extracting desired information from the compressed data stream by

determining a reentry point based upon the reentry information that provides most efficient access to the start of the encoded data corresponding to the desired information, and

restarting decoding of the compressed data stream at a pointer associated with the reentry point, the pointers including a bit-level pointer to the start of a code, the position of the output, and a DC predictor for each component of the MCU, and where the reentry information is for an MCU boundary of the compressed data stream.

46. The method of claim **45**, further comprising restoring state information associated with the reentry point.

47. The method of claim **45**, wherein said preserving reentry information into a compressed data stream comprises: receiving the reentry information; and keeping all or part of reentry information.

48. The method of claim **45**, wherein said preserving reentry information into a compressed data stream comprises:

generating the reentry information; and

storing associated state information, if necessary, and a pointer of the one or more pointers associated with each of the one or more reentry points.

49. The method of claim **48**, wherein said generating comprises decoding the compressed data stream sufficiently

to identify one or more pointers into the compressed data stream at a finer level of granularity than a byte.

50. A machine-readable medium having stored thereon data representing sequences of instructions, said sequences of instructions which, when executed by a processor, cause said processor to generate reentry information into a compressed data stream, wherein source data corresponding to the compressed data stream comprises image data that is compressed according to a transform based data compression technique, and wherein the reentry information for an MCU boundary of the compressed data stream comprises a bit-level pointer to the start of a code, the position of the output, and a DC predictor for each component of the MCU.

51. The machine-readable medium of claim **50**, wherein said sequences of instructions further cause the processor to use the reentry information to decode one or more selected pieces of compressed data to generate corresponding regions of reconstructed data.

52. The machine-readable medium of claim **50**, wherein said sequences of instructions further cause the processor to use the reentry information to partially decode one or more selected pieces of compressed data to generate intermediate data.

53. The machine-readable medium of claim **52**, wherein said sequences of instructions further cause the processor to manipulate the intermediate data to generate manipulated intermediate data.

54. The machine-readable medium of claim **53**, wherein said sequences of instructions further cause the processor to create one or more new pieces of compressed data based upon the manipulated intermediate data.

55. The machine-readable medium of claim **53**, wherein said sequences of instructions further cause the processor to create one or more regions of reconstructed data based upon the manipulated intermediate data.

56. A machine-readable medium having stored thereon data representing sequences of instructions, said sequences of instructions which, when executed by a processor, cause said processor to selectively extract information from a compressed data stream by using reentry information for an MCU boundary of the compressed data stream, source data corresponding to the compressed data stream comprising image data that is compressed according to a transform based compression technique, the reentry information including one or more pointers into the compressed data stream, the pointers including a bit-level pointer to the start of a code, the position of the output, and a DC predictor for each component of the MCU, and each of the one or more pointers capable of identifying a point in the compressed data stream with a level of granularity that is finer than a byte boundary.

57. The machine-readable medium of claim **56**, wherein the reentry information is available with the compressed data stream.

58. The machine-readable medium of claim **56**, wherein the reentry information is generated during the process of decoding to reconstructed data.

59. The machine-readable medium of claim **56**, wherein the reentry information is generating during a partial decoding of the compressed data stream.

60. The machine-readable medium of claim **56**, wherein the reentry information is generating during entropy encoding or decoding of the compressed data stream.