

# JPEG 2000: Good Things in Smaller Packages

Jonathan W. Lowe

**T**he emerging international standard for image compression and decompression that is commonly known as JPEG 2000 (JP2) has already been adopted by such players as BAE Systems, General Dynamics, International Land Systems, LizardTech, LuraTech, and Northrup Grumman as an efficient means of handling large imagery collections. And, thanks to input from a variety of industry experts, the standard is not just a more tightly compressed version of the familiar JPEG image: its 14 parts also include designs that support secure military radio transmissions, regional decoding, inclusion of GML (Geography Markup Language) for geographic coordinates and object descriptions, and predefined resampling for rapid viewing at different scales. Tuning a JPEG 2000 image file's encodings to match a specific workflow can improve rendering performance even when that file is very large.

### A Complement to Current JPEG Standards

Initially intended as a complement — not a replacement — to the widespread JPEG image format, JPEG 2000's core benefit is far-greater image compression with a minimally detectable loss of image quality. The standard (which is more formally

Modern image formats now come with a huge array of options. One format in particular, however, is offering just the features that many image-using industries need, including the geospatial industry.

known as ISO-15444) is the result of input from a wide variety of industries, so it includes optional features designed to benefit medical imaging, digital photography, animation, military transmission, and, especially relevant, remote sensing and GIS. The Internet Society's memo on *MIME Type Registrations for JPEG 2000* (<http://rfc3745.x42.com/>) offers a concise description.

JPEG 2000 is a relatively new standard intended to create an image coding system for many types of still images (bi-level, gray-level, color, multi-component) with different characteristics (natural images, scientific, medical, remote sensing, imagery, text, rendered graphics. . .) allowing different imaging models (client/server, real-time transmission, image library archival, limited buffer and bandwidth resources, etc.) within a unified system.

Part 1 of JPEG 2000's formal ISO-15444 specification is what one might expect from an image standard based on The Internet Society's description. Namely, it details such core image-compression issues as the specifics of decoding between compressed and reconstructed image data, code-stream syntax for interpreting the compressed image data, file format details, guidance on encoding source data for compression,


and practical advice on implementation.

Compared with the original JPEG specification, JPEG 2000 achieves greater levels of compression with minimal loss thanks to wavelet technology. Original JPEG implementations use a compression technique called discrete cosine transformation (DCT), which breaks up and attempts to compress raw images one 8-kilobyte block at a time, then reassembles them as one image. DCT algorithms process each block independently, so if a line crosses two blocks, that line may appear broken or jagged when the blocks are reassembled. Aggressive compression using DCT results in mathematical data losses — once compressed, the exact original image can never be faithfully reconstructed.

Instead of DCT compression, JPEG 2000 relies on discrete wavelet transformation (DWT) algorithms (see Figure 1). DWT works on the complete image (not 8-kilobyte blocks), leaving no block borders or artifacts. An image compressed using a DWT process can also incur data loss if so specified, but losses are stored in places where the human eye has difficulty recognizing them: DWT conceals the loss using less sharpness and less accuracy of detail. Often capable of compression ratios as great as 200:1 (and frequently 300:1 with such color images as photos), JPEG 2000 consequently speeds the performance of image-processing software by decreasing the volume of data being



Net Results columnist **Jonathan W. Lowe** covers the role of emerging technologies in the exchange of spatial information. Lowe is the owner of Local Knowledge Consulting (Berkeley, California), where he designs and implements spatial Web sites. He can be contacted at [info@giswebsites.com](mailto:info@giswebsites.com).



Courtesy of LizardTech

▲ **Figure 1.** Because the JPEG 2000 wavelet-based technique encodes an image at multiple resolution levels (comparable to pyramid-based schemes but without the file-size overhead), a JP2 image can be decoded in different ways, such as at full resolution, half resolution, and quarter resolution.

transferred between disk and memory, for instance.

Regarding this advance in compression capabilities, it's important to understand both what the JPEG 2000 standard is and what it is not. Specifically, the JPEG 2000 standard specifies how a decoder must work and how a JPEG 2000 file should be structured for both to be compliant. It does not, however, describe how to encode the file in the first place. Devising an encoding technique is left to individuals or vendors once they are given a standard structure target. As an example, the JPEG 2000 specification requires layering of different color versions of documents, but it does not specify how to layer them.

In a recent interview, Carsten Heiermann, president of LuraTech, Inc. of Redwood City, California, explained what JPEG 2000 contains and how the

standard benefits the geospatial industry. He also explained why LuraTech participates in the expensive and lengthy ISO (International Organization for Standardization) standardization process. To help us understand JPEG 2000, Heiermann illustrated the limits of a published standard.

“One base image, say a 300-megabyte TIFF, is converted to the JPEG 2000 format by two different vendors’ image-processing tools with the stipulation that the compressed JPEG 2000 files be no more than 10K in size,” Heiermann explained. “The encoding and compression algorithms that each vendor uses probably differ, but the final two JPEG 2000 files will both meet the JPEG 2000 standard, and so can be read by any other JPEG 2000 viewer or image handler. What’s different? The quality of the compressed

JPEG 2000 images.”

So how do we measure quality? Heiermann explained that people measure quality by two criteria: PSNR (or the signal-to-noise ratio) and subjective criteria (such as what looks better). Evaluating images mathematically will reveal different PSNR levels, but judging an image’s looks remains beyond any artificial intelligence. Each viewer is his or her own judge.

### Emerging Geo-JPEG 2000

Advances in compression are noteworthy, but the remaining 13 parts of ISO-15444 take JPEG imagery functionality quite a bit further. For instance, one part addresses efficient image handling, such as how to open, save, store, and transmit JPEG 2000 images. What’s so complex about opening an image? The answer depends on the size and view scale. What if, for example, the image is several gigabytes in size? Opening the whole file could take some time.

Imagery users with large collections, such as NGA (the National Geospatial-Intelligence Agency, which predicts it will adopt the JPEG 2000 standard by 2010), are particularly sensitive to image-viewing performance. Karen Morley, vice-president of global marketing for LizardTech, clarified that how an image is initially encoded influences how it performs in whichever workflow then utilizes it. For instance, Morley cited NGA’s two specifications for the use of JPEG 2000, one called NSIF/BIIF/NITF Preferred JPEG 2000 Encoding (NPJE) covering a workflow in which users pan through large images at 1:1 resolution, and another called EPJE (“E” for exploitation) involving more distant zooms. In support of differing image-processing workflows, LizardTech software allows users to establish, save, and share any encoding profiles that perform well for certain workflows. The decoding and embedding options detailed below illustrate that the JPEG 2000 standard supports both of these NGA encoding approaches. However, responsibility

rests on vendors and users to implement and select the most appropriate one based on intended workflow.

**Progressive Decoding.** The JPEG 2000 format anticipates variations in image view scales by reserving space in the file for “progressive decoding” (with more efficient sorting than the original JPEG standard). Progressive decoding means that the first bytes of the file contain a viewable image, such as a thumbnail or low-resolution screen view, formally termed a “lower decoding-progression order.” For instance, the first 1–7 percent of the file can contain a screen-fitting, best-resolution, initial view such that users can quickly assess even very large 5–6-GB files at fully zoomed-out scales (see Figure 2a).

**Region Decoding.** JPEG 2000 also anticipates the need for fast performance when zooming to the full raster resolution. The structure of JPEG 2000 files supports “region decoding” such that only the image data needed to render the view extent is decompressed for display (see Figure 2b). In older formats, region decoding was typically accomplished by referencing pre-established tiles or pyramids built into the image file itself. JPEG 2000 instead uses markers and indexes also built into the image file to enable the faster zooming resulting from selective decompression. Explained Heiermann, “It’s still possible to use tiles and mosaics, but this is no longer necessary, though there’s even a way to specify overlapping tiles without matching problems if that’s what you want to do. And with this new format, it’s all inside one file.”

**Embedded GML.** Beyond core compression, the marriage of JPEG 2000 with GML is potentially the most important element of the JPEG 2000 standard for the geospatial industry. The idea is that a JPEG 2000 file can contain not only image data, but also a subset of the textual definitions in GML, allowing image-viewing software to identify the coordinate system in use for a given image; the coordinates of each pixel in the image; and names or types of objects, such as

bridges, houses, cars, tanks, and so on (see Figure 2c). Hovering a mouse over an area of the image linked to GML attributes allows a compliant viewer to access the GML text linked to that location.

Establishing the specific subset of GML for embedding within a JPEG 2000 file is admittedly beyond the scope of the ISO-15444 committee. Fortunately, there is a cooperative relationship between ISO and the Open Geospatial Consortium (OGC, [www.opengeospatial.org](http://www.opengeospatial.org)) surrounding the integration of imagery and attribution.

Says Heiermann, “The JPEG 2000 specification sets aside space for XML content. Now, OGC will decide what to put there, and ISO will reference the OGC decision.”

To this end, geospatial experts such as the European Union Satellite Center, Galdos Systems, and LizardTech are conducting an OGC interoperability experiment that embeds GML within JPEG 2000. The committee will encode a JPEG 2000 image with embedded GML, then each participant will attempt to view the file and its embedded attributes with their own software. This experiment ensures that each participant’s interpretation of the standard works as expected against the same input, and that end users can exchange JPEG 2000 images no matter whose software they deploy.

**Extra Baggage?** After emphasizing the importance of ever-greater compression ratios in the JPEG 2000 effort, these aforementioned options may sound like extra baggage. And the geospatial options are just a subset of the entire collection of options. Even though there are many features, however, the JPEG 2000 files themselves can remain small. Likewise, the specification allows for embedding GML within the JPEG 2000 file or keeping the image and GML separate for sizing or maintenance reasons.

“ISO standard-compliant decoders have to be aware that not all markers are required,” Heiermann explained. For example, biometric criteria in passports are emerging in Germany, with the aim of

storing small, highly compressed passport photos on a chip inside the German passport as soon as October 2005. The image format for these photos will be JPEG 2000, in part because there is very little overhead in the standard. JPEG 2000 image decoders must handle the fact that some markers are missing but that the file is still compliant. Use of small digital images in passports is an emerging ISO standard also surfacing in other European countries and Japan, for instance, because all countries will want to be able to read the passports of visitors using such technology.

### ISO: Small Name, Broad Reach

The JPEG 2000 standard is one of many ISO efforts, all of which carry considerable weight given the organization’s wide reach. Based in Geneva, Switzerland, ISO was founded in 1946 to establish voluntary standards and promote global trade among the approximately 100 member countries. The name ISO is itself a standard of sorts, in that it is not an abbreviation but a word derived from the Greek word *isos*, which means “equal.” By choosing a single-word name, ISO’s founders prevented the creation of multiple abbreviations for their organization as a result of worldwide language variations. For example, the English would have abbreviated “International Organization for Standardization” as IOS, but the French would have abbreviated “Organisation Internationale de Normalisation” as OIN.

LuraTech’s experience as a participant in forging ISO-15444 was one of considerable commitment — both of money and time. In Heiermann’s words, “It was a long way.” The effort began seven years ago and resulted in a published ISO standard in 2001, more than four years later. As a founding member, LuraTech performed compliance testing, co-authored Part 6, and contributed to parts 1 and 2. Participation required LuraTech’s Chief Technology Officer Klaus Jung to travel all over the world to attend workshops with participating

members. Jung continues to work with ISO as the head of delegation for Germany, bringing new suggestions to both ISO and OGC for integration across industries that will maintain compliance with the JPEG 2000 standard.

The original LuraTech product was proprietary but became open following the 2001 ISO launch. As president, Heiermann reports to LuraTech's shareholders, many of whom initially asked, "Why invest in an open standard? Why not stick with your proprietary software?" Heiermann said he was confident that LuraTech's customers would want their technology investment secured by compliance with an ISO standard. "They can always go back to the standard and interoperate with other users without forced reliance on manufacturer-dependent viewers or fear of compliance issues," he said. LuraTech's customers were already asking for open-standard solutions, so the shareholders approved of Heiermann's strategy.

**OGC Vertical Standards.** LizardTech also participated in the ISO-15444 effort initially, but later determined their most effective point of contact to be OGC membership. Since geospatial imagery users often must manage very large files, LizardTech has been able to establish a core competency as a geospatial image-processing technology vendor, without having to serve the much broader general imagery market. In contrast to the OGC's focus, the ISO-15444 JPEG 2000 effort

addresses not only geospatial imagery, but medical imaging, digital photography, motion JPEG, and other markets not directly relevant to LizardTech's business plan.

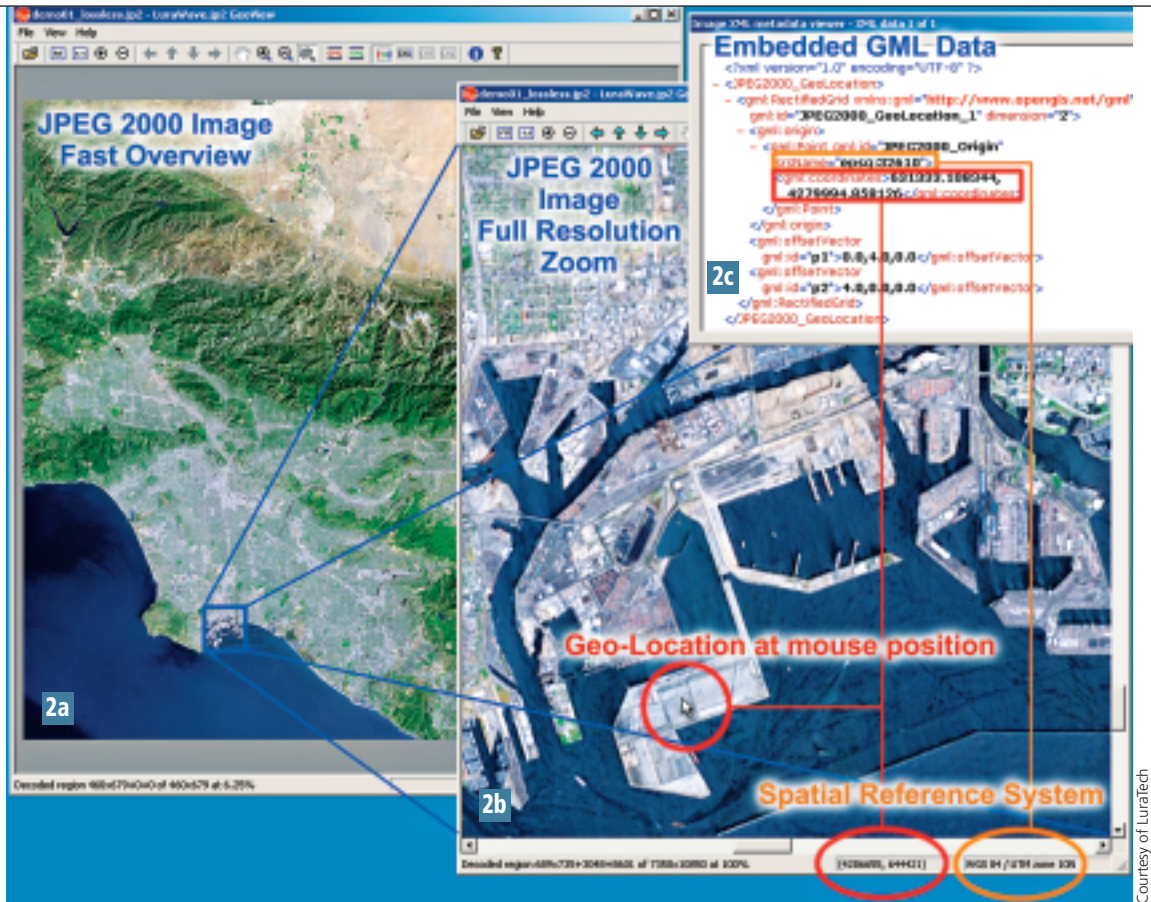
"Participation in the OGC is very expensive," noted Morley, "but it's in LizardTech's and our industry's best interests to make sure that the most advanced technology becomes part of any new imagery standards."

Like Heiermann, Morley also had to explain the value of standards participation to LizardTech's board, which supported a standards investment within LizardTech's area of core competency.

Whether LizardTech's and LuraTech's investments in imagery and geospatial standards will pay off or not depends on their customers' perceptions of the value of ISO and/or OGC compliance and, most importantly, the market timing of those perceptions. (Both LizardTech and LuraTech participate in standards efforts with

the hope of being first to market with standards-compliant products.)

Though JPEG 2000 is not yet widely adopted within the geospatial industry, the trend in that direction seems to be approaching a tipping point. Morley noted initial uptake of JPEG 2000 mainly by Department of Defense customers. For instance, the military's National Image Transmission Format (NITF) version 2.1 supports an embedded JPEG 2000 format (even though some downstream military applications don't yet support this inclusion). And the JPEG 2000 provisions for secure radio transmission seem to dovetail well with the NITF provisions for classification of sensitive intelligence data. According to Heiermann, industries such as medical imaging now simply assume use of JPEG 2000 within such larger standards as DICOM. The geospatial industry may be approaching that same point in 2005, with the military as an early adopter. ☉



**Figure 2a-2c.** A screenshot of LuraTech's LuraWave JP2 GeoView product demonstrating a fast overview of the whole large image (2a), a fast zoom-in on one region (2b), and a box showing GML code embedded in a JPEG 2000 file, including GML coordinates corresponding to the cursor position (2c).

Courtesy of LuraTech