Technology Watch Report

JPEG 2000 - a Practical Digital Preservation Standard?

Robert Buckley, Ph.D.
Xerox Research Center Webster
Webster, New York

DPC Technology Watch Series Report 08-01
February 2008

Executive Summary

JPEG 2000 is a wavelet-based standard for the compression of still digital images. It was developed by the ISO JPEG committee to improve on the performance of JPEG while adding significant new features and capabilities to enable new imaging applications. The JPEG 2000 compression method is part of a multi-part standard that defines a compression architecture, file format family, client-server protocol and other components for advanced applications.

Instead of replacing JPEG, JPEG 2000 has created new opportunities in geospatial and medical imaging, digital cinema, image repositories and networked image access. These opportunities are enabled by the JPEG 2000 feature set:

- A single architecture for lossless and visually lossless image compression
- A single JPEG 2000 master image can supply multiple derivative images
- Progressive display, multi-resolution imaging and scalable image quality
- The ability to handle large and high-dynamic range images
- Generous metadata support

With JPEG 2000, an application can access and decode only as much of the compressed image as needed to perform the task at hand. This means a viewer, for example, can open a gigapixel image almost instantly by retrieving and decompressing a low resolution, display-sized image from the JPEG 2000 codestream.

JPEG 2000 also improves a user's ability to interact with an image. The zoom, pan, and rotate operations that users increasingly expect in networked image systems are performed dynamically by accessing and decompressing just those parts of the JPEG 2000 codestream containing the compressed image data for the region of interest. The JPEG 2000 data can be either converted to JPEG and delivered for viewing with a standard image browser or delivered to a native JPEG 2000 viewer using the JPIP client-server protocol, developed to support the JPEG 2000 feature set.

Using a single JPEG 2000 master to satisfy user requests for dynamic viewing reduces storage costs and management overhead by eliminating the need to maintain multiple derivatives in a repository.

Beyond image access and distribution, JPEG 2000 is being used increasingly as a repository and archival image format. What is remarkable is that many repositories are storing "visually lossless" JPEG 2000 files: the compression is lossy and irreversible but the artefacts are not noticeable and do not interfere with the performance of applications. Compared to uncompressed TIFF, visually lossless JPEG 2000 compression can reduce the amount of storage by an order of magnitude or more.

Keywords: JPEG 2000, image compression, access format, archival format, distribution format, scalability, visually lossless, smart decoding, digital preservation.

What is JPEG 2000?

At its core, JPEG 2000 is an international standard for the compression of still digital images. But this is only the beginning of the story. Modern signal-based approaches to image compression have reached the point where it is more and more difficult to obtain dramatic improvements in compression performance and efficiency. As a result, other factors often enter into the discussion of the pros and cons of a compression method. This is particularly true for JPEG 2000. So while it improves on the compression performance of previous methods, JPEG 2000 offers significant new features and capabilities. It is these new features and capabilities that are largely responsible for the growing use of JPEG 2000 and the interest it has generated in several application areas, including digital preservation.

While JPEG 2000 began as a standard for image compression, it has evolved into a suite of standards that goes beyond the core method and includes the components needed for a JPEG 2000 ecosystem. This suite includes file format standards for JPEG 2000-compressed image data, network protocols for JPEG 2000 data, conformance specifications, reference implementations and advanced applications.

JPEG 2000 is the most recent standard from the JPEG committee. JPEG, which stands for Joint Photographic Experts Group, is a joint committee which brings together interests and participants from multiple standards developing organizations. The standards it creates are in most cases published jointly by ISO/IEC and ITU. So the official designation for the JPEG 2000 standard, and more specifically for Part 1 of the standard, is ISO/IEC 15444-1 and ITU-T T.800[1]. Part 1 describes the core decoder.

An Introduction to the JPEG 2000 Standard

The first standard created by the JPEG committee was the eponymous JPEG standard, published in 1992 and widely used on the web and in digital cameras. The Baseline JPEG standard is a lossy compression method. This means that the original, uncompressed image that existed before JPEG compression was applied cannot be exactly recovered from the JPEG-compressed data or codestream. The differences in most cases are adjusted so that they are small and either not noticeable or unobtrusive.

As a result, this type of compression is called visually lossless: the before and after images are mathematically but not visually distinguishable. However since the effects of compression, however slight, are not reversible and the original image is not recoverable, this is irreversible compression.

There are cases where the requirements are stricter and mathematically lossless behaviour is desired. In these cases, the image recovered from the compressed data needs to be identical, bit for bit, with the original, uncompressed image before compression. This type of compression is lossless or reversible compression.

As something of an afterthought, a simple lossless compression method was added to the JPEG suite of standards. Although it was called Lossless JPEG, it was based on a

[1] The current version is ISO/IEC 15444-1:2004 Information technology – JPEG 2000 image coding system – Part 1: Core coding system

completely different algorithm than Baseline JPEG: all the two had in common was their name[2]. As a result, lossless and lossy JPEG compressions were incompatible choices and a user had to select between them. By comparison, JPEG 2000 offers a unified approach to lossless and lossy compression, and the difference between them is one of degree, not of method.

Believing they could do better, the JPEG committee initiated the development of a more advanced method for lossless image compression. The standard that emerged from this activity was JPEG-LS. In connection with this work Ricoh proposed a wavelet-based method that offered both lossless and lossy compression, the ability to extract images with different resolutions and quality, and the ability to extract regions of interest[3]. While this proposal was not accepted for JPEG-LS, it did generate enough interest to initiate a new standards activity, aimed at producing a more general and widely applicable image coding system. It was this activity that produced the JPEG 2000 standard.

This new standard had several goals: improve on existing standards; serve the needs of high-end and emerging applications; and open up new markets and opportunities for image compression. Seven areas were identified that would be addressed by the new standard[4]:

- Low bit-rate compression performance
- Lossless and lossy compression in a single codestream.
- Large images, noting that JPEG was limited to 64K by 64K images without tiling.
- Single decompression architecture to integrate features and applications in a unified framework
- Transmission in noisy environments
- Computer-generated imagery, which is smoother and less noisy compared to natural imagery
- Compound documents, which combine images and text.

JPEG 2000 has successfully addressed all the areas on the list. It is interesting to note that in the beginning low bit-rate performance and doing a better job than JPEG at medium to high compression was cited as a primary driver for the activity. However, in the end it is the performance of JPEG 2000 in other areas on this list that have attracted applications to it.

With this list in mind, the call for contributions for this new standard was issued in March 1997. About two dozen proposals were received and evaluated, and in December 1997 the JPEG committee decided on a wavelet-based approach as the basis for the new standard. Between then and December 2000, the coding method was developed and refined, with reference to a verification model that was used to test and assess various technology proposals.

---

[2] While its use has been reported in medical imaging, Lossless JPEG was not widely adopted.

[3] M.J. Gormish, JPEG 2000: Worth the Wait?, 42nd Midwest Symposium on Circuits and Systems, Volume 2, pp. 766-769 (1999)

[4] From original Call for Contributions, 21 March 1997

The result was Part 1 of the JPEG 2000 standard, which defines the core decoder. This part of the standard focuses on the codestream, which is the collection of bits that contain the compressed image data and the parameters needed for interpreting it. In particular, Part 1 defines the codestream syntax, the methods for interpreting the codestream and for decoding the compressed image data it contains, and an optional file format for encapsulating the codestream with associated metadata including colour encoding.

Since Part 1 was completed, several other parts have emerged—13 in all[5]. Table 1 lists the parts and a brief description of each.

Part	Description
1	Core Coding System, including optional file format (JP2)
2	Extensions to Part 1 algorithm, including extended file format (JPX)
3	Motion JPEG 2000 file format (MJ2) for timed sequence of images, such as video
4	Conformance Testing
5	Reference Software for implementing JPEG 2000
6	Compound Image File Format (JPM) for paged sequence of compound images
8	Secure JPEG 2000 (JPSEC) for the encryption, authentication and conditional access of JPEG 2000 image data and metadata; an amendment due in 2008 will extend the notion from the codestream to the file formats.
9	Interactivity tools, APIs and protocols (JPIP) - a client-server protocol for JPEG 2000 files and data
10	JPEG 2000 3D (JP3D) - extensions to volumetric data
11	Wireless (JPWL) - extensions for error protection and correction in wireless and error-prone networks.
12	ISO Base Media File Format - Joint JPEG-MPEG[6] activity that defined a base media format, MJ2 in Part 3 is an application of it
13	An entry level JPEG 2000 encoder – an encoder equivalent of Part 1
14	XML Structure Representation of JPEG 2000 file formats and codestreams

It is evident from this table that the JPEG 2000 standard covers not only the compression algorithm, but also what’s needed to support the use of JPEG 2000 in a variety of applications.

[5] All parts are complete and published, except 10, which is due soon and 14. Part 7 was started but abandoned and the work subsumed by another part.
[6] MPEG is the Motion Picture Experts Group, a sister group of JPEG that is focused on digital audio and video

Most applications, including those in the digital preservation domain, can meet their needs with JPEG 2000 codestreams created with reference to Part 1. Part 1 was designed so that implementations of it would not require licensing fees or royalties. Part 2 was developed for more advanced applications and defines specialized extensions to the Part 1 algorithm, not all of which are free of fees for their use. It also defines an extension to the Part 1 file format.

JPEG 2000 is an open standard. This means that it was developed within the documented JPEG process, which was open to a wide variety of participants who, besides their time and energy, contributed technology, software, intellectual property and documentation.

One consequence of JPEG 2000 being an open standard is that implementations are available from several independent sources. Reference software is defined in Part 5 of the standard and available via the web: one is written in C and is open source; the other is written in Java. Also, another open source codec (coder-decoder) written in C is available. Finally, both software and hardware JPEG 2000 implementations are available from several vendors.

JPEG 2000 is also supported by several popular image processing toolkits, editing programs and stand-alone viewers[7]. It has been pointed out that JPEG 2000 is not natively supported by web browsers. In this respect it is like TIFF and PDF, neither of which is natively supported by browsers. And like them, plug-ins are available.

Features

While JPEG 2000 did improve on the performance of JPEG, it did not replace JPEG in most of the applications where JPEG was being used. What happened instead is that the JPEG 2000 feature set opened up new opportunities and enabled new applications. It is its feature set that is responsible for the growing use of JPEG 2000.

First of all, JPEG 2000 defines a single compression algorithm for both lossless (reversible) *and* lossy (irreversible) compression. Whether the compression is lossy or lossless depends on how the compression parameters are set[8]. The core decoder specified in Part 1 of the JPEG 2000 standard, and all implementations that comply with it, can decompress both lossless and lossy codestreams.

So how well does JPEG 2000 compression perform? Lossless compression tests have shown that, for most images, JPEG 2000 gives compression ratios a few percentage points less than JPEG-LS, but a few points better than PNG. While the results vary from image to image and depend on image content, lossless compression ratios are usually around 2:1. For the full size colour version of the image in Figure 1, which was one of the JPEG 2000 test images, the JPEG 2000 lossless compression ratio is 2.08:1[9]. Appendix A compares several lossless compression methods, including JPEG-LS, JPEG 2000 and PNG, applied to four test images.

---

[7] Among them are Adobe® Photoshop® CS3, Corel® Paint Shop Pro®, Irfanview, ER Viewer and Apple® Preview and QuickTime® 7.
[8] It is even possible to obtain a lossy codestream from a lossless one.
[9] The coding parameters include reversible component and wavelet transforms, five resolution levels, one quality layer and 64x64 code-blocks.

In lossy compression tests performed by the JPEG committee, print versions of JPEG 2000-compressed images were judged visually lossless at a bit rate of around 1 bit per pixel, which corresponds to a compression ratio of 24:1[10]. For this particular image set and test conditions, this was the boundary between visually lossless and visually lossy compression, past which compression artefacts began to become visible in the decompressed image. At this and higher compression ratios (or lower bit rates), JPEG 2000 consistently outperformed JPEG: JPEG 2000 images were judged to have the same quality as JPEG images that had higher bit rates or lower compression ratios.

At lower, more conservative compression ratios, JPEG 2000 and JPEG both are judged visually lossless. In this range, the performance advantage of JPEG 2000 persists. While the differences may not be visible, JPEG 2000 images have lower errors than the JPEG images with the same compression ratio or bit rate. In this case, the error is the difference between the original uncompressed and decompressed images[11].

From the standpoint of compression alone, JPEG 2000 is a single algorithm that is almost as good as JPEG-LS for lossless compression and better than JPEG for visually lossless and lossy, but more complex than both. What the complexity buys you is a set of features than neither offer.

If one had to name the single most important feature of JPEG 2000, it would probably be "smart decoding." With smart decoding, a JPEG 2000 codestream can support multiple decompression options and enable an application to access and decode only as much of the codestream as is needed for the task at hand. This makes interacting with the image easier and makes it possible for a single JPEG 2000 compressed image to supply multiple derivatives of the original image. For example, it can supply multiple reduced resolution versions of the original, a compressed image with a specified file size or error rate, or a high resolution, high quality view of a portion of the image with a region of interest.

The ability of a single JPEG 2000 image to support multiple derivatives has significant consequences for online image collections. Applications that offer access to online collections typically provide multiple, different-sized image views, each supported by a separate, pre-computed image file. This leads to increased costs, as there are now multiple image files to store, maintain and synchronize. The wavelet transform in JPEG 2000 as a matter of course provides multiple resolution images within a JPEG 2000 codestream.

For example, progressive decompression of a JPEG 2000 codestream can provide the sequence of images shown in Figure 1, starting with the lowest resolution, thumbnail image on the left, with each subsequent image derived from the one on its left. Because of the nature of the wavelet transform in JPEG 2000, each image in the

---

[10] Bit rate is the number of bits in the compressed image divided by the number of pixels in the original image; compression ratio is the number of bits in the original image divided by the number of bits in the compressed image.

[11] In the technical literature, the error is typically reported as PSNR or Peak Signal-to-Noise Ratio, which is 20 times the logarithm of the ratio of the maximum pixel value to the root mean square error. PSNR is given in dB.

sequence has twice the height and width, and twice the resolution of the image to its left. The final image in the sequence would have the same size and resolution as the original.

Figure 1: Multiple Resolution Levels; from left – 1/8, 1/4, 1/2 and full resolution

One of the consequences of this is that a gigapixel image compressed with JPEG 2000 can be opened and viewed with little if any delay since all that is accessed from the codestream is the lowest resolution image needed for creating a display-sized image. So instead of gigabytes, the application deals with only as many megabytes or even kilobytes as are needed for a display view. This means that the time it takes to display an image compressed using JPEG 2000 is roughly independent of the size of the original. Successive views would access approximately the same number of bytes, trading off the resolution with the amount of the image shown, from a low resolution view of the entire image to a high resolution view of a portion of it.

The example in Figure 1 showed a progression in resolution; Figure 2 shows a progression in quality, enabled by partitioning a JPEG 2000 compressed image into multiple quality levels or layers. The partitions can be set so that the intermediate images have specified compression ratios (or bit rates) or error rates. In the example of Figure 2, the quality levels of the intermediate images were selected to correspond to compression ratios of 400:1, 100:1 and 24:1; the final image in the sequence is losslessly compressed. Analogous to the resolution progression in Figure 1, the image at a particular quality level is derived from the images at lower quality levels preceding it in the sequence.

Figure 2: Multiple Quality Levels; clockwise from upper left – lossless, 24:1, 100:1 and 400:1 compression

Scaling the images in Figure 2 to make them fit on the page diminished somewhat the visibility of the differences between them. To better compare the lossless or original image with the visually lossless image compressed at 24:1, the same region has been cropped from both images and reproduced in Figure 3, at 200 pixels per inch. The ability to access just these regions of the image in the codestream is one of the features of JPEG 2000.

Figure 3: Cropped regions of lossless (left) and 24:1 compressed images (right)

The ability to create JPEG 2000 compressed images that contain different quality levels has a consequence worth noting here. It makes it possible to obtain lower quality image versions from a higher quality compressed image without having to recompress or even decompress the image. For example, if a losslessly compressed master image is stored in an archive, service requests can be met with visually lossless derivatives of it, which can save time and bandwidth. It makes it possible to distinguish the needs of image archive from image delivery but have them both satisfied by a single file.

In some applications, the master can be created in the first place to meet a specified compressed file size (or bit rate) or error rate. Specifying the file size or bit rate is useful in bandwidth constrained applications. The result would be compressed images with the same size, independent of image content. Specifying the error rate would create compressed images whose size varies with image content in a way that provides approximately constant image quality[12], since the error rate correlates, if imperfectly, with image quality.

Besides a wide array of output options, JPEG 2000 also supports a variety of input images. It can handle large images: images with hundreds of megabytes are common and at least one vendor claims support for terabyte images. JPEG 2000 can also handle high-bit depth and high-dynamic range images, unlike JPEG which is practically limited to 8 bits per component. This is useful as colour encodings with 12 and 16 bits per component become more prevalent. There is no requirement that all components have the same bit depth. Finally, JPEG 2000 can easily handle multi-component images, going beyond greyscale and colour images to include multi-spectral image data. If there are limitations in JPEG 2000 applications, they are more likely to be due to the source, the format or the implementation than to the JPEG 2000 algorithm itself.

[12] Bill Comstock, Using PSNR thresholds to modulate the degree of lossy compression in JPEG2000 files, DPC-British Library Joint JPEG 2000 Workshop, 25 June 2007, http://www.dpconline.org/graphics/events/0706jpeg2000wkshop.html

It is not possible to discuss digital preservation without some discussion of metadata[13]. The JPEG 2000 standard has generous support for metadata. However this support is at the level of the file format rather than the codestream, so this is a good point to introduce the JPEG 2000 file formats.

File Formats

When it was published, the original JPEG standard defined only the compression and decompression algorithms and the representation of the compressed data. Before the JPEG committee defined a file format a few years later in Part 3, several appeared in the market in the meantime, some creating interoperability problems. Taking this as a lesson learned, the JPEG committee this time around defined an optional file format for JPEG 2000 compressed data and included it in Part 1 of the JPEG 2000 standard. The file format associates metadata with the codestream and serves as an interface for accessing and rendering the compressed image data.

The JPEG 2000 standard defines a family of four binary file formats:

- JP2 is a basic format for single images and a limited set of colour encodings; defined in Part 1, file extension .jp2
- JPX is an extended version of JP2, with support for one or more images, a larger set of colour encodings and other advanced features; defined in Part 2, file extension .jpf
- MJ2 is a format for motion images, where each frame is a JPEG 2000-compressed image; defined in Part 3, file extension .mj2
- JPM is a format for compound images, which can consist of images, graphics and text; defined in Part 6, file extension .jpm

The file format defined in Part 1 is JP2. It is a basic file format for JPEG 2000 data; it consists of a sequence of "boxes," shown in the schematic in Figure 4. The boxes in JPEG 2000 files have the same structure as atoms in QuickTime and MPEG-4 files. Each box has length, type and data fields: the length field can be used to skip boxes and navigate a file; the type field determines how the data field is interpreted.

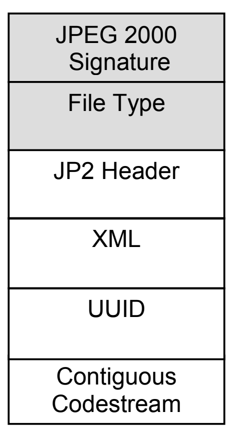

Figure 4: Schematic of a JP2 file.

[13] Preservation Metadata, Brian Lavoie and Richard Gartner, DPC Technology Watch Series Report 05-01, September 2005

The JPEG 2000 Signature box identifies the file as a member of the JPEG 2000 file format family and is the first box in all format family files. The File Type box identifies which member of the family the file is, and gives version, compatibility and profile information. In this case, the File Type is JP2. The JP2 Header box contains the image and colour parameters for rendering the image. It may also specify the capture or display resolution. The JP2 Header box is a superbox, which means its data field contains other boxes.

The JP2 file format was designed with digital photography in mind. This is reflected in the colour encodings defined in the standard. JP2 supports two colour encoding methods: the sRGB family of colour spaces and restricted ICC profiles. The sRGB encodings are sRGB, sYCC and a greyscale encoding defined in Part 1 that applies the sRGB non-linearity to luminance values. The use of ICC profiles is restricted to the Monochrome Input or Three-Component Matrix-Based Input profile[14]. These profiles are adequate to represent density values and RGB encodings such as ProPhotoRGB and Adobe RGB, which have a wider gamut than sRGB and are more suited to digital preservation. However, JP2 supports these encodings only indirectly via ICC profiles rather than directly as named colour encodings. Named encodings for these and other colour spaces require extensions, some of which are found in the extended JPX format defined in Part 2 of the JPEG 2000 standard.

The XML and UUID (Universally Unique Identifier) boxes are boxes for optionally including metadata in the file. The XML box contains XML-formatted metadata; the UUID box contains non-XML formatted or vendor specific metadata. Part 1 provides the framework for including metadata in the file; it is organizations external to the JPEG committee that have defined the metadata elements and encodings, and applications that determine what metadata is used.

For example, a UUID box was defined for adding GeoTIFF-based geospatial metadata to a JP2 file. A more recent method for adding GML-based[15] geographic information to a JPEG 2000 file uses XML boxes plus features found in JPX, the extended version of JP2. A UUID box is also used to incorporate IPTC tags directly in a JP2 file. These metadata definitions are not part of the JPEG 2000 standard. There is the possibility that a reader may not recognize the identifier in a UUID box so that it contents are unavailable to the application. To help avoid this, a JP2 file may include a UUID Info Box, which contains a URL where an application can obtain information about the UUID boxes in the file.

Interpreting XML-encoded data in an XML box is usually not an issue, as the XML data is defined with reference to a published schema and namespace. As noted earlier, metadata elements are defined external to the JPEG 2000 standard. An exception to this is the extended metadata definitions in Part 2 of the JPEG 2000 standard. These definitions are based on the DIG 35 specification[16] and specify XML-encoded metadata for:

---

[14] The ICC profiles in Part 1 of the JPEG 2000 standard are defined with respect to the ICC.1:1998-09 specification. The tag values and computational models for these profiles are supported in the current ICC.1:2004-10 specification, which also adds new tags and tag types so that backward compatibility is not guaranteed.

[15] GML is the Geographic Markup Language, http://www.opengeospatial.org/standards/gml

[16] DIG35 Specification - Metadata for Digital Images. Version 1.0. August 2000

- Image Creation
- Content Description
- History
- Intellectual Property Rights
- Image Identifier

The image creation metadata elements apply mainly to images from digital cameras. Applications of course are free to choose whatever metadata schema suit their needs. For example, *Chronicling America* uses XML boxes that contain Dublin Core elements in the JP2 files that contain images of historic newspapers.

An amendment to Part 1 of the JPEG 2000 standard would enable integrating EXIF metadata in a JPEG 2000 file using a UUID box. The amendment is due in late 2008.

While the standard specifies the order of the first three boxes in a JP2 file, it is good practice to have the metadata boxes precede the Contiguous Codestream box to make access to them faster. A JP2 file has a single Contiguous Codestream box, which contains the JPEG 2000 codestream and is normally the last box in the file. The codestream may contain multiple components, including an alpha channel. The JP2 header box contains information that links the components and colour channels, so for example there is no requirement that red, green and blue channels occur in that order in the codestream. The codestream components, when decompressed are interpreted with reference to the colour specification contained in the JP2 Header Box.

Part 2[17] defines the JPX file format, which is the extended version of the JP2 file format defined in Part 1. A JPX file supports any component set for which there is a corresponding colour specification. It also allows for other compression methods, including uncompressed image data, which raises the interesting possibility of a JPX file being an alternative container for uncompressed image data.

JPX adds a large numbers of features and a great deal more capability to what's available with JP2. For example, it is possible to take the codestream in a JPX file and fragment it, with the fragments distributed around the file, across several disks or over the Internet. Also, a JPX file can contain multiple images, along with instructions on how they are to be combined to create either a still image or an animated sequence.

In addition, two other JPX features have application to digital preservation: expanded colour support and the ability to link metadata and labels to regions-of-interest in an image. The expanded colour support includes full ICC profiles, vendor-specified colour spaces and a larger number of named colour spaces, including ProPhoto RGB, CIELAB, extended sRGB spaces and PhotoYCC.

JPX introduced a label box, which simply contains a text string, and an ROI Description box, which gives the location and size of a rectangular or elliptical-shaped region of interest (ROI) in the image. JPX also introduced an Association box,

---

http://www.i3a.org/i_dig35.html

[17] ISO/IEC 15444-2:2004, Information technology – JPEG 2000 image coding system: Extensions

a superbox that links two or more boxes in its data field. This means an Association box can make semantic associations between Label, ROI Description and XML boxes and their contents. This makes it possible to attach a label to a region of interest, label a metadata field, or link metadata with a specific region in the image. These associations are part of the file and travel with it, so they are available to viewers and applications.

In geospatial imaging for example, GMLJP2 files[18] use the Association box at the top level in the file to package GML-encoded geographic metadata in a JPX file. The Association box links the label "gml.data" to one or more Association boxes, each of which in turn associates a label or name with GML-encoded metadata.

A JPX reader is not expected to support all the features that a JPX file can have. A file contains a Reader Requirements box that signals which features it uses. However to promote interoperability, the standard defines a baseline JPX file, a subset of the full JPX specification which limits the scope of some features and eliminates others to arrive at a set that a baseline JPX reader is required to support. The File Type box identifies a file as a Baseline JPX file. As it is currently defined, Baseline JPX does not include the Association, Label or ROI Description boxes. It may be that a significant set of digital preservation applications could be served by a modest extension of JP2, adding a few, select features from JPX.

Part 3 defines the Motion JPEG 2000 or MJ2 format for motion sequences of JPEG 2000 images[19]. While MJ2 files can contain audio Part 3 itself does not define an audio coding method. What is significant about the approach MJ2 takes to video coding, compared to MPEG, is that it uses intraframe compression only, where each frame is compressed independently and without reference to other frames. While this means that MJ2 does not take advantage of the correlation between frames to improve compression, it also means reduced latency in the video stream and that "smart decoding" techniques apply as well to individual frames in a video sequence as they do to still images.

Part 6 of the JPEG 2000 standard defines JPM, the JPEG 2000 standard file format for compound images and image documents[20]. What's significant about a compound image is that it combines images and text, which have different characteristics and quality requirements. For example, text can be binary with high spatial resolution to preserve the edge information that is important for text quality and legibility. Text preferably would use lossless or near-lossless compression. By comparison, images can have lower spatial resolution than text, but need high tone scale resolution for colour fidelity; they can use visually lossless compression.

In a JPM file, each type of content type is represented and compressed in a way matched to its characteristics so as to give the best results for overall quality and compressed file size. To meet these differing requirements, JPM, which stands for

---

[18] GML in JPEG 2000 for Georgraphic Imagery Encoding Specification, http://www.opengeospatial.org/standards/gmljp2
[19] ISO/IEC 15444-3:2007, Information technology – JPEG 2000 image coding system: Motion JPEG 2000
[20] ISO/IEC 15444-6:2003, Information technology – JPEG 2000 image coding system: Compound image file format

Multi-Layer JPEG 2000, uses the MRC or Mixed Raster Content approach. MRC defines a multi-layer imaging model in which the text and image content of the original compound image are decomposed into separate layers[21]. Each layer can then be encoded independently of the other layers, with a spatial resolution, tone resolution and compression method matched to its content. A JPM file organizes these layers in pairs called layout objects.

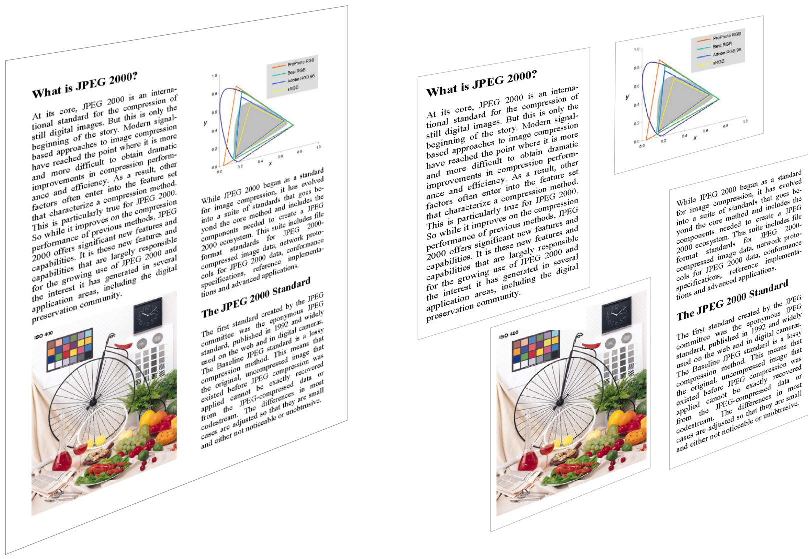

Figure 5: Compound Image and JPM Layout Objects

Figure 5 illustrates the decomposition of a sample compound image into four layout objects. A layout object consists of a binary mask layer and a colour layer, and is rendered by applying the colour through the mask layer. Mask layers would preferably use a binary compression method, such as Group 4, JBIG or JBIG2, which were designed for text[22]. The colour layers would use JPEG 2000 compression, which was designed for continuous-tone images. The decomposition of a page into layout objects is not unique and depends on the application and the implementation of the JPM encoder.

A JPM file organizes layout objects into pages and pages into page collections, leading up to the main page collection, which is the starting point for navigating the contents of a JPM file. JPM also supports auxiliary page collections, which provide alternative views of the document or a subset of it. For example, an auxiliary page collection could reference all the pages that contained a specified search term. An amendment to the JPM standard published this year defines an XML-based format for storing OCR results in a JPM file.

While other formats also support a multi-layer format for representing images[23], what is significant about JPM is that it extends the paradigm of smart decoding to the level

[21] Not to be confused with the quality layers that occur in the JPEG 2000 codestream.
[22] Although JPEG 2000 can also compress binary images, JBIG2, JBIG and Group 4 are superior methods for the lossless compression of binary images.
[23] PDF, TIFF-FX and DjVu®

of the document, allowing access to individual pages within a document and to individual layout objects within a page. JPM also makes it possible to interactively access pages and layout objects and progressively render them so that only those pages and parts of a page needed for the task at hand are accessed[24,25].

Performance

A short tour of a JPEG 2000 coder will give an idea of where the features of JPEG 2000 come from and what to watch for and be aware of when using JPEG 2000[26]. Figure 6 shows the steps in the JPEG 2000 compression of an image.

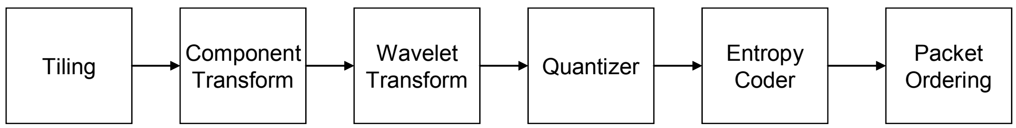

Figure 6: From left to right, the steps in the JPEG 2000 compression of an image.

Conceptually the first thing that happens when an image is JPEG 2000 compressed is that it is split into rectangular tiles. Since each tile is compressed independently of every other tile, the usual rationale for tiling is to limit the amount of memory needed to implement JPEG 2000 and to provide independent access to regions in an image. Some implementations are designed for tiling and perform best with tiled images; other implementations can compress megabyte and gigabyte images without tiling. And there are other mechanisms for independent access to image regions.

If the components in a multi-component image are red, green and blue, then an optional component transform is available to convert them to luminance and chrominance[27]. Two transforms are supported by Part 1 decoders. One is the same as the conversion that JPEG uses. Because its transform coefficients are floating point, there is some small round-off error in a round trip conversion to and from luminance-chrominance values. As a result, the original red, green and blue component values are not recoverable and this is an irreversible component transform (ICT). The second transform was designed with coefficients so that it is a reversible component transform (RCT) and the original red, green and blue values are recoverable from the luminance-chrominance values.

The purpose of these transforms is to decorrelate the red, green and blue image components, which improves compression performance by redistributing the energy across the image components. In this respect, the ICT does a better job at decorrelating the red, green and blue values than the RCT, which leads to better

---

[24] R. Buckley and J. Reid, A JPEG 2000 compound image file reader/writer and interactive viewer, Applications of Digital Image Processing XXVI. Edited by A. Tescher, Proceedings of the SPIE, Volume 5203, pp. 236-247 (2003)

[25] M. Boliek and M. Gormish, Network Access to Parts of Archived Document Image Files, Proceeding of IS&T 2006 Archiving Conference, pp. 147-150 (2006)

[26] For a more detailed technical description of JPEG 2000 compression, the reader is referred to M. Rabbani and R. Joshi, An overview of the JPEG2000 still image compression standard, Signal Processing: Image Communication 17, pp. 3-48 (2002), and D. Taubman and M. Marcellin, JPEG2000: Standard for interactive imaging, Proceedings of the IEEE, volume 90 (number 8), pages 1336-1357 (2002).

[27] Alternatively, luminance and chrominance components can be precomputed, subsampled and tiled before being wavelet transformed.

compression. Whichever transform is used before compression; the inverse transform is applied after decompression to restore the red, green and blue values.

Unlike JPEG, a JPEG 2000 coder does not subsample the chrominance components, but subsampled chrominance components can be prepared in advance and input to a JPEG 2000 coder along with a full resolution luminance component.

With these preliminaries out of the way, next come the core operations of the JPEG 2000 coder, starting with the wavelet transform. The wavelet transform is applied to each component independently after the optional component transform.

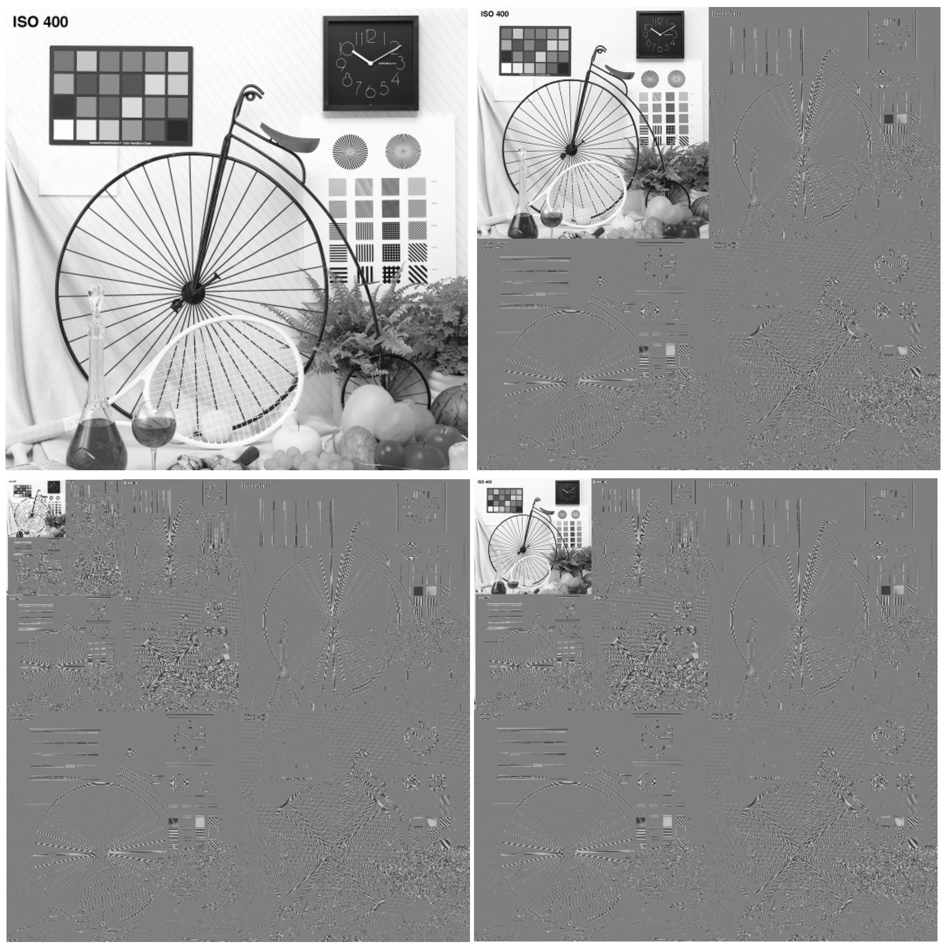

Figure 7: Three-level wavelet transform of luminance component.

Figure 7 shows the result of applying the wavelet transform to the luminance component (a) of the original image. One level of waveform transform produces the four sub-images or subbands, shown in (b), each half the width and height of the original. The upper left subband is a lower-resolution version of the original; the other three represent various combinations of edge data in the original image. The four subband image can be recombined to reconstruct the original image.

The wavelet transform can be applied again to the upper left subband image, to create four more subband images, each now one quarter the width and height of the original. The result after two levels of wavelet transform is the seven subband images shown in (c). The result after three levels is the ten subband images shown in (d). Running this operation in reverse will generate a progression of resolution levels, analogous to those shown in Figure 1.

The standard allows for as many as 32 levels of wavelet transform, although in practice only as many levels are applied as are needed to create a conveniently-sized thumbnail from the subband image at the lowest resolution. For example, five levels of wavelet transform applied to a 10000-by-8000 original will create at the lowest resolution a 312-by-250 or approximately QVGA-size subband image. Furthermore, little if any improvement in compression is obtained with more than four or five levels of wavelet transform.

As with the component transform, the JPEG 2000 standard defines two wavelet transforms. One is the 9-7 transform, which uses floating point operations and, like the ICT, is irreversible. The other is the 5-3 transform, which, like the RCT, is reversible[28]. And as with the component transform, compression performance is better with the irreversible wavelet transform than with the reversible one.

For example, when the image in Figure 7(a) is losslessly compressed with the reversible component and wavelet transforms, the resulting compression ratio is 2.10. When the same operation is performed with the irreversible component and wavelet transforms, the compression ratio is 2.97. Because of round-off errors in the floating point transforms, the result is lossy. However the error, which is the difference between the decompressed and original images, is small: 68% of the luminance values are the same between the two images and there is no visible difference between the images. In many applications, the cost of the difference would be outweighed by the gain in compression.

The next step after the wavelet transform is the quantization of the subband images, which is a lossy step that reduces their precision in order to improve their compressibility in the following step, which is the arithmetic coder. In lossless compression, the subband images are passed unchanged to the arithmetic coder.

After quantization comes the entropy coder, which takes advantage of the statistical properties of the quantized subband images to reduce the number of bits used to represent them. This is the stage where the actual compression occurs. While baseline JPEG use Huffman coding, JPEG 2000 uses a more sophisticated and computationally expensive method known as adaptive arithmetic coding.

The subband images are partitioned into fixed-size codeblocks and the arithmetic coder applied independently to each bitplane[29] of each subband image within a codeblock. This has two important consequences for JPEG 2000 performance. First the codeblock and collections of them called precincts allow for regions of the compressed image to be accessed and decoded independently of other regions. Second, the coded bitplanes from different subbands and components can be partitioned into quality layers.

---

[28] 9-7 and 5-3 refer to the sizes of the filters used to implement the wavelet transform.
[29] Because arithmetic coding can become less effective for lower bitplanes, JPEG 2000 has an optional Bypass mode that skips the coding of the lower bitplanes, which saves time with little reduction in compression efficiency.

A quality layer is one increment in quality for the entire image. For example, the image in Figure 2 has four quality layers. The first gives the coarsest rendition, equivalent to 400:1 compression. Each quality layer adds more coded data and improves the quality of the decompressed image. In the example of Figure 2, the intermediate quality images have specified compression ratios or equivalently bit rates. It is also possible for the intermediate images to have specified error rates.

While adding resolution levels improves compression up to a point, adding layers takes away from it, if only slightly because of the small overhead their use requires in the codestream. Many applications specify the use of only one layer. However, some applications use both multiple resolutions and multiple layers, since reduced resolution images can be rendered more quickly with fewer layers and without noticeable loss in quality[30].

Another way to think of a layer is as a collection of packets. Packets are the fundamental building blocks of a JPEG 2000 codestream. While a layer is an increment in quality for the entire image, a packet is an increment in quality for a specific position of a given resolution of a component of a tile. The interleaving of packets in a codestream determines the progression order in which compressed data is received and decompressed.

JPEG 2000 defines five progression orders or packet orderings. In resolution-major progression orders, the packets for all layers, positions and components of the lowest resolution come in the codestream before all those for the next higher resolution level. Figure 1 shows the sequence of images in a codestream with a resolution-major progression order. Figure 2 is an example of the sequence from the layer-progression order, where all the packets for one quality layer occur in the codestream before all the packets for the next quality layer. Progression order is another design choice in the use of JPEG 2000.

Figure 3 is an example of the image that is obtained when the packets retrieved are the ones for all quality levels, resolution levels and components, but just for this position in the image.

Variable length coders, such as the arithmetic coder in JPEG 2000 and the Huffman coder in JPEG, can be vulnerable to errors. For example, errors that appear in the middle of the codestream can cause the coder to lose synchronization and make all the data from there on in is not decodable. JPEG 2000 has several mechanisms to limit the effect of errors and their propagation in the codestream. One is the Start-of-Packet markers defined in Part 1 that provide reference points for periodic resynchronization of the decoder. Part 11 defines more sophisticated methods, such as unequal error protection, for protecting the data from transmission errors in noisy environments.

Multiple compression-decompression cycles can cause generation loss when lossy compression is used. For example, image quality can suffer when an image is

[30] R. Buckley, W. Stumbo and J. Reid, The Use of JPEG2000 in the Information Packages of the OAIS Reference Model, Proceeding of IS&T 2007 Archiving Conference, pp. 24-28 (2007)

decompressed, cropped and then recompressed. In this situation it has been observed that the generation loss with JPEG 2000 is slight and less than it is with JPEG[31]. This is so if the codeblock boundaries are unchanged, which is possible with JPEG 2000, except of course around where the image was cropped.

Altogether JPEG 2000 offers significantly more features and flexibility than JPEG and other compression methods at the cost of more complexity in its implementation and use. While shielded from the complexity of the implementation, the user may notice that it takes longer to compress the same size image with JPEG 2000 than with JPEG, as JPEG 2000 is computationally more intensive than JPEG. The same thing may be true with decompression as well, except that with JPEG 2000 the user only needs to touch and decompress only as much of the compressed image as needed. As a result, decompressing and viewing a JPEG 2000 image can be faster than decompressing and viewing the same image compressed with JPEG[32].

With all the choices offered by JPEG 2000, the JPEG 2000 standard did two things to help promote interoperability. The first was to define profiles in Part 1 that restrict the size and use of codestream parameters, such as image size, tile size and codeblock size. Five profiles have been defined to date, including two for Digital Cinema. The second thing was to define compliance classes in Part 4 that limit parameters, such as image size, number of components and number of bits per components, which a decoder has to support.

Within these limits, there is still scope for an application to choose coding parameters such as tile size, wavelet filter, number of resolution levels, number of quality layers, progression order, codeblock size, and whether or not to use lossless or visually lossless compression. The next section describes the experience of some JPEG 2000 practitioners in making some of these choices.

JPEG 2000 Practice in Access and Preservation

JPEG 2000 is being used for geospatial imaging, medical imaging and by the cultural heritage and digital preservation communities. Many digital collection and library systems support JPEG 2000, and several institutions use it in their collections. This section will discuss the experiences of a few of those institutions, chosen as they highlight the issues and questions many have about JPEG 2000.

An institution that has done much work in the use of JPEG 2000 and is now one of the leaders in its adoption is the Harvard University Library (HUL). The move to JPEG 2000 was driven in part by institutional clients who wanted features such as interactive zoom, pan, and rotate. These requirements are not easily implemented with TIFF, GIF, or JPEG, but are easily enabled by JPEG 2000. In 2006, Harvard reported

[31] R.L. Joshi, M. Rabbani and M.A. Lepley, Comparison of multiple compression cycle performance for JPEG and JPEG 2000, Applications of Digital Image Processing XXIII. Edited by A. Tescher, Proceedings of the SPIE, Volume 4115, pp. 492-501 (2000)
[32] The next standard in the JPEG committee pipeline is JPEG-XR, based on the Microsoft HD Photo compressed file format and aimed at the next generation of digital cameras. JPEG-XR has computational complexity close to JPEG with some of the features of JPEG 2000.

the successful test migration of more than 10,000 TIFF, GIF, and JPEG images to equivalent lossless and lossy JPEG 2000 form[33].

Over the past several years, the rate of acquisition of new JPEG 2000 images into the HUL Digital Repository Service (DRS) has steadily increased, while that for JPEG and TIFF has decreased. The DRS now manages about two million JPEG 2000 images, and JPEG 2000 is becoming the default format for image conversion and acquisition. A single JPEG 2000 master image in the repository enables the dynamic delivery of arbitrarily-sized use images (transcoded to JPEG for rendering by the client browser), all computed on demand from the master, thereby eliminating the need to maintain multiple variants in anticipation of client requests. In addition, JPEG 2000 enables an interactive interface that lets users perform the zoom, pan, and rotation operations that now form the common user expectation for web-based image delivery[34].

Library and Archives Canada ran a year-long JPEG 2000 pilot project over 2006 and 2007, the results of which were described at the Museums and Web 2007 conference[35]. This pilot was undertaken to address many of the questions that cultural institutions have regarding JPEG 2000. One of their main results was to show that the use of JPEG 2000 could reduce file sizes significantly without loss to image quality. In the case of lossless archival masters, the compression ratio was typically around 2:1. For production or access masters, they specified a recommended compression ratio of 24:1 for colour images, which included photographs, prints, drawings and maps, and 8:1 for greyscale images, which included newspapers, microfilm and textual materials. They found that the JPEG 2000 codec they used performed best when images were tiled, and they recommended tiles sizes of 512 by 512 and 1024 by 1024. They also observed that the use of JPEG 2000 meant that derivative files were no longer required. The JP2 files they created in this pilot contained XML boxes with MODS-based metadata records.

The Library of Congress already makes use of JPEG 2000. For example, Civil War maps in the American Memory collection are compressed using JPEG 2000[36]. A client's pan and zoom requests are served with reference to a JPEG 2000 image; the resulting views are transcoded to JPEG for delivery to a standard web browser. The site also offers the option of downloading the JPEG 2000 image of the map. The Library's collection still has some maps compressed using MrSID®, a proprietary wavelet-based compression method that predates JPEG 2000.

JPEG 2000 is also used by the Chronicling America website, which was officially announced in March 2007 and currently provides access to US newspaper pages from 1900 to 1910[37]. Chronicling America is sponsored jointly by the National Endowment for the Humanities and the Library of Congress as part of the National Digital Newspaper Program (NDNP). It is worth noting the project objectives and constraints, since they are ones with broad applicability:

---

[33] S. Abrams, Automated Migration for Image Preservation, Proceedings of IS&T 2006 Archiving Conference, pp. 113-116 (2006)
[34] http://hul.harvard.edu/ldi/html/delivery.html
[35] http://www.archimuse.com/mw2007/papers/desrochers/desrochers.html
[36] http://rs6.loc.gov/ammem/collections/civil_war_maps/
[37] http://www.loc.gov/chroniclingamerica/home.html

- Convenient accessibility over the World Wide Web for the general public
- Page images of sufficient spatial and tonal resolution to support effective OCR performance
- Use of digital formats with a high probability of sustainability
- Attention to the cost of digital conversion and maintenance

With these objectives in mind, the program adopted a solution based on JPEG 2000 production masters. These masters are processed to meet user requests and then converted to JPEG files. Pages are delivered and viewed as JPEG images in a Flash-based application that lets the user zoom in to and out of the image and pan by click and drag. The JPEG files are dynamically created on demand from JPEG 2000 production masters. These masters contain a lossy but visually lossless, tiled JPEG2000-compressed greyscale image, with multiple resolution levels and multiple quality layers, encapsulated in a JP2 file with Dublin Core-compliant metadata[38]. While the production masters are visually lossless JPEG 2000 files, the archival masters are still uncompressed greyscale TIFF images from microfilm scans. Some microfilm scanners now have the ability to export JPEG 2000 files.

To look at the applicability of JPEG 2000 to a broader collection of content, the Library of Congress in October 2007 announced collaboration with Xerox Corporation to study the use of the JPEG 2000 format in large repositories of digital cultural heritage materials[39]. This study will build on the work that led previously to the JPEG 2000 profiles for newspapers, and extend it to cover prints, photographs and maps. It will pay attention to the preservation, access and performance issues associated with large image repositories, and how the JPEG 2000 standard can address those issues. The work, due to be completed in 2008, is expected to lead to specifications and best practices for the use of JPEG 2000. This is all part of the general ferment around JPEG 2000 in the cultural heritage and digital preservation communities.

The Harvard Image Delivery Service and Library of Congress Chronicling America site both satisfy user requests by retrieving JPEG 2000-compressed data from a master file on a server and converting it to visually lossless JPEG for delivery to the user's browser. By comparison, the National Archives of Japan Digital Gallery[40] offers both JPEG and JPEG 2000 delivery options. The latter requires motivated users to download and install a JPEG 2000 plug-in on Internet Explorer. Offering both lets users compare and, within the context of these implementations, get an appreciation for the performance of a native JPEG 2000 viewer.

The plug-in uses JPIP, the client-server protocol defined in Part 9 of the JPEG 2000 standard. JPIP was designed to take advantage of the features of JPEG 2000 in a networked environment. Besides supporting the exchange and negotiation of capabilities, it allows selective access to JPEG 2000 codestreams and files for incremental updates, metadata downloads, progressive rendering and interactive viewing. For example, JPIP is used to view visually lossless JP2 images from the

---

[38] The NDNP JPEG profile and the rationale for the choices made are described in http://www.loc.gov/ndnp/pdf/NDNP_JP2HistNewsProfile.pdf
[39] http://www.loc.gov/today/pr/2007/07-213.html
[40] http://jpimg.digital.archives.go.jp/kouseisai/index_e.html

Mars Reconnaissance Orbiter on the HiRISE website[41]. At the cost of downloading a plug-in, JPIP makes it possible to pan and zoom megapixel and gigapixel images in real time over the network.

Currently underway are several large-scale projects to digitize books. Participants in three of them—the Open Content Alliance, Google Library Project and Gallica at the Bibliothèque Nationale de France—co-authored a paper[42] at the IS&T 2007 Archiving Conference that described their coordinated study of the format for page image masters. What was most significant about their results is that they all agreed on the use of lossy JPEG 2000 for archival masters. First of all, the industrial scale of these projects could not support the storage costs of lossless compression. More than that, lossless compression was found to be unnecessary and their performance sub-optimal in this application: it was reported that lossy files were readable and gave better OCR performance than lossless files. This last observation was attributed to the reduction in image noise by the JPEG 2000 wavelet transform. In short, the reflex to use losslessly compressed files as archival masters was not consistent with the use scenarios for the files. As it is, saving scanned book pages as lossy compressed JPEG 2000 files is current practice at the Internet Archive.

As JPEG 2000 formats become more prevalent in digital repositories, it would be valuable to have tools that can validate JPEG 2000 files and metadata standards that are JPEG 2000-aware. Fortunately, both have appeared. JHOVE, the extensible framework for format validation developed jointly by JSTOR and Harvard University Library, has a module that recognizes and tests the well-formedness of JP2 and JPX files[43].

As for metadata standards, the latest version of the ANSI/NISO Z39.87 Data Dictionary for Technical Metadata for Digital Still Images[44], released in 2006, includes JPEG 2000. It recognizes JP2 as a format name and JPEG 2000 Lossless and Lossy as compression schemes. It also defines a JPEG 2000 data container for information about the JPEG codec and encoding. The three encoding options currently recorded are tile size, resolution levels and quality layers.

JPEG 2000 compression was added to PDF in Version 1.5, published in August 2003[45]. (The current version of PDF is 1.7, which will be the basis of the ISO standard for PDF.) What a PDF file does is encapsulate an entire JP2 or JPX file in a PDF data stream. The PDF Reference recommends using Baseline JPX, defined in Part 2 of the standard, with some restrictions that mostly have to do with colour space usage. It drops support for one colour space (CIEJab) and adds it for another (CMYK)[46].

---

[41] http://hirise.lpl.arizona.edu/

[42] S. Chapman et al., Page Image Compression for Mass Digitization, Proceedings of IS&T 2007 Archiving Conference, pp. 37-42 (2007)

[43] http://hul.harvard.edu/jhove/jpeg2000-hul.html

[44] http://www.niso.org/standards/index.html

[45] PDF Reference, Fourth Edition, Adobe® Portable Document Format Version 1.5, 2003 http://www.adobe.com/devnet/pdf/pdf_reference.html. .

[46] Part 7 of ISO 15930, the Graphic Technology Standard for Prepress Digital Data Exchange, includes the use of JPEG 2000 compression.

Before JPEG 2000 was added to PDF, the Association for Information and Image Management (AIIM) and the Association for Suppliers of Printing, Publishing and Converting Technologies (NPES) had begun the development of a file format for long-term document preservation based on existing PDF practices at the time. This activity led to PDF/A-1, which was published in 2005[47]. The primary purpose of PDF/A-1 is to define a format that preserves the visual appearance of a document over time. Based on a subset of PDF Version 1.4, PDF/A-1 allows, restricts or prohibits the use of PDF 1.4 elements according to how they are perceived to affect long term preservation.

Because it is based on PDF Version 1.4, PDF/A-1 supports only JPEG for continuous-tone image compression, although RunLength, Flate and LZW could also be used to losslessly compress continuous-tone images. Currently in working draft form is PDF/A-2, which would be based on PDF Version 1.7 and which would add, among other things and with some restrictions, JPEG 2000 as an image compression method. While the PDF/A standard recommends not using lossy compression, if it alters the content or degrades the quality of original data, in the end the standard leaves it up to the applications that generate PDF/A file to ensure that the contents are an accurate representation of the original.

The JPEG 2000 application that has probably been in the news the most lately has been digital cinema. At the most recent meeting of the JPEG committee[48], it was reported that over 5000 theatres now support digital cinema. In 2005, the Digital Cinema Initiative, a joint venture of six Hollywood movie studios, adopted JPEG 2000 image compression in the package used to distribute digital movies to theatres. The Digital Cinema Package uses an MXF[49] wrapper for the JPEG 2000-compressed images, audio and other movie data. The images use a 12-bit non-linear encoding of CIE XYZ values.

To support this application, the Digital Cinema Ad Hoc Group within the JPEG Committee developed two codestream profiles, one for 2K and the other for 4K images[50]. Visually lossless intraframe compression is used. Although each frame is compressed independently, the overall compressed data rate has to be within the 250 Megabits per second bandwidth limit of digital cinema. While this limit could be met with constant bit rate coding, several Digital Cinema coders use variable bit rate coding, allocating more compressed data to complex frames in a way that keeps the quality roughly constant over frames.

The growing use and adoption of JPEG 2000 in digital preservation and other domains enhances its standing as a sustainable format. The increased application means increased commitment on the part of both users and suppliers to the persistence of the format, the image data committed to it and the resources and tools for supporting it. As an open format, with a publicly accessible specification developed in an open forum using a documented process, JPEG 2000 has attracted multiple

---

[47] ISO 19005-1: 2005, Document management – Electronic document file format for long-term preservation - Part 1: Use of PDF 1.4 (PDF/A-1)
[48] 43rd Meeting, Kobe, Japan, 12-16 Nov 2007
[49] MXF or Material Exchange Format is a container format defined by SMPTE, the Society of Motion Picture and Television Engineers
[50] 2K images are up to 2048 by 1080; 4K images are up to 4096 by 2160

vendors and open source implementations. As a standard, JPEG 2000 is platform independent, and therefore not tied to specific hardware, software or operating system.

Adoption and openness are two key properties of a sustainable format. Other important properties have to do with the specification of the format itself and the ability to maintain and preserve files conforming to the specification. This makes the metadata support offered by the JPEG 2000 suite of file format standards significant; it enables a JPEG 2000 file to be self documenting with descriptive and preservation metadata. Various error protection mechanisms defined in the JPEG 2000 standard improve the robustness of the files that use them in potentially noisy environments.

The ability of JPEG 2000 to encode a wide range of images, including the number and type of components and bits per component, suggests that JPEG 2000 can meet longer term needs for image representation, beyond those of the image types that influenced its design and are typical of its current application. This may help manage diversity and promote uniform practices in an image collection, contributing to the overall sustainability of the collection. Finally, there are the questions about the advisability of compression, even lossless compression, in digital preservation. This question has come to the fore recently[51], driven by the economics of digital preservation and the needs of applications such as the ones that have been described here.

[51] For example, see J. Rog, Compression and Digital Preservation: do they go together?, Proceedings of IS&T 2007 Archiving Conference, pp. 80-83 (2007)

Conclusions

JPEG 2000 is a multi-vendor, open standard for the compression of still digital images. Applications and file formats have extended its use to motion and document imaging as well. Just as JPEG came along at about the same time as the web and digital still cameras and its use grew with them, JPEG 2000 has come at a time when there is the need for collecting and distributing an ever-growing volume of digital images, coupled with demands for interactive viewing and online access.

While initial interest in JPEG 2000 focused on access and distribution, more and more it is appearing as an archival format in large repositories. What is especially remarkable is that many of these applications use visually lossless JPEG 2000, which can reduce storage requirements by an order of magnitude compared to uncompressed TIFF. While visually lossless JPEG 2000 is lossy and irreversible, the losses are not visually noticeable and so far have not interfered with applications. So the overall result has been a cost savings with no adverse effect on performance.

JPEG 2000 offers significantly more features and flexibility than JPEG and other compression methods at the cost of more complexity in its use. What things then should a user be aware of? The codestream parameters with the most obvious implications for performance are the number of resolution levels, the number of quality layers, the size of the tiles, progression order and of course how much compression to use. Other codestream parameters have more subtle effects on performance. The settings of all these parameters would be collected together into one or more profiles. So far, there has not been a systematic approach to developing JPEG 2000 profiles for digital preservation and archival applications.

There profiles should include the file format as well. The main format elements to be determined by the profile are the colour encoding and the metadata content. It may be that a modest extension of JP2, with a few select features from JPX, would address a significant set of applications in digital preservation.

In the end, a clear understanding of the use cases and scenarios and the resulting requirements for image use will be key in developing the profiles and achieving the desired performance in a JPEG 2000-based applications.

Appendix A – Comparison of Lossless Compression Methods

	Compression Ratio			
Method	Image 1	Image 2	Page 1	Text
TIFF – Packbits	1.003	1.02	1.68	3.00
TIFF – LZW	1.32	1.38	3.10	3.56
ZIP	1.37	1.49	2.91	4.51
Lossless JPEG	1.45	1.52	2.77	3.07
PNG	1.57	1.59	3.88	4.35
Lossless JPEG 2000	2.08	2.00	4.25	4.36
JPEG-LS	2.11	2.05	4.86	5.20

The performance of any one lossless compression method and the comparative performance of several depend on the image content. Image 1 is the original of the image used in Figures 1-3; Image 2 is the original of the image used in Figure 7. Both originals were test images used by the JPEG committee.

Page 1 is a 300 dpi colour scan of Page 1 of the DPC 2006-2007 Annual Report; Text is the greyscale text cropped from Page 1. The Page 1 image is shown below on the left; the Text image on the right. About 77% of the pixels in the Text image are white.


Chairman's Introduction

Ronald Milne (Chairman) Director, Scholarship & Collections, British Library.

Membership Subscriptions from 2001 to 2007


Chairman's Introduction

Ronald Milne (Chairman) Director, Scholarship & Collections, British Library.

The images were available initially as uncompressed 24-bit RGB TIFF files, except in the case of the Text image, which started as an 8-bit Greyscale TIFF file. All TIFF files had 8 rows per strip. The compression ratios are computed with respect to the uncompressed TIFF file.

Lossless JPEG 2000 compression used reversible component and wavelet transforms, five resolution levels, one quality layer and 64x64 code-blocks, except in the case of the Text image, where no resolution levels were used.

When applied to the colour images, JPEG-LS used the same reversible component transform as Lossless JPEG 2000. Each component was then compressed independently.

About the Author

Robert Buckley is a Research Fellow with the Xerox Research Center Webster in Webster, NY. He has been with Xerox since 1981, when he joined the Xerox Palo Alto Research Center after receiving a PhD in Electrical Engineering from MIT. He also has an MA in Psychology and Physiology from the University of Oxford, which he attended as a Rhodes Scholar, and a BSc in Electrical Engineering from the University of New Brunswick. During his career at Xerox, he has held research management and project leadership positions in colour imaging and systems and has worked on colour printing, image processing, enterprise coherence and standards for colour documents and images.

He is the Xerox representative on the US JPEG 2000 committee and was the Project Editor for Part 6 of the JPEG 2000 standard, which defines the JPEG 2000 file format for compound and document images. He currently chairs the CIE Technical Committee on Archival Colour Imaging. He was founding co-chair of the IS&T Archiving Conference and has given several invited talks on JPEG 2000 to the cultural heritage community. He is the Xerox Principal in the collaboration with the Library of Congress described in the report. He is a Fellow of the Society for Imaging Science & Technology (IS&T) and President of the Inter-Society Color Council.