

ISO/IEC JTC1/SC29/WG1
(ITU-T SG8)

Coding of Still Pictures

JBIG

Joint Bi-level Image
Experts Group

JPEG

Joint Photographic
Experts Group

TITLE: An analytical study of JPEG 2000 functionalities
Paper to be published in the proceedings of ICIP 2000.

SOURCE: Diego Santa Cruz and Touradj Ebrahimi
Diego.SantaCruz@epfl.ch
Touradj.Ebrahimi@epfl.ch

PROJECT: JPEG 2000

STATUS: Information

**REQUESTED
ACTION:**

DISTRIBUTION: WG1 delegates, WG1 website and reflectors

Contact:

ISO/IEC JTC1/SC29/WG1 Congener - Dr. Daniel Lee
Hewlett-Packard Company, 11000 Wolfe Road, MS 42U0, Cupertino, CA 95014
Tel: +1 408 447 4160, Fax: +1 408 447 2842, E-mail: daniel_lee@

Diego Santa-Cruz and Touradj Ebrahimi, **An analytical study of JPEG 2000 functionalities**. To be published in Proc. of the International Conference on Image Processing (ICIP), Vancouver, Canada, September 10-13, 2000.

Copyright 2000 IEEE. Published in the 2000 International Conference on Image Processing (ICIP-2000), scheduled for September 10-13, 2000 in Vancouver, BC. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the IEEE. Contact: Manager, Copyrights and Permissions / IEEE Service Center / 445 Hoes Lane / P.O. Box 1331 / Piscataway, NJ 08855-1331, USA.
Telephone: + Intl. 908-562-3966.

AN ANALYTICAL STUDY OF JPEG 2000 FUNCTIONALITIES

Diego Santa-Cruz and Touradj Ebrahimi

Signal Processing Laboratory
Swiss Federal Institute of Technology, CH-1015 Lausanne, Switzerland

ABSTRACT

JPEG 2000, the new ISO/ITU-T standard for still image coding, is about to be finished. Other new standards have been recently introduced, namely JPEG-LS and MPEG-4 VTC. This paper compares the set of features offered by JPEG 2000, and how well they are fulfilled, versus JPEG-LS and MPEG-4 VTC, as well as the older but widely used JPEG and more recent PNG. The study concentrates on the set of supported features, although lossless and lossy progressive compression efficiency results are also reported. Each standard, and the principles of the algorithms behind them, are also briefly described. As the results show, JPEG 2000 supports the widest set of features among the evaluated standards, while providing superior rate-distortion performance.

1. INTRODUCTION

JPEG 2000 [1] will be the next ISO/ITU-T standard for compression of still images. Effort has been made to make this new standard suitable for today's and tomorrow's applications by providing features unavailable in previous standards, but also by providing more efficient support for features that are covered by them. A legitimate question would be: What are the features offered by JPEG 2000 but also how well are they fulfilled when compared to other standards offering the same features. This paper aims at providing an answer to this simple but somewhat complex question.

2. OVERVIEW OF STILL IMAGE CODING STANDARDS

For the purpose of this study we compare the coding algorithm in the JPEG 2000 standard to the following three standards: JPEG [2], MPEG-4 Visual Texture Coding (VTC) [3] and JPEG-LS [4]. In addition, we also include PNG [5]. The reasons behind this choice are as follows. JPEG is one of the most popular coding techniques in imaging applications ranging from Internet to digital photography. Both MPEG-4 VTC and JPEG-LS are very recent standards that start appearing in various applications. It is only logical to compare the set of features offered by JPEG 2000 standard not only to those offered in a popular but older standard (JPEG), but also to those offered in most recent ones using newer state-of-the-art technologies. Although PNG is not formally a standard and is not based on state-of-the-art techniques, it is becoming increasingly popular for Internet based applications. PNG is also undergoing standardization by ISO/IEC JTC1/SC24 and will eventually become ISO/IEC international standard 15948.

Although JPEG 2000 supports coding of bi-level and paletted color images, we restrict ourselves to continuous tone, since it is one of the most popular image types. Other image coding standards are JBIG [6] and JBIG2 [7]. Although these are known for

providing very good performance for bi-level images, they do not support an efficient coding of continuous tone images with a large enough number of levels. Since this paper concentrates on the latter, JBIG and JBIG2 are not considered. Other popular defacto standards for coding continuous tone images are GIF and FlashPix. GIF is limited to 8 bit paletted images and therefore is not considered here. FlashPix is based on JPEG and is therefore more of a file format than a coding standard and is not considered in this paper either. In the following a brief explanation of the principles behind the algorithms used in each of these standards is given.

2.1. JPEG

This is the very well known ISO/ITU-T standard created in the late 1980s. There are several modes defined for JPEG [2], including baseline, lossless, progressive and hierarchical. Baseline mode is the most popular and supports lossy coding only. It is based on the 8x8 block DCT, zig-zag scanning, uniform scalar quantization and Huffman coding. The lossless mode is not popular but provides for lossless coding, but not lossy. It is based on a predictive scheme and Huffman coding. It should be noted that the lossy and lossless modes of JPEG are based on totally different algorithms. The progressive and hierarchical modes of JPEG are both lossy and differ only in the way the DCT coefficients are coded or computed, respectively, when compared to the baseline mode. They allow a reconstruction of a lower quality or lower resolution version of the image, respectively, by partial decoding of the compressed bitstream. Progressive mode encodes the quantized coefficients by a mixture of spectral selection and successive approximation, while hierarchical mode uses a pyramidal approach to computing the DCT coefficients in a multi-resolution way.

2.2. MPEG-4 VTC

MPEG-4 Visual Texture Coding (VTC) is the algorithm used in MPEG-4 standard [3] in order to compress the texture information in photo realistic 3D models. As the texture in a 3D model is similar to a still picture, this algorithm can also be used for compression of still images. It is based on the discrete wavelet transform (DWT), scalar quantization, zero-tree coding and arithmetic coding. MPEG-4 VTC supports SNR scalability through the use of different quantization strategies: single (SQ), multiple (MQ) and bi-level (BQ). SQ provides no SNR scalability, MQ provides limited SNR scalability and BQ provides generic SNR scalability. Resolution scalability is supported by the use of band-by-band scanning (BB), instead of the traditional zero-tree scanning (tree-depth, TD), which is also supported. MPEG-4 VTC also supports coding of arbitrarily shaped objects, by the means of a shape adaptive DWT, but does not support lossless coding.

2.3. JPEG-LS

JPEG-LS [4] is the latest ISO/ITU-T standard for lossless coding of still images and which also provides for “near-lossless” coding. Part-I, the baseline system, is based on adaptive prediction, context modeling and Golomb coding. In addition, it features a flat region detector to encode these in run-lengths. Part-II will introduce extensions such as an arithmetic coder, but is still under preparation. This algorithm was designed for low-complexity while providing high compression. However it does not provide for scalability, error resilience or other additional functionality.

2.4. PNG

Portable Network Graphics (PNG) [5] is a W3C recommendation for coding of still images which has been elaborated as a patent free replacement for GIF, while incorporating more features than this last one. It is based on a predictive scheme and entropy coding. The entropy coding uses the Deflate algorithm of the popular Zip file compression utility, which is based on LZ77 coupled with Huffman coding. PNG is capable of lossless compression only and supports gray scale, paletted color and true color, an optional alpha plane, interlacing and other features.

2.5. JPEG 2000

JPEG 2000 is still under development, although Part I (the core system) [1] is technically frozen and scheduled to reach International Standard (IS) status in December 2000. It is based on the discrete wavelet transform (DWT), scalar quantization, context modeling, arithmetic coding and post-compression rate allocation. The entropy coding is done in blocks, typically 64x64, inside each sub-band. The DWT can be performed with reversible filters, which provide for lossless coding, or non-reversible filters, which provide for higher coding efficiency without the possibility to do lossless. The coded data is organized in so called *layers*, which are quality levels, using post-compression rate allocation and then output to the code-stream in packets. JPEG 2000 provides for resolution, SNR and position progressivity, or any combination of them, parseable code-streams, error-resilience, arbitrarily shaped region of interest, random access (to the sub-band block level), lossy and lossless coding, etc., all in a unified algorithm.

3. COMPARISON METHODOLOGY

One important concern in coding techniques is that of compression efficiency which is still one of the top priorities in the design of imaging products. In a previous study, we devoted a special attention to compression efficiency [8]. However, we report lossless and lossy progressive compression efficiency results to evaluate how well the algorithms code different types of imagery and how well progressive coding is supported. Most applications also require other features in a coding algorithm than simple compression efficiency. This is often referred to as functionalities. Examples of such functionalities are resiliency to residual transmission errors that occur in mobile channels for instance. In the next section we summarize the results of the study as long as the considered functionalities are concerned.

4. RESULTS

The algorithms have been evaluated with seven images from the JPEG 2000 test set, covering various types of imagery. The images “bike” (2048x2560) and “cafe” (2048x2560) are natural, “cmpnd1” (512x768) and “chart” (1688x2347) are compound documents consisting of text, photographs and computer graphics, “aerial2” (2048x2048) is an aerial photography, “target” (512x512) is a computer generated image and “us” (512x448) an ultra scan. All these images have a depth of 8 bits per pixel.

The software implementations used for coding the images are the JPEG 2000 Verification Model (VM) 6.1 (ISO/IEC JTC1/SC29/WG1 N 1580), the MPEG-4 MoMuSys VM of Aug. 1999 (ISO/IEC JTC1/SC29/WG11 N 2805), the Independent JPEG Group JPEG implementation (<http://www.ijg.org>), version 6b, the SPMG JPEG-LS implementation of the University of British Columbia (<http://spm.ece.ubc.ca>), version 2.2, the Lossless JPEG codec of Cornell University (<ftp://ftp.cs.cornell.edu/pub/multimed>), version 1.0, and the libpng implementation of PNG (<ftp://ftp.uu.net/graphics/png>), version 1.0.3.

4.1. Lossless compression

Table 1 summarizes the lossless compression efficiency of lossless JPEG (L-JPEG), JPEG-LS, PNG and JPEG 2000 for all the test images. For JPEG 2000 the reversible DWT filter, referred to as J2K_R, has been used. In the case of L-JPEG optimized Huffman tables and the predictor yielding the best compression performance have been used for each image. For PNG the maximum compression setting has been used, while for JPEG-LS the default options were chosen. MPEG-4 VTC is not considered, as it does not provide a lossless functionality.

It can be seen that in almost all cases the best performance is obtained by JPEG-LS. JPEG 2000 provides, in most cases, competitive compression ratios with the added benefit of scalability. PNG performance is similar to the one of JPEG 2000. As for lossless JPEG, it does not perform as well as the other, more recent, standards. One notable exception to the general trend is the “target” image, which contains mostly patches of constant gray level as well as gradients. For this type of images, PNG provides the best results, probably because of the use of LZ77. Another exception is the “cmpnd1” image, in which JPEG-LS and PNG achieve much larger compression ratios. This image contains, for the most part black text on a white background. In average PNG performs the best, although this is solely due to the very large compression ratio it achieves on “target”. However, JPEG-LS provides the best compression ratio for most images.


This shows that as far as lossless compression is concerned, JPEG 2000 seems to perform reasonably well in terms of its ability to efficiently deal with various types of images. However, in specific types of images such as “cmpnd1” JPEG 2000 is outperformed by far in JPEG-LS. This result is even more striking noting that JPEG-LS is a significantly less complex algorithm.

4.2. Progressive compression

Figure 1 depicts the average rate-distortion behavior obtained by applying progressive compression schemes studied in this paper on the set of tested images. For JPEG 2000, results for reversible and non-reversible DWT filters are shown, referred to as J2K_R and J2K_{NR} respectively, with a SNR scalable bitstream. For MPEG-4 VTC the results have been generated using multiple quantization

Table 1. Lossless compression ratios.

	J2K _R	JPEG-LS	L-JPEG	PNG
bike	1.77	1.84	1.61	1.66
cafe	1.49	1.57	1.36	1.44
cmpnd1	3.77	6.44	3.23	6.02
chart	2.60	2.82	2.00	2.41
aerial2	1.47	1.51	1.43	1.48
target	3.76	3.66	2.59	8.70
us	2.63	3.04	2.41	2.94
average	2.50	2.98	2.09	3.52

**Fig. 1.** PSNR corresponding to average RMSE, of all test images, for each algorithm when performing lossy decoding at 0.25, 0.5, 1 and 2 bpp of the same progressive bitstream.

(MQ) and targeting the tested bitrates. In the case of JPEG the progressive mode has been used where the coefficients are encoded by successive refinement, and is referred to as P-JPEG. The results shown are also typical, for all the tested images.

As it can be clearly seen, progressive lossy JPEG 2000 outperforms all other schemes. The progressive lossless JPEG 2000 does not perform as well, mainly due to the use of reversible wavelet filters, however a lossless version of the image remains available after compression, which can be of significant value to many applications (archiving, medical, etc.). MPEG-4 VTC provides results comparable to those of JPEG 2000 with the reversible filter at high bitrates, although at lower ones a difference appears. As for progressive JPEG it is outperformed by far by the other algorithms, as expected for a relatively old standard. Because of the nature of the default rate allocation algorithm in the JPEG 2000 VM software the non-progressive variants of JPEG 2000 would be practically identical to the progressive ones.

It is also worth noting that when the above results are compared to the non-progressive ones shown in [8] it is clearly seen that JPEG 2000's and MPEG-4 VTC's compression performance is not adversely affected when SNR scalable bitstreams are generated. On the contrary, this can not be said about JPEG.

4.3. Error resilience

In order to evaluate the error resilience features offered by the different standards, we have simulated a transmission channel with random errors and evaluated the average reconstructed image qual-

Table 2. PSNR, in dB, corresponding to average RMSE, of 200 runs, of the decoded "cafe" image when transmitted over a noisy channel with various bit error rates (ber) and compression bitrates, for JPEG baseline and JPEG 2000 (J2K).

bpp		ber: 0	ber: 1e-6	ber: 1e-5	ber: 1e-4
0.25	J2K	23.06	23.00	21.62	16.59
	JPEG	21.94	21.79	20.77	16.43
0.5	J2K	26.71	26.42	23.96	17.09
	JPEG	25.40	25.12	22.95	15.73
1.0	J2K	31.90	30.75	27.08	16.92
	JPEG	30.34	29.24	23.65	14.80
2.0	J2K	38.91	36.38	27.23	17.33
	JPEG	37.22	30.68	20.78	12.09

ity after decompression. Table 2 shows the results for JPEG 2000, with the non-reversible filter, and JPEG baseline. JPEG-LS and MPEG-4 VTC could not be evaluated since the software did not offer proper error resilience support. As for PNG, the comparison is not applicable, since this format only supports error detection, not concealment. Due to lack of space only the results of the cafe image are shown, however, the behavior is very similar for the other images. In the case of JPEG the results have been obtained by using the maximum amount of restart markers, which amounts to an overhead of less than 1%. In the case of JPEG 2000 the sensitive packet head information has been moved to the bitstream header (using a PPM marker) and the entropy coded data has been protected by the regular termination of the arithmetic coder combined with the error resilient termination and segment symbols. The overhead of these protections amount also to less than 1%. In both cases the bitstream header is transmitted without errors.

As it can be seen, the reconstructed image quality under transmission errors is higher for JPEG 2000 than JPEG, across all encoding bitrates and error rates. However, at low bitrates (0.25 and 0.5) the quality of JPEG 2000 decreases more rapidly than JPEG as the error rate increases, although the absolute quality is always higher. Concerning the visual quality at moderately low error rates (i.e. 1e-6), that of JPEG 2000 is much higher when compared to JPEG. In fact, the artifacts created by transmission errors under JPEG 2000 are of the same nature as those created by quantization. In the case of JPEG, when a transmission error occurs it is often entire 8x8 blocks that will be missing and/or misplaced and the bottom of the image will often be missing as well.

It should also be noted that at higher error rates (i.e. 1e-4), the reconstructed image quality in JPEG 2000 is almost constant across all bitrates. This is due to the fact that in JPEG 2000 each sub-band block is coded by bitplanes. When the error rate is high enough almost all blocks are affected in the most significant bitplanes, which are transmitted first. When a particular bitplane is affected in a block, lower bitplanes can not be decoded and are therefore useless. In the case of JPEG the problem is even worse: the higher the encoding bitrate the lower the decoded quality. This can be explained by the fact that when a 8x8 block is affected by a transmission error the entire block is basically lost. The higher the encoding bitrate, the more bits it takes to code a block, and therefore the probability of a block being hit by an error and lost is higher, for the same bit error rate. In other words, in JPEG the density of error protection decreases with an increase in bitrate.

Table 3. Functionality matrix. A “+” indicates that it is supported, the more “+” the more efficiently or better it is supported. A “-” indicates that it is not supported.

	JPEG 2000	JPEG-LS	JPEG	MPEG-4 VTC	PNG
lossless compression performance	+++	++++	+ ^a	-	+++
lossy compression performance	+++++	+	+++	++++	-
progressive bitstreams	+++++	-	++ ^b	+++	+
Region Of Interest (ROI) coding	+++	-	-	+ ^c	-
arbitrary shaped objects	-	-	-	++	-
random access	++	-	-	-	-
low complexity	++	+++++	+++++	+	+++
error resilience	+++	++	++	+++	+
non-iterative rate control	+++	-	-	+	-
genericity ^d	+++	+++	++	++	+++

^aOnly using the lossless mode of JPEG.

^bOnly in the progressive mode of JPEG.

^cTile-based only.

^dAbility to efficiently compress different types of imagery across a wide range of bitrates.

4.4. Functionality

Table 3 above summarizes the results of the comparison of different algorithms from a functionality point of view. The table clearly shows that from this perspective, JPEG 2000 is the standard offering the richest set of features in an efficient manner and within an integrated algorithmic approach. Although some of the rows in this table are self-explanatory, others deserve some comments.

MPEG-4 VTC, as JPEG 2000, is able to produce progressive bitstreams without any noticeable overhead. However, the latter provides more progressive options and produces bitstreams that are parseable and that can be rather easily reorganized by a transcoder on the fly. Along the same lines, JPEG 2000 also provides random access (i.e. involving a minimal decoding) to the block level in each sub-band, thus making possible to decode a region of the image without having to decode it as a whole. These two features could be very advantageous in applications such as digital libraries.

Concerning error resilience JPEG 2000 offers higher protection than JPEG, as shown in the previous section. MPEG-4 VTC also offers error resilience features and although it could not be evaluated the support should be in between JPEG and JPEG 2000. JPEG-LS does not offer any particular support for error resilience, besides restart markers, and has not been designed with it in mind. As for PNG, it offers error detection, but no concealment possibilities.

Overall, one can say that JPEG 2000 offers the richest set of features and provides superior rate-distortion performance. However, this comes at the price of additional complexity when compared to JPEG and JPEG-LS, which might be currently perceived as a disadvantage for some applications, as was the case for JPEG when it was first introduced.

5. CONCLUSIONS

This work aims at providing a comparison of the efficiency of various features that can be expected from a number of recent as well as most popular still image coding algorithms. To do so, many aspects have been considered including genericity of the algorithm to code different types of data in lossless and lossy way, and features

such as error resiliency, complexity, scalability, region of interest, embedded bitstream and so on.

The results show in a quantitative way how much improvement can be expected from various points of view (genericity and other functionalities) from JPEG 2000 standard. At the same time, it puts into the same perspective many existing standards one can efficiently choose from, based on the needs of the underlying product.

6. REFERENCES

- [1] ISO/IEC JTC 1/SC 29/WG 1, *ISO/IEC FCD 15444-1: Information technology — JPEG 2000 image coding system: Core coding system [WG 1 N 1646]*, Mar. 2000, <http://www.jpeg.org/FCD15444-1.htm>.
- [2] William B. Pennebaker and Joan L. Mitchell, *JPEG: Still Image Data Compression Standard*, Van Nostrand Reinhold, New York, 1992.
- [3] ISO/IEC, *ISO/IEC 14496-2:1999: Information technology — Coding of audio-visual objects — Part 2: Visual*, Dec. 1999.
- [4] ISO/IEC, *ISO/IEC 14495-1:1999: Information technology — Lossless and near-lossless compression of continuous-tone still images: Baseline*, Dec. 1999.
- [5] W3C, *PNG (Portable Network Graphics) Specification*, Oct. 1996, <http://www.w3.org/TR/REC-png>.
- [6] ISO/IEC, *ISO/IEC 11544:1993 Information technology — Coded representation of picture and audio information — Progressive bi-level image compression*, Mar. 1993.
- [7] ISO/IEC JTC 1/SC 29/WG 1, *ISO/IEC FCD 14492: Information technology — Coded representation of picture and audio information — Lossy/Lossless coding of bi-level images [WG 1 N 1359]*, July 1999, <http://www.jpeg.org/public/jbigpt2.htm>.
- [8] Diego Santa-Cruz and Touradj Ebrahimi, “A study of JPEG 2000 still image coding versus other standards,” in *Proc. of the X European Signal Processing Conference*, Tampere, Finland, Sept. 2000.