

University of Connecticut Libraries
UConn Libraries Published Works

University of Connecticut

Year 2009

A Status Report on JPEG 2000
Implementation for Still Images: The
UConn Survey

David Lowe*

Michael J. Bennett†

*University of Connecticut - Storrs, david.lowe@uconn.edu

†University of Connecticut - Storrs, michael.bennett@uconn.edu

This paper is posted at DigitalCommons@UConn.

http://digitalcommons.uconn.edu/libr_pubs/19

A Status Report on JPEG 2000 Implementation for Still Images: The UConn Survey

David B. Lowe and Michael J. Bennett; University of Connecticut Libraries; Storrs, CT/USA

Abstract

JPEG 2000 is the product of thorough efforts toward an open standard by experts in the imaging field. With its key components for still images published officially by the ISO/IEC by 2002, it has been solidly stable for several years now, yet its adoption has been considered tenuous enough to cause imaging software developers to question the need for continued support. Digital archiving and preservation professionals must rely on solid standards, so in the fall of 2008 the authors undertook a survey among implementers (and potential implementers) to capture a snapshot of JPEG 2000's status, with an eye toward gauging its perception within this community.

The survey results revealed several key areas that JPEG 2000's user community will need to have addressed in order to further enhance adoption of the standard, including perspectives from cultural institutions that have adopted it already, as well as insights from institutions that do not have it in their workflows to date. Current users were concerned about limited compatible software capabilities with an eye toward needed enhancements. They realized also that there is much room for improvement in the area of educating and informing the cultural heritage community about the advantages of JPEG 2000. A small set of users, in addition, perceived problems of cross-codec consistency and future file migration issues.

Responses from non-users disclosed that there were lingering questions surrounding the format and its stability and permanence. This was stoked largely by a dearth of currently available software functionality, from the point of initial capture and manipulation on through to delivery to online users.

Background

In the fall of 2008, the authors surveyed the status of JPEG 2000 implementation as a still image format among cultural heritage institutions involved in digitization. This sample was taken from August 27, 2008 through October 31, 2008. Respondents totaled 161, with the overwhelming majority coming from academic research libraries [1]. The following focuses primarily on the major issues broached by respondents, examines current use and perceived barriers to the standard's adoption, and proposes recommendations towards JPEG 2000's greater utilization within the cultural heritage community.

Migration Concerns

Codec Inconsistency

An interesting opinion among respondents focused on perceived codec inconsistencies among software vendors. Coupled with this were migration concerns based upon such

inconsistencies and also the general nature of JPEG 2000's currently limited adoption, and future migration toolkits:

"Lack of consistency across codecs (e.g. Aware, Kakadu) for creating JPEG 2000s." [written in response to the question of drawbacks to JPEG 2000 implementation]

"JPEG 2000 is a great format, but the main problem resides in acceptance not only in the repository level but also commercially. To have a fully robust digital archival format we will require good migration software for when it becomes obsolete. If it becomes commonly used (such as TIFF) migration software will work smoother with less errors as they will not have to necessarily be homegrown."

"It's a new format with an unproven history or migration."

Codec concerns may be ameliorated to some extent when put into the larger context of the standard itself. JPEG 2000 is a fully documented and open standard and as such is available for software developers of all types (vendors and freeware authors) to write encoders for. Much like capture hardware's vagaries of unique sensor filters and device-specific profiles, software encoders are similarly geared around their creator's best perceptions of fidelity in the production of these files.

Perhaps the most important residual of the standard's openness in this regard, however, is the fact that decoding of valid JPEG 2000 files remains transparent regardless of the encoder used. In this way, migration concerns may be mitigated to a degree as developers today and into the future can be assured access to the standard in order to write such applications. Yet, the questions of future prevalence and quality of software toolkits for JPEG 2000 mass migration remain foremost in many practitioners' minds.

Visually Lossless, Mathematically Lossless

The possibility of visually lossless (mathematically lossy) JPEG 2000 compression as an archival storage option has recently begun to gain traction particularly in areas of large scale TIFF migration at both Harvard and the Library of Congress [2][3]. Mass digitization projects such as the Open Content Alliance have also adopted visually lossless JPEG 2000 as their archival standard [4]. Among survey respondents there was a divergence of opinion on the idea as many felt that possible future migration costs for moving out of the standard may not make up for the real benefits of storage efficiencies realized today:

"I have some concerns that once we start going down a slope of compromising images what the potential of it being accentuated after multiple migrations possibly with different

lossy compression schemes. Considering the relative cost of space I don't think it is a worthwhile risk."

"...visually lossless but technically lossy compression is not a good basis for later format migrations."

Here, notions of fidelity between hardware and software, born digital vs. converted digital were used to try to strike a balance in the decision making process:

"...visually lossless is fine. The only reason to use mathematically lossless would be in conversion of born digital materials where there hasn't already been loss due to analog to digital conversion. For analog materials, the loss inherent in lossy jp2 is minimal compared to sensor and sampling error from the original scanning."

Lossless JPEG 2000 compression, though in fact lossless at the bit-level upon decoding, still elicited its own migration anxieties:

"I have little concern over lossless compression other than prominence and easy migration. It adds another level of encoding which could very well complicate future migrations (especially if one is missed) unless it is common and well documented. Again the availability of good migration software is useful."

Finally, as one respondent phrased it, nothing may ever be perfect:

"One problem with the widespread acceptance of .jp2 is the fear of future migration. However, I have heard that migration projects of tiff formats haven't gone smoothly either."

Current Use Scenarios Point to Advantages

JPEG 2000's support of both true lossless and a wide array of visually lossless (lossy) compression enjoyed broad use at many of the responding institutions. Scalable storage savings through the standard's comparative file size economy to TIFF and JPEG 2000's flexible individual file rendering on the web were focal reason for its favor. A sample of the more intricate use scenarios included the following responses:

"We produce TIFF files for our new photography, and for some projects we produce lossless JP2 files that we class as "master". In these cases we discard the original TIFF and the losslessly-encode file serves as master and delivery image. For some projects we save uncompressed-TIFF files, classes as "masters", and also produce a lossy compressed JP2 file for delivery purposes. A third common workflow produces TIFF files that are used to produce conservatively, but lossy-compressed JP2 files for delivery. The TIFF files are not saved and the lossy JP2 files serve as masters and as delivery images."

"Yes, we have started to make lossless JPEG 2000 images for some collections where we would have previously saved (LONG term) uncompressed TIFF and lossy JP2 for delivery.

We like keeping a single file that can be used as master and deliverable, and the fidelity is equivalent to an uncompressed TIFF."

"For our large-scale book scanning projects (published materials from circulating collections) we are saving conservatively, but lossy compressed color JP2 images. This is a high quality, but lower cost and high volume service and we need to take advantage of the power of lossy compression to reduce our file storage costs."

Misperceived Disadvantages Affect Adoption

Incorrect assumptions on the standard were, however, common throughout the survey and revealed a real need for better education and understanding. Common threads included a lack of trust in JPEG 2000's lossless compression as being truly lossless. Others believed that such lossless compression did not confer significant file size savings in comparison to uncompressed TIFF or that JPEG 2000 did not support higher bit depth images. A small number of respondents continued to make the unfortunate association of JPEG 2000 with JPEG as two lossy-only standards, a belief that has hounded JPEG 2000 in particular since its inception. Also false notions of JPEG 2000 as being proprietary in nature and not fully documented lead some to believe that software tools would forever be scarce, expensive, and never open source.

Compression Choices in the Context of other Standards and Best Practices

As part of the "visually lossless" (that is, slightly lossy) vs. mathematically lossless compression debate, it is important to emphasize that, although any whiff of lossiness may raise eyebrows among some colleagues, there may be perfectly reasonable situations for choosing the visually lossless route. The major case in point is in mass digitization efforts, converting print pages to digital images.

Consider the Digital Library Federation's (DLF) "Benchmark for Faithful Digital Reproductions of Monographs and Serials," which specifies 600dpi TIFF, compressed losslessly, but bitonal for text or line art, which represents the bulk of historical print materials, but far from everything. For more complex print situations, such as grayscale photos or color illustrations, the Benchmark recommends (but does not require) progressing up to grayscale and color as needed, albeit at only 300dpi [5]. Genealogically, this print capture standard essentially developed from the joint Michigan and Cornell Making of America Project (Phase I, circa 1996) and it has been the one implemented in major book digitization projects among DLF members and beyond [6].

More recently, the Internet Archive has developed its visually lossless JPEG 2000-based benchmark in concert with partner institutions, who helped settle on an all-color alternative, which eliminates the human factor of bit-level decisions from the moment of capture [7]. Surely anyone familiar with an activity even remotely as repetitive as scanning a book page-by-page can appreciate the fact that, in a three-level system, many color or grayscale pages of material will often slip through as the default bitonal by mistake. Thus, considering the choice of having either a visually lossless color JPEG 2000 or a losslessly compressed TIFF bitonal image of a page that should have been in color in

order to render features faithfully, then clearly the visually lossless compromise comes out ahead. This is not to say that in certain situations some may not still opt for bitonal, on the low end, or, on the high end, for mathematically lossless throughout for a particular object or set of objects. In a nutshell, the visually lossless color is at least as satisfactory as bitonal in the grand scheme of things, where realistically files need to fit on the servers allotted.

One of the primary goals driving the development of the JPEG 2000 standard was the unfilled need for scalability that would extend from a high resolution archival master to a lower resolution, web-deliverable browser image. The bifurcating path of preserving a bulky, high resolution TIFF for the master, then running processes to extract derivative files for user access is inherently inefficient. The survey results were striking in that, of the survey respondents who considered themselves implementers, the majority reported that they use JPEG 2000 to provide web access to material, while only a minority used it for archiving images. This was despite the fact that one of the main problems to date with the standard is the lack of browser support, while one of the chief advantages is its more efficient file size for high resolution mastering. A more efficient model than the TIFF/derivative method would be for the JPEG 2000 format to do both the heavy lifting of high resolution archiving as well as the delivery to the user.

Current Tools & Browser Support

Adobe Photoshop with its free optional plug-in proved to be the most utilized JPEG 2000 file creation tool among practitioners. Feelings expressed on this score were that the plug-in was easy to use, could be integrated into batch processing, but could also be slow.

Interestingly however, beginning in 2007, Adobe themselves have questioned their own continued development of the plug-in in light of cameras not entering the market with native JPEG 2000 support, coupled with the standard's assumed minimal adoption among Photoshop users [8]. To date, Adobe plans to keep shipping the plug-in with its newest Photoshop versions, but will do so most likely as an optional installation (personal communication with John Nack of Adobe, February 19, 2009).

The digital collection management software, CONTENTdm, with its built-in JPEG 2000 converter was also a popular utility. In this case the tool's primary reported use, the ingest and subsequent conversion of pre-created high quality JPEG or TIFF archivals into access JPEG 2000 files, pointed to the fact that much of JPEG 2000's use at least within this community was as an access format.

Frustration was expressed on the current lack of native browser support for the standard. This focused primarily on the resulting server-side requirements that are needed in order to take advantage of the standard's flexible, zoom and panning capabilities from single JPEG 2000 files. In most cases, this dedicated server layer interprets a zoom scale request from the browser, then converts the stored JPEG 2000 into a format like JPEG or BMP that the browser can support and finally render to users. Respondent's comments included:

“Currently, very little client software and very few repositories seem to take advantage of the jpip protocol. This means that jpeg2000 images either need to be transformed on the server (for different regions, resolutions, etc.) to jpeg for example, or the whole image downloaded by the client before displaying native jpeg2000. It also means that features such as quality layers and region of interest are less likely to be taken advantage of as this information should be client/user preference and is difficult to efficiently communicate without a dedicated client/server protocol.”

Yet, among some there was confidence that the browsers would eventually come around. Indeed, though native support is currently absent from Internet Explorer, Mozilla/Firefox, Safari, and Chrome, the QuickTime plug-in for each can render JPEG 2000, though only at one zoom level. The authors feel that it is imperative for browser developers to bake into their code native JPEG 2000 support that includes the full range of image manipulations that the standard enables, such as broad panning and deep zooming. Since part of the design aspect of wavelet compression schemes, like that of JPEG 2000, involves pushing more computing to the user's viewing device and its software, and since the major developers involved tend to want their browsers' codebases to travel light as a competitive advantage, there exists a threshold for implementation that has not been kind to the image standard. Setting aside for a moment the unhandy option of dedicated JPEG 2000 servers, the extra code required in the browser has relegated JPEG 2000 to the realm of extensions and add-ons, most of which, like QuickTime, serve a much broader audience and do not take the potential functionality much beyond a zoom level or two.

IPR Barriers: Genuine Paper Chase Threats vs. Paper Tigers

To a limited extent, the UConn survey revealed that patent claims surrounding JPEG 2000 remain a concern to some in our community. The blogosphere, not surprisingly, can go even farther, for example: “JPEG 2000 is doomed to failure because of the patent issue [9].” Yet even professionals much closer to the inner workings of the standard have viewed the issue as a significant barrier as recently as 2005 [10]. It is important to keep in mind several points while considering the legal implications of choosing JPEG 2000.

First, the Intellectual Property Rights (IPR) landscape of networked computing is fraught with or at least constantly borders potentially litigious issues with practically any conceivable standard. For still images alone, the earlier JPEG and GIF file formats have been no stranger to legal entanglements. In the case of GIF, what had been a free and open format became litigious when the patent holders changed their minds about that formerly open model [11]. Technically, it was the Lempel-Ziv-Welch (LZW) compression algorithm that was the specific patent card played, but the format, including its LZW compression scheme, had been freely open since 1987 when the patent holders suddenly exerted their rights in 1993, resulting in what is referred to as a “submarine patent claim.” In the case of JPEG, in 2002, a company claimed patent rights despite the existence of “prior art,” or public evidence that the company's claim was not in fact

original. Since then, in 2007, a patent mill has sought to squeeze the last drops of revenue from the final remaining patent recognized by the U.S. Patent Office from the mostly expired JPEG chest, and again prior art appears to be rendering this last claim invalid as well. In the final analysis, unpredictable human behavior will always threaten progress, and the best an organization can do is to take prudent steps to minimize the threat.

Fortunately, the JPEG Committee has indeed taken such proactive measures by having all contributors to the JPEG 2000 standard itself sign...

agreements by which they provide free use of their patented technology for JPEG 2000 Part 1 applications. During the standardization process, some technologies were even removed from Part 1 because of unclear implications in this regard. Although it is never possible to guarantee that no other company has some patent on some technology, even in the case of JPEG, unencumbered implementations of JPEG 2000 should be possible [12].

However, as it so happened, one of the signers did file a suit against a competitor, claiming patent infringement, but the District Court judge in the case ruled the patent claim invalid [13]. JPEG 2000, then, not only has the benefit of a foundation cleared for patent issues by its designers, but it has also thwarted an offensive by one of those designers—one who had been most likely to succeed in the crucial prior art category—and now enjoys a record of patent claims resistance. Moreover, the JPEG committee remains vigilant, seeking to identify any IPR claims regarding JPEG 2000, and regularly solicits information toward this end at each triannual meeting in their ongoing standards work. By documenting these claims (or more accurately the lack thereof) via regular updates, the case for future GIF-like submarine patent claims is severely curtailed, if not nullified.

Recommendations

1. Compile an implementation registry, which would include contact information, of JPEG 2000 related digital projects in cultural heritage institutions (similar to the current METS and PREMIS implementation registries at the Library of Congress) [14][15].

2. Suggest the creation of a new set of JPEG 2000 benchmarks (e.g. NDNP's profile for newspapers) that could be referenced in collaborative projects, vendor RFPs, and grant applications [16]. Outline the standard's appropriate use as an archival and access solution by format, including:

- a. General collections books (e.g. Internet Archive: lossy color) [17]
- b. Special collections books (e.g. lossless color)
- c. Photographs
- d. Maps
- e. Image files migrated from other raster formats

3. Vet the above suggested benchmarks through a competent collaborative body, such as DLF, and pursue its stamp of approval.

4. Have the collaborative body identify and empower a liaison from among imaging experts to serve as an advocate for

JPEG 2000 to browser developers and imaging software developers.

5. Better educate the cultural heritage community about the soundness and advantages of JPEG 2000 in the context of possible format benchmarks.

References

- [1] D. Lowe, and M. Bennett, "Digital Project Staff Survey of JPEG 2000 Implementation in Libraries" (2009), (retrieved 3.20.09). *UConn Libraries Published Works*. Paper 16. http://digitalcommons.uconn.edu/libr_pubs/16
- [2] S. Abrams, S. Chapman, D. Flecker, S. Kreigsmann, J. Marinus, G. McGath, and R. Wendler, "Harvard's Perspective on the Archive Ingest and Handling Test," *DLib Magazine*, 11, 12 (2005), (retrieved 3.18.09) <http://dlib.org/dlib/december05/abrams/12abrams.html>
- [3] "Preserving Digital Images (January-February 2008)" - Library of Congress Information Bulletin. (retrieved 3.17.2009) <http://www.loc.gov/loc/lcib/08012/preserve.html>
- [4] R. Miller, "Internet Archive (IA): Book Digitization and Quality Assurance Processes," confidential document for IA partners; permission to reference from R. Miller 3/16/2009.
- [5] Digital Library Federation, "Benchmark for Faithful Digital Reproductions of Monographs and Serials," (retrieved 3.18.2009), <http://purl.oclc.org/DLF/benchrepro0212>
- [6] E. Shaw and S. Blumson, "Making of America: Online Searching and Page Presentation at the University of Michigan," *D-Lib Magazine*, July/August (1997), (retrieved 3.18.2009), <http://www.dlib.org/dlib/july97/america/07shaw.html>
- [7] *ibid.*, R. Miller.
- [8] "John Nack on Adobe: JPEG 2000 - Do you use it?" (retrieved 3.17.2009), http://blogs.adobe.com/jnack/2007/04/jpeg_2000_do_yo.html
- [9] B. Dowling on February 18, 2007 04:07 AM on J. Atwood's blog "Coding Horror: programming and human factors," (retrieved 3.18.2009), <http://www.codinghorror.com/blog/archives/000794.html>
- [10] Y. Nomizu, "Current status of JPEG 2000 Encoder Standard," (retrieved 3.18.2009), http://www.itscj.ipsj.or.jp/forum/forum20050907/5Current_status_of_JPEG%202000_Encoder_standard.pdf
- [11] M. C. Battilana, "The GIF Controversy: A Software Developer's Perspective," (retrieved 3.18.2009), <http://www.cloanto.com/users/mcb/19950127giflwz.html>
- [12] D. Santa Cruz, T. Ebrahimi, and C. Christopoulos, "The JPEG 2000 Image Coding Standard," (retrieved 3.18.2009), <http://www.ddj.com/184404561>
- [13] "Gray Cary Successfully Represents Earth Resource Mapping in Patent Dispute," from <http://www.businesswire.com> dated March 27, 2004, (accessed 3.18.2009 via LexisNexis: <http://www.lexisnexis.com/>)
- [14] "METS Implementation Registry: Metadata Encoding and Transmission Standard (METS) Official Web Site," (retrieved 3.23.09), <http://www.loc.gov/standards/mets/mets-registry.html>
- [15] "PREMIS Implementation Registry - PREMIS: Preservation Metadata Maintenance Activity (Library of Congress)," (retrieved 3.23.09), <http://www.loc.gov/standards/premis/premis-registry.php>
- [16] "NDNP JPEG 2000 Profile," (retrieved 3.23.09), <http://www.loc.gov/ndnp/pdf/JPEG2kSpecs.pdf>
- [17] *ibid.*, R. Miller.

Authors Biographies

David Lowe serves as Preservation Librarian at the University of Connecticut. A 1996 graduate of the School of Information at the University of Michigan, Lowe has worked with both analog and digital media as a Preservation staff member at the University of Michigan (1994-1997) and at Columbia University (1997-2005), before coming to UConn (2005-), where his primary duties now involve preservation issues surrounding local digital collections.

Michael J. Bennett is Digital Projects Librarian & Institutional Repository Coordinator at the University of Connecticut. There he manages digital reformatting operations while overseeing the University's institutional repository. Previously he has served as project manager of Digital Treasures, a digital repository of the cultural history of Central and Western Massachusetts and as executive committee member for Massachusetts' Digital Commonwealth portal. He holds a BA from Connecticut College and an MLIS from the University of Rhode Island.