JPEG 2000: Get the Bit You Want

Michael J. Gormish
Ricoh Innovations, Inc.
Menlo Park, California

Abstract

JPEG 2000 is an international standard for the compression of still images.[1] Although the new standard provides better compression at any given quality level, the real advantage over the DCT based "original" JPEG is the feature set. JPEG 2000 can compress different types of image data from black and white graphics to multi-spectral high bit depth imagery at extremely low bit rates or extremely high quality (including lossless). Perhaps more importantly, decisions about compression ratio or quality or region of interest can be made after the image is compressed. This wide range of capabilities make the standard useful in a diverse set of applications including: internet imaging, printing, scanning, digital photography, remote sensing, mobile, color facsimile, medical imagery, and digital libraries.

This paper provides a brief overview technologies used in Part I of the standard and sample applications which make use of those features.

Introduction

For a long time the sole goal of image compression has been transmitting or storing data in the fewest number of bits. Different subfields of image compression addressed specific types of imagery or quality requirements.

Facsimile compression, for example, was concerned with the lossless compression of black and white "document images." Of course documents could contain line art, tables, and halftones and any fax standard would need to handle these, but first and foremost *text,* both hand written and printed, had to be reproduced in an extremely readable format. The technology developed for this application consisted of more and more sophisticated statistical predictors of pixels, efficient methods to encode these predictions in bits, and methods to recover after potential errors in transmission. Special prediction techniques were developed for portions of documents e.g. halftones, and even the method of prediction could be adaptive. The JBIG compression standard is representative of this effort.[2]

Other compression systems were needed for "natural imagery" which was generally of a much higher data rate (8 bits of red, green, and blue data instead of one bit for black/ white). While "document" images could often be compressed losslessly by a large factor (25:1 compression is not atypical), natural images could not be. Natural images typically have a much higher variance and contain noise in the acquisition process which limits lossless compression to about 2:1. Thus development of natural image compression focused on technology for approximating signals (lossy compression/quantization) and on models of the human visual system so that approximations would not be visible. The JPEG standard is the most common of these systems.[3]

JPEG 2000 was spurred on in part by the development of wavelet technology which provided better compression than the discrete cosine transform (DCT) used in the JPEG standard. More importantly one of the goals of JPEG 2000 is to be a single solution to a large range of image compression problems. The standard handles black and white *and* color imagery, the standard provides lossy *and* lossless compression, and the standard allows the same compressed image to be useful with a monitor *and* a printer. These goals are described in the JPEG committee's "Call for contributions"[4] and in the CREW system[5] which led the committee to issue the call.

JPEG 2000 has not displaced JPEG as the standard of choice for lossy image compression. In part this is because JPEG 2000 is more computationally complex and more "feature" complex than JPEG. But primarily it is because JPEG is a good lossy image compression standard. In fact, JPEG is likely to continue as the image compression standard of choice for some applications for the foreseeable future. However, great interest and development effort of JPEG 2000 has gone into those applications which are not well served by JPEG. In general these applications are those where access to portions of an image are important.

During the course of development of JPEG and JPEG 2000, the amount of data storage available for a fixed cost has risen exponentially. Availability of bandwidth in most cases has increased as well, though usually no where near as much as disk storage increases. Because of these increases in storage and bandwidth, an improvement of compression efficiency of 10% (a remarkable feat), is less important in end to end product design. At best a better compression algorithm saves only a few months time vs. waiting for better network and storage capabilities. Of course, during this time the resolution and bit depth of image acquisition devices has also increased in part using up the additional network and storage capacity.

For these reasons, the single most important aspect of JPEG 2000 is arguably the ability to access portions of the compressed data. An application wishing to display a 600 dpi JPEG compressed image on a 75 dpi monitor filling a 1024x1024 region of the screen, must decompress the entire

JPEG file. This JPEG file might be 10 megabytes even if compressed very highly (and thus low quality). A JPEG 2000 decompressor might only need to access 150 kilobytes in order to display an equivalent image. The JPEG 2000 decompressor in this case makes far fewer disk accesses and takes far less CPU time (memory requirements are likely to be similar for both JPEG and JPEG 2000 in this case, both being dominated by the output image buffer).

If the same image is accessed by a printer which desires to print only the lower right hand 1024x1024 portion of the high resolution image, the same comparison occurs. The JPEG image must be decoded entirely, and only a very small although *different* portion of the JPEG 2000 image must be decoded.

The technologies of JPEG 2000 which impact the ability to access a portion of the compressed data are described in the next section of this paper. An applications section describes how those technologies can be used for specific purposes.

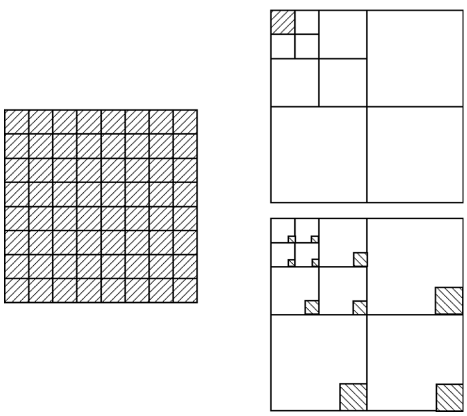

Figure 1. Data accessed to decode low resolution or lower right corner for JPEG image (left). Data accessed to decode low resolution JPEG 2000 image (top right). Data accessed to decode lower right corner for JPEG 2000 image (lower right).

JPEG 2000 Technologies

Wavelets Provide Access to Multiple Resolutions

Use of wavelets in JPEG 2000 instead of the discrete cosine transform of JPEG allowed better compression at low bitrates, and allowed lossless compression (with the use of a reversible wavelet).

Most importantly from an image access point of view is that one application of a forward wavelet transform leads to a high quality image reduced in size by a factor of two both horizontally and vertically. If a low resolution image is desired, only the compressed data corresponding to this subband need to be accessed. For compression purposes, the wavelet transform is repeated several times on the low pass subband, leading to several low resolution images at intermediate points in the compression process.

A decoder may decode only those portions of the data corresponding to the desired resolution. Alternatively, if a decoder was interested only in edge locations, it could decode only the high pass subbands.

Figure 2. Wavelets provide lower resolution images

Tiles, Precincts, and Codeblocks Provide Access to Regions

Transform images may be divided into regular rectangular regions called tiles before the forward wavelet. Each tile is wavelet transformed and compressed independently of the other tiles in the image. Tiles correspond to the same spatial region of the image at all resolutions. The length of the compressed data for every tile-part is contained in the tile-part header, and may optionally be collected in the main header of the codestream. Thus it is easy to skip to the desired tile and decode only the lower right corner for example. For poorly compressed images tiles can lead to rectangular artifacts in addition to the wavelet artifacts.

Precincts divide wavelet coefficients into regular rectangular regions. Once again, to decode the lower right corner a PLT marker segment can be used to skip intermediate data. However, because a single wavelet coefficient affects many pixels it is sometimes necessary to decode precincts beyond the desired region of interest or an image of uneven quality may result.

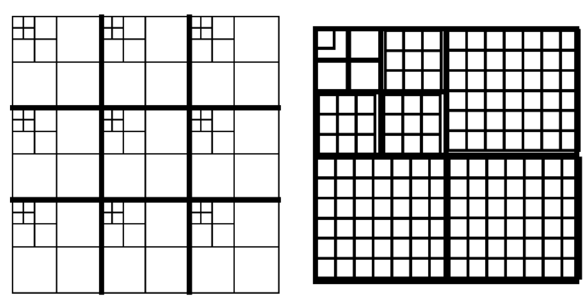

Figure 3. Image divided into tiles then subbands (left), Tile divided into subbands then precincts (right)

Precincts are divided into even smaller units called code blocks. Each code block is compressed independently. Codeblock lengths are stored in the packet headers but they are typically so short that "skipping" codeblocks does not reduce the number of seeks when accessing a disk. Skipping the decoding of unneeded codeblocks does lead to savings in CPU time.

File Format and Syntax Provide Access to Data

Several file formats have been defined for wrapping JPEG 2000 compressed data. These file formats all consist of a set of boxes with different boxes for each type of data stored e.g. copyright information, color space information, audio information. The boxes all begin with four bytes indicating the length of the box, and four bytes indicating the type of the box. The different file formats define new types of boxes and specify which boxes must be present in the file. Because most unknown boxes can safely be skipped, it is possible for a JPEG 2000 file to be a legal file for many different file formats. Incidentally, the box definition is remarkably similar to the "atoms" of Quicktime™ and MPEG-4, allowing JPEG 2000 to make use of several previously defined multimedia boxes/atoms (and other multimedia players to relatively easily add JPEG 2000 support). As far as access is concerned, the important aspect of the file format is the ability to skip boxes which are not of interest after accessing the length and type.

At least one of the boxes in the JPEG 2000 file will contain a "codestream." The codestream contains all the information necessary to recreate sample values. The codestream has a syntax based on the syntax of the original JPEG. There are two byte "markers" which begin with the value 0xFF and have an additional byte indicating the type of marker. Some markers contain two bytes of length information others are of fixed size. Thus some segments of the codestream can be skipped. Many of the marker codes are prevented from appearing in the compressed data, so it is possible to examine bytes looking for a specific marker code, even if the length information has become corrupted. Some optional marker segments (TLM, PLT, and PLM) contain collected length information so that individual lengths need not be read.

Layers Provide Access to Different Quality Images

A JPEG 2000 encoder is permitted to represent different wavelet coefficients to different precision. Indeed including only the top few bit-planes of a wavelet coefficient is the primary manner to obtain compression. Clearly it takes fewer bits to represent and transmit the value 0100 1010 using two bits as 01 than with six bits using 0100 10. Of course if too few bits are used to represent a coefficient a low quality image will result. In JPEG 2000 each codeblock must be represented to the same precision, but different codeblocks can be stored to different precision (within the same subband and across subbands).

In JPEG 2000 an initial precision can be used for coefficients. This becomes the first layer. Each coefficient can be further refined (or not) in other layers. Because each packet contains only one layer, one resolution, one precinct and one component, the marker segments can be used to locate data for the number of layers desired by the decoder. This provides access to low quality (high compression) or high quality low compression images.

Context Model and Entropy Coder Provide Compression

A context model examines the coefficients within a codeblock and makes a prediction of each bit of each coefficient one bit at a time. These contexts are used with the Mqcoder to produce a segment of compressed data.[6] This is a key component in representing the wavelet coefficients in fewer bits, i.e. achieving compression.

JPEG 2000 Applications

Digital Cameras

Current digital cameras using JPEG compression typically compress the full resolution of the sensor in a JPEG file. They also typically store a reduced resolution "icon" within the JPEG file as required by the Digital Camera File System standard.[7] The low resolution icon is used for the LCD display in the camera (or sometimes yet another copy of the image is stored at a different resolution for the LCD display). When the digital camera images are uploaded to a catalogue/manipulation program like Apple's iPhoto™, an additional low resolution version of the image is created which iPhoto™ uses to enable browsing. Then if images are emailed often the full resolution images are too large, and iPhoto™ will reduce them and recompress them before sending. If the images are saved to a web page two resolutions are saved: one for initial viewing and one for high quality.

With JPEG 2000 there is little need to create and store multiple copies of the image at different resolutions. The camera LCD can use the lowest resolution version of the image (which can be stored first in the file). Web pages and emails can use higher resolutions (more of the same file). Photo quality printers can access the entire image. The ability to access the various resolutions reduces both storage requirements and computational requirements.

Printing

The process of previewing an image on a monitor and then printing is part of many workflows. The large difference between monitor and print resolutions makes the ability to access a low resolution image for the monitor important. Viewing a preview image of a page requires a low resolution image. Checking a small detail in a preview requires the ability to selectively access a higher resolution image. Printing a color document to a black and white printer is facilitated in JPEG 2000 by the ability to access just the luminance component. Because printers may print a document top to bottom or left to right, access to different regions of the image is important. Eventually the whole document is printed, but the printer needs only one row or column of tiles at a time, for example.

JPEG 2000 Extensions

Document Imaging

Part 6 of the JPEG 2000 standard is defining a file format for documents, called JPEG 2000 Multilayer[8] or JPM. This JPM file will typically contain many JPEG 2000 images (or layout objects). A page might be divided into a the text portion plus a couple of different images. Each of these objects could be accessed independently. These objects could be stored at different resolutions or quality. Further, this file format defines ways to store and access collections of pages. Thus an entire book, or even an entire set of encyclopedias could be stored as a single JPM file. The "page collection" boxes in the JPM file define a typical access (perhaps in "page number" order), but other boxes can define other access methods e.g. all pages with a figure. It is even possible to define a JPM file which is just an index to another file. Such a file might contain the result of a search on the set of encyclopedias.

Video

JPEG 2000 does not provide a method for utilizing correlation between multiple images. This means that a DCT based compression system like MPEG provides better overall compression for video sequences than JPEG 2000. However, compression of one frame using another frame does lead to some difficulties: memory is required to buffer multiple frames, latency is increased, and quality is not uniform among the frames (MPEG codes every *n*th frame independently and these are much higher quality than the intermediate frames).

In situations where access to individual frames is important, JPEG 2000 provides a useful solution. JPEG 2000 Part 3 defines a file format for storing video as a sequence of JPEG 2000 images.[9] This is valuable if access is desired to any frame without decoding other frames. It is especially useful in editing situations. If a frame must be changed in a frame to frame compression system several adjacent frames will need to be modified either because the changed frame depends on them, or they depend on the changed frame. Clearly JPEG 2000's access to individual frames is an advantage. (Of course, MPEG allows frames to be stored independently, but JPEG 2000 provides progression by quality or bitrate for the video frames).

JPEG 2000 Interactive Protocol (JPIP)

Access to various portions of the JPEG 2000 file is possible when the file is stored locally e.g. on a compact flash card or disk drive. However, access to images of different resolutions, regions, and quality is even more useful over a network.[10] The JPEG committee is currently working on a protocol to describe portions of a JPEG 2000 image and transmit them over the internet. This will allow development of the same features available with local images over the network.[11]

Conclusion

JPEG 2000 provides just "the bit" of the compressed file necessary for any application. This access to the data is more important than the exact PSNR achieved at a given bitrate. JPEG 2000 files can be compressed without knowledge of the many different ways the file will be used.

References

1. TU-T Rec. T.800 | IS 15444-1:2000, Information Technology—JPEG 2000 image coding system: Core coding system.
2. ISO/IEC 11544 and ITU-T Recommendation T.82. JBIG bilevel image compression standard, 1993.
3. W. B. Pennebaker and J. L. Mitchell, *JPEG still image datacompression standard*, Van Nostrand Reinhold, New York, NY, 1992.
4. Call for contributions for JPEG2000: Image Coding System, ISO/IEC JTC1/SC29/WG1 N505, March 1997.
5. M. Boliek, M. Gormish, E. L. Schwartz, and A. F. Keith, "Decoding compression with reversible embedded wavelets (CREW) codestreams," Journal of Electronic Imaging, Vol. 7. No. 3, pp. 402-209, July 1998.
6. D. S. Taubman and M.W. Marcellin, "JPEG2000: Image compression fundamentals standards, and practice," Kluwer Academic Publishers, September 1, 2001.
7. Japan Electronic Industry Development Association (JEIDA), *Design Rule for Camera File System (DCF)*, adopted December 1998, available (in English) at http://www.pima.net/standards/it10/PIMA15740/dcf.htm
8. JPEG2000 Part 6 Study Text A for FDIS, ISO/IEC JTC1/ SC29/WG1 N2588, April 2002.
9. ITU-T Rec. T.802 | IS 15444-3:2000, Information Technology—JPEG 2000 image coding system: Motion JPEG 2000.
10. Martin Boliek, Gene K. Wu, Michael J. Gormish, "JPEG 2000 for Efficient Imaging in a Client/Server Environment," Proc. SPIE Vol. 4472, p. 212-223.
11. Prandolini, Houchin, Clark, "Call for JPIP technology proposals," ISO/IEC JTC1/SC29/WG1 N2550, March 2002.

Biography

Michael Gormish earned bachelor's degrees in mathematics and electrical engineering from the University of Kansas, and a masters and Ph.D. in electrical engineering from Stanford University finishing in 1994. He currently works on image compression, image enhancement, and image quality at Ricoh Innovations, Inc. Michael has contributed to JPEG 2000 in a variety of ways, most recently he was the coeditor of JPEG 2000 Part 4, the conformance tests. He maintains a web site on JPEG 2000 at: www.crc.ricoh.com/ ~gormish/JPEG2000.html