Color Imaging with JPEG 2000

Robert Buckley
Xerox Architecture Center, Xerox Corp.
Webster, New York

Abstract

The JPEG 2000 image compression standard is a wavelet-based follow-on to JPEG. While it offers better compression performance, JPEG 2000 is not expected to replace JPEG. What it is expected to do is to offer new features and capabilities, so that an encoder/decoder (client/server) will have processing choices not possible with previous compression methods. This is expected to lead to unique color imaging applications for the Web, digital cameras and handheld devices. Part 1 of the JPEG 2000 standard defines the minimal decoder and was approved as an international standard in January 2001. Still in the standards pipeline are several more parts that will specify decoder extensions, file formats, conformance and reference software. This paper will give the background and an overview of JPEG 2000, demonstrate its performance and then describe its implications for color imaging.

Background

JPEG 2000 is an image-compression standard, developed by the same committee (ISO/IEC JTC1 SC29/WG1) that previously developed JPEG, although with a different group of participants and contributors. JPEG 2000 was conceived as a next generation image compression standard that would improve on the performance of JPEG while, more significantly, adding features and capabilities not available with Baseline JPEG compression. The formal call for contributions for JPEG 2000 was issued in March 1997. Among the criteria it gave for the next generation compression system were:

- Superior low bit-rate performance
- Continuous-tone and bi-level compression
- Lossless and lossy compression
- Progressive transmission by pixel accuracy and resolution
- Fixed-rate, fixed-size, limited workspace memory
- Random codestream access and processing
- Robustness to bit-errors

By December 1997, the committee had evaluated 24 proposals and selected a wavelet-based approach as the basis for the JPEG 2000 standard. Over the next 3 years, the original proposal was revised and developed so that by December 2000 a final proposal for the JPEG 2000 codestream syntax and decoder semantics had passed the committee review process and was sent out for balloting as an international standard. Part 1 of the JPEG 2000 standard (ISO/IEC 15444-1) defines the core JPEG 2000 decoder and was approved in January 2001. The official web site for the JPEG 2000 standards committee is http://www.jpeg.org.

JPEG 2000 Overview

Like most image compression methods, JPEG 2000 is an example of the general 3-step compression model: Transform–Quantize–Encode. The Encode step, using a lossless entropy coder, is what actually reduces the data volume and achieves the compression. The Transform and Quantize steps leading up to it condition the signal so that it can be efficiently coded (and of course the Encode step is designed to efficiently code the output of the Quantize step).

In JPEG 2000, the Transform step uses a wavelet transform, high-and low-pass (H and L) filtering the original image vertically and horizontally to decompose it into 4 quarter-size subband images (Figure 1). The Part 1 JPEG 2000 standard specifies 2 wavelet filters: a 9x7 floating point and a 5x3 integer. (Part 2 extensions will allow user-specified wavelet filters.) The LL subband image, generated by low-pass vertical and horizontal wavelet filters, can be wavelet transformed again to create 4 more subband images, and so on. Figure 1 only shows 2 decomposition levels (3 resolution levels), but the JPEG 2000 standard itself allows up to 15 decomposition levels. The subband images are then quantized using a scalar quantizer. Finally an arithmetic coder is applied to each bitplane of each subband to compress the quantized subband image data. For details on JPEG 2000 compression, see Refs. 1 and 2.

Figure 2 compares JPEG and JPEG 2000 for a color image compressed at 2 and 0.125 bits per pixel (bpp). For the same image, Figure 3 compares the PSNR (Peak Signal-to-Noise Ratio) obtained using JPEG and JPEG 2000 at a range of compressed bit rates. Although there is a significant PSNR difference between JPEG and JPEG 2000 at 2 bpp there is little visual difference between them for this image.

Where JPEG 2000 is superior is at low bit rates. With a luminance-chrominance image, chominance subsampling by 2 in each direction and default Huffman tables, the upper bound on Baseline JPEG is 0.125 bpp or 192:1 for a 24-bit color image. Lower bit rates are possible with JPEG 2000. Figure 4 shows a JPEG 2000-compressed color image at 0.0625 bpp, corresponding to a compression ratio of 384:1.

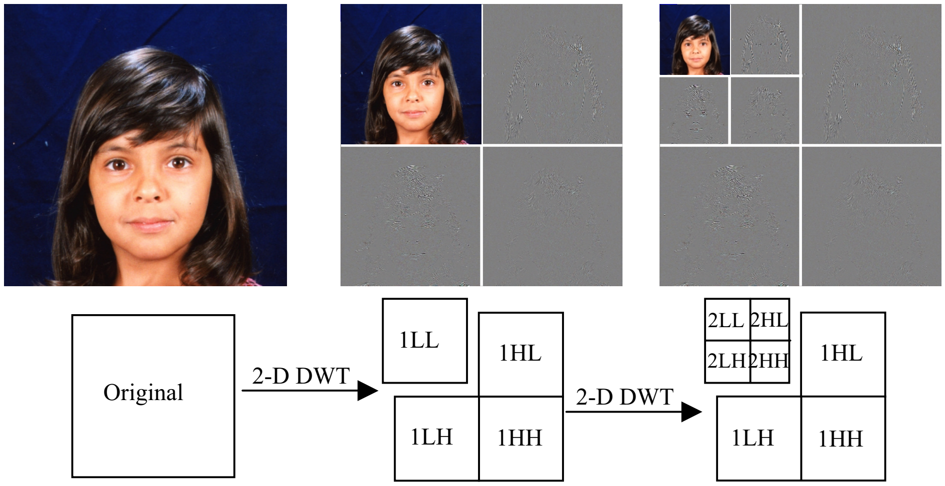

Figure 1. 2-level Wavelet Decomposition via 2-D Discrete Wavelet Transform (DWT)

Figure 2. Comparison of JPEG and JPEG 2000 image compression at 2.0 and 0.125 bpp. Top row: 2.0 bpp; Bottom row 0.125 bpp; Left column JPEG; Right column JPEG 2000

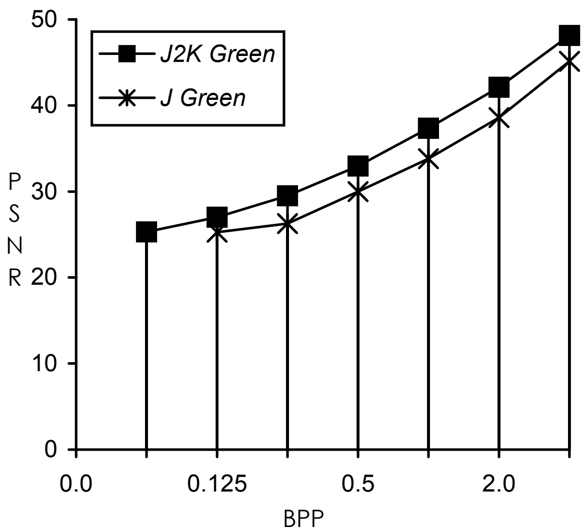

Figure 3. PSNR (dB) vs Bit Rate (bpp) for JPEG (J) and JPEG 2000 (J2K) compression of Green component of Figure 2.

Figure 4. JPEG 2000 image compression at 0.0625 bpp

The compression advantages of JPEG 2000 come at a cost of added complexity. Using the committee's unoptimized reference software, compression and decompression of the images in Figure 2 took about 5 times longer for JPEG 2000 than for JPEG, using the Independent JPEG Group software, on a Sun Ultra 5 with a 333 MHz processor. Software JPEG 2000 implementations are now available in C, C++ and Java.[3,4,5] were available when this was originally written, but this is an area of feverish activity and many more will be available by the time of the conference. Altogether, they significantly improve the software performance, although none are as fast as JPEG. The first commercial JPEG 2000 chip was announced in Spring 2001; more are expected and will start appearing in late 2001.

Binary and Lossless Image Compression

The JPEG 2000 architecture supports binary image compression. For binary compression, the Transform and Quantize steps are skipped and only the Encode step is used. As a binary lossless compressor, JPEG 2000 is better than Group 3 1-dimensional (MH) but not as good as Group 4 (MMR) compression. On the 8 CCITT Test Documents at 400 dpi, JPEG 2000 compression is about twice that obtained with MH but 25% less than Group 4 and 50% less than that obtained with JBIG.

The JPEG 2000 architecture also support lossless continuous-tone image compression. By using the reversible 5x3 integer wavelet filter with no quantization, the image in Figure 2 can be losslessly compressed with a bit rate as low as 11.5 bpp. However, JPEG-LS on average is superior to JPEG 2000 for lossless continuous-tone compression.[6]

Color Image Compression

A compressor typically operates on image data without regard to whether it represents a red, luminance or alpha channel component, although compression parameters, such as JPEG Q tables, can be selected differently for each component. The compressor treats the components simply as continuous-tone data with some pixel-to-pixel correlation that it can exploit to compress the image. When the components are red, green and blue, there is also component-to-component correlation, and decorrelating the components improves multi-component image compression.

With JPEG 2000, a user has the option to use one of two 3x3 component transforms in the encoder when compressing a multi-component image. One is reversible and the other irreversible. "Reversible" means the transform and its inverse together give an identity transform and the original pixel values are recoverable. While an "irreversible" transform may be reversible in practical terms with no noticeable differences, the original pixel values are not recoverable with limited precision arithmetic and so the transform is strictly speaking not reversible.

Reversible:
$$Y = (R + 2G + B) / 4$$
$$C1 = B - G$$
$$C2 = R - G$$

Irreversible:
$$Y = 0.299\,R + 0.587\,G + 0.114\,B$$
$$C1 = -0.16875\,R - 0.33126\,G + 0.5\,B$$
$$C2 = 0.5\,R - 0.41869\,G - 0.08131\,B$$

The irreversible transform is identical to the RGB-to-YCbCr transform in ITU-R 601. Although the effect of these transforms is to convert RGB-type components to luminance-chrominance-type components, they are not intended to be color transforms but rather decorrelating transforms. The JPEG 2000 encoder itself makes no provision for subsampling the "chrominance" components generated by these transforms.

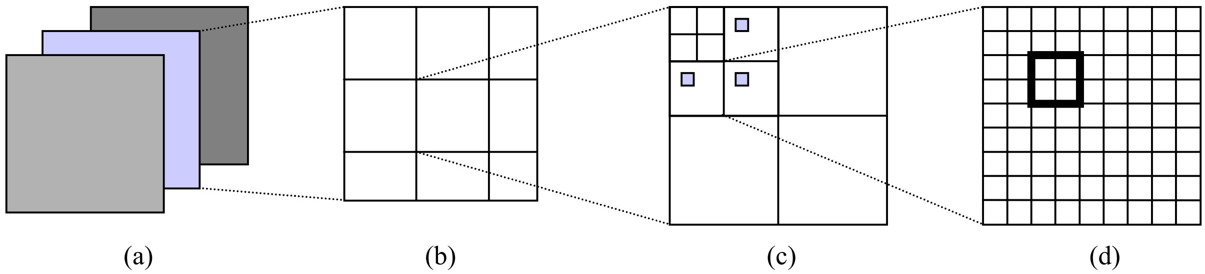

Figure 5. JPEG 2000 codestream structure: (a) component images, (b) tiles, (c) decomposition levels, (d) codeblocks with precinct

Without a component transform, the RGB image in Figure 2 can be losslessly compressed to 13.1 bpp. Both components transforms reduce the lossless bitrate for this image: to 11.5 bpp for the reversible component transform and 8.1 bpp for the irreversible. The difference is due to the extra bit depth carried by the reversible transform and to their relative ability to decorrelate the RGB components of the original image.

JPEG 2000 Codestream Structure

JPEG 2000 derives many of its features and capabilities from how it structures the compression and organizes the resulting codestream. An original image consists of components, such as luminance and chrominance, or red, green and blue. Each component can be optionally divided into the same size rectangular tiles, where each tile independently wavelet transformed; see Figure 5. Successive applications of the wavelet transform create multiple resolution or decomposition levels, where a decomposition level consists of 3 subbands. These 3 subbands when combined with the next lower resolution-level image produce the next higher resolution-level image; see Figure 1.

The subbands at a given resolution level are identically divided into code-blocks. A rectangular array of code-blocks in the same location in all 3 subbands of a decomposition level can be grouped together to form a precinct or packet partition location. Figure 5 shows one precinct in the second decomposition level.

Just as the wavelet transform operates independently on tiles, the entropy coder operates independently on code-blocks. The results of coding the bitplanes within a code-block are collected into one or more layers.[7] A chunk of the codestream can therefore represent the coded data from a layer within a precinct from a particular decomposition level of the tile for a component. These chunks occur in the JPEG 2000 codestream as contiguous segments known as packets. For a more comprehensive discussion of the codestream structure, see Refs. 1 and 2. Later on it will be observed that the JPEG 2000 file format can break up the codestream and its packets into fragments.

The notion of packets is significant to a JPEG 2000 user for two reasons. The first is that the order of packets in the codestream determines the behavior of the decoder and the way in which an image is decoded and displayed. If all the packets for a given decomposition level and component occur together, followed by the packets for other decomposition levels and components, then the image would be decoded and displayed progressively by resolution. If the packets are ordered by layer across all decomposition levels and components, then the image would be decoded and displayed progressively by quality or precision. Altogether, the JPEG 2000 standard specifies 5 progression orders, according to the ordering of packets.

The second reason is that more or less packets can be used to represent or subsequently decompress an image. For example, JPEG 2000 supports lossy-to-lossless decompression. If the decoder decompresses all packets in a losslessly compressed image, it will obtain an exact reconstruction of the original image. However, if the decoder decides to decompress only a subset of the packets and truncates the codestream, it will instead obtain a lossy image, with the loss determined by the number of packets it uses and the way in which they were ordered.

In other words, a decoder doesn't have to decompress the entire codestream to obtain a useable image. More than that, a decoder can stop decoding packets when it has obtained an image whose quality is adequate for the application at hand. This is in contrast to JPEG, where partial decompression is not an option. For example, the two JPEG 2000 compressed images in Figure 2 need not be two independently compressed images. They can be two intermediate steps in the decompression of a single JPEG 2000-compressed image, one decompressing more packets than the other to get a higher quality image. Unlike a JPEG codestream, a JPEG 2000 codestream is scalable.

This changes the compression paradigm. In the past with JPEG, all the choices were made at the encoder. With JPEG 2000, the decoder now has choices when it comes to decompressing a codestream and rendering an image.

JPEG 2000 Applications

The JPEG 2000 codestream structure allows a single codestream to support multiple applications. For example, depending on the layering and number of decomposition levels, a single, losslessly-compressed codestream can serve the needs of hand-held devices such as cell phones or

PDAs, low-resolution displays, mid-resolution printers, and high-resolution imaging devices. It is all a question of how many packets are used and how they are delivered to the application.

In FlashPix-type applications, a single tiled and multi-decomposition level JPEG 2000 codestream can support zooming and panning in client-server applications. FlashPix would need multiple multi-resolution and tiled JPEG codestreams to do the same thing, with an overhead of 33% to accommodate multiple resolutions.

In constrained-memory devices, such as digital cameras, JPEG 2000 can be used to specify in advance, and even after the memory is partly or completely full, the storage capacity in numbers of pictures. This is possible for two reasons. First, a user can directly specify the size of a JPEG 2000 codestream (or compression ratio), unlike JPEG, where a user directly specifies the Quality or Q factor, which is related to the codestream size in an image-dependent way. (The JPEG bit rates in Figure 2 are as close to 2.0 and 0.125 bpp as it was possible to get by specifying the Q factor.) Second, a user can go back and delete packets from already-captured images to make room for more images, thus trading off quality for capacity after the images have been captured.

Another significant feature of JPEG 2000 is its support for Region-of-Interest or ROI coding. For example, a user may wish to encode some parts of an image with more fidelity than other parts. In Figure 6, 0.5 bpp were allocated to the entire image, but a "boost" was applied to the region around the face, so that it uses proportionally more of the allocated bits and has less loss compared to the surrounding background region.

Figure 6. JPEG 2000 image compression with ROI coding applied around the face region

It is possible to arrange for the ROI to be encoded in different layers of the codestream compared to the rest of the image. In this way and with progressive decompression, the ROI can be decompressed earlier and with more fidelity than the background of the image. As more packets are received, then the fidelity of the background region improves until, when all packets have been received and decompressed, the quality of the decompressed image is the same in both ROI and background. In effect, the ROI has higher priority in the codestream. The ROI can be explicitly defined in the codestream by means of a binary mask, or implicitly by arranging for the ROI wavelet coefficients to be scaled so that they are in an entirely separate set of layers than the background wavelet coefficients.

JPEG 2000 File Formats

In general, the file format that encapsulates the codestream is significant since most applications deal directly with the file format rather than with the codestream. One of the perceived shortcomings of the original JPEG standard was that it did not define a file format that could be used to interpret the codestream, e.g. identify the color space, and associate metadata with it. A subsequent part of the JPEG standard did define the SPIFF file format, but by then several formats were already in use for JPEG, some more successful than others, and SPIFF was ignored.

To help avoid the same problem with JPEG 2000, Part 1 of the standard defines the optional JP2 file format for JPEG 2000 compressed data.[8] JP2 is a simple file format. However, it establishes the architecture for a family of JPEG 2000 file formats, members of which are being specified in subsequent parts of the JPEG 2000 standard to meet the needs of more advanced and specific applications.

Architecturally, the JPEG 2000 family file formats are sequences of "boxes," much like the atoms used in the Apple QuickTime format. A box has three fields: a 32-bit Length, a 32-bit Type, and a variable-size Data field containing data interpreted according to the Type value. Boxes can be nested so that the Data field of a box can be another box.

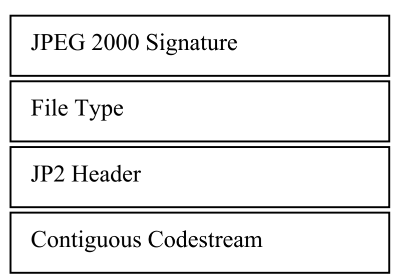

Figure 7. JP2 File Organization

The simplest JP2 file would contain the 4 boxes in the order shown in Figure 7. The JPEG2000 Signature Box identifies the file as a member of the JPEG 2000 file format family and contains some characters that an application can use to detect some common file transmission errors. The File Type Box identifies the file as a JP2 file and contains version number and conformance information. The JP2 Header Box Date field contains information about the image, such as its color space, resolution, number of components and so on. Finally the Contiguous Codestream Box contains the JPEG 2000 codestream. Other optional boxes allow a file writer to add intellectual property, metadata and other vendor-specific information.

Color Spaces and Color Management

A conforming JP2 file supports a limited number of color spaces—sRGB and an sRGB-based grayscale space—and a limited set of ICC profiles from the Monochome Input and Three-Component Matrix-Based Input profile classes.[9] A JP2 file can also support opacity components and color palettes. While these choices meet the needs of some applications, this approach has created two issues.

Most compressed-image file formats, image communication systems and now devices such as digital still cameras support some sort of luminance-chrominance color space. Such spaces improve compression by decorrelating RGB components and allowing chrominance subsampling. Some of the benefits of luminance-chrominance compression can be obtained with the component transform described earlier in the JPEG 2000 encoder, but this assumes RGB and not luminance-chrominance components. An amendment to the JPEG 2000 standard has been proposed that would add support for an sRGB-based YCC color space color space to the JP2 file format, and is working its way through the ISO standards process. This color space would apply the YCbCr coefficients to sRGB color data.

The second issue had to do with ICC support. Tying the support for ICC profiles to the current version of the ICC Profile Format specification[10] can create problems for future compatibility, especially since the ICC is currently updating its spec. Writers can easily write a fixed subset of the current version of the spec, but readers may have to support that version of the spec as well as future versions in a more general Color Management System. An amendment was proposed that would update the ICC profile support in JP2; in particular, it would add support for a new ICC profile introduced in the latest version of the ICC profile specification. This amendment was considered but dropped, partly because it would have required major revisions to several portions of the JPEG 2000 specification.

JPEG 2000 File Format Family

Parts of the JPEG 2000 standard still under development will specify 3 other members of the JPEG 2000 file format family: JPX, MJ2 and JPM. JPX is an extension of JP2 that allows broader color space support, including any ICC profile; fragmented, non-contiguous codestreams; references to external data; and multiple codestreams for alpha blending. Fragments are the results of splitting up the codestream into collections of packets, which the file format exposes at the application level. Though targeted at JPEG 2000 data, JPX also supports other compression methods using the JPEG 2000 file format architecture.

The Motion JPEG 2000 or MJ2 file format supports multiple images and timing information.[8] It is intended for use in digital still cameras, where the user can shoot frames either one at a time or in sequence.

The JPEG 2000 Multi-Layer or JPM file format is intended for the representation of compound images using the Mixed Raster Content imaging model with JPEG 2000, and possibly other compression methods.[11] JPM also uses the notion of multiple frames or pages and is intended to take advantage of the features of JPEG 2000 for document imaging on the web.

Schedule for the JPEG 2000 Standard

Part 1 of the JPEG 2000 standard, which specifies the core decoder and the JP2 file format, was approved in January 2001. Still to come at the time this was written are Parts 2-6. Their status is summarized in the following table.

Table 1. Schedule for JPEG 2000 Standard (ISO/IEC 15444)

Part	Content	Due
1	Core decoder and file format (JP2)	Done
2	Extensions to Part 1 decoder & file format (JPX)	11/01
3	Motion JPEG 2000 File Format (MJ2)	11/01
4	Conformance and profiles	02/02
5	Reference Software	11/01
6	JPEG 2000 Compound Image File Format (JPM)	02/02

A significant topic of discussion in the JPEG 2000 committee (at the time this was written) is "profiling" or defining subsets of the Part 1 JPEG 2000 decoder. At issue is defining a subset of the capabilities of the codestream and decoder, especially with regard to tile and code-block size and filter use, so as to define specified subsets of the JPEG 2000 decoder that would enable efficient, cost-effective and compliant implementations.

Summary

JPEG 2000 is a wavelet-based image compression standard, designed as the next generation follow-on to the DCT-based JPEG compression standard. The JPEG 2000 core decoder standard was approved as an international standard in January 2001. It offers better low-bit rate and rate-distortion performance than existing compression methods. While it is not the best binary or lossless image compression method, JPEG 2000 does have the advantage of a single decoder architecture that supports them both besides lossy continuous-tone image compression.

The structure and organization of a JPEG 2000 codestream enables decoder options and multiple modes for progressively decompressing and rendering a compressed image, including lossy-to-lossless and Region-of-Interest decompression. The JPEG 2000 encoder includes an optional decorrelating component transform to improve the compression of multi-component images, such as RGB. The JPEG 2000 standard also defines a file format architecture with color space and color management support.

The JPEG 2000 file format family is intended to enable a range of potential applications, including digital photography, web imaging, and compound and document imaging.

Acknowledgments

The author would like to thank Jim Reid for his help in generating the images and results reported here and Alexander Reid for his help in generating the plots.

References

1. M.W. Marcellin, M.J. Gormish, A. Bilgin and M.P. Boliek, An Overview of JPEG-2000, pp. 523-541, Proc. IEEE Data Compression Conf., March 2000, Snowbird, UT
2. C. Christopoulos, A. Skodras and T. Ebrahimi, The JPEG2000 Still Image Coding System: An Overview, IEEE Trans. Consumer Electronics, vol. 46, no. 4, pp. 1103-1127, Nov. 2000
3. Jasper, http://www.ece.ubc.ca/~mdadams/jasper/
4. Kakadu, http://maestro.ee.unsw.edu.au/~taubman/kakadu/
5. JJ2000, http://jj2000.epfl.ch/
6. D. Santa-Cruz, T. Ebrahimi, J. Askelöf, M. Larsson and C. A. Christopoulos, JPEG 2000 still image coding versus other standards, SPIE Applications of Digital Image Processing XXIII, Vol. 4115, SPIE Annual Meeting, San Diego, August 2000
7. D. Taubman, High Performance Scalable Image Compression with EBCOT, IEEE Trans. Image Processing, vol. 9, no. 7, pp. 1158-1170, July 2000
8. J.S. Houchin and D.W. Singer, JPEG 2000 file format: an imaging architecture for today and tomorrow, SPIE Applications of Digital Image Processing XXIII, Vol. 4115, p. 455, SPIE Annual Meeting, San Diego, August 2000
9. J.S. Houchin and C.R Hauf, JPEG 2000 file format provides flexible architecture for color encoding, SPIE Applications of Digital Image Processing XXIII, Vol. 4115, p. 476, SPIE Annual Meeting, San Diego, August 2000
10. International Color Consortium, ICC Profile Format Specification ICC.1:1998-09, www.color.org
11. L.H. Sharpe II and R.R. Buckley, JPEG 2000 .jpm File Format: a layered imaging architecture for document imaging on the web, SPIE Applications of Digital Image Processing XXIII, Vol. 4115, SPIE Annual Meeting, San Diego, August 2000