This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th-12th, pp. 359-366, 2000.

JPEG2000: The Upcoming Still Image Compression Standard

A. N. Skodras[1], C. A. Christopoulos[2] and T. Ebrahimi[3]

[1]*Electronics Laboratory, University of Patras, GR-26110 Patras, GREECE*

Email: skodras@cti.gr

[2]*Media Lab, Ericsson Research, Ericsson Radio Systems AB, S-16480 Stockholm, SWEDEN*

Email: charilaos.christopoulos@era.ericsson.se

[3]*Signal Processing Laboratory, EPFL, CH-1015 Lausanne, SWITZERLAND*

Email: Touradj.Ebrahimi@epfl.ch

Abstract: *With the increasing use of multimedia technologies, image compression requires higher performance as well as new features. To address this need in the specific area of still image encoding, a new standard is currently being developed the JPEG2000. It is not only intended to provide rate-distortion and subjective image quality performance superior to existing standards, but also to provide functionality that current standards can either not address efficiently or not address at all. Lossless and lossy compression, encoding of very large images, progressive transmission by pixel accuracy and by resolution, robustness to the presence of bit-errors and region-of-interest coding, are some representative examples of its features. It is interesting to note that JPEG2000 is being designed to address the requirements of a diversity of applications, e.g. Internet, colour facsimile, printing, scanning, digital photography, remote sensing, mobile applications, medical imagery, digital library and E-commerce.*

Keywords*:* *JPEG, colour image coding, data compression, source coding, subbabnd coding, wavelet transform.*

1. INTRODUCTION

Since the mid-80s, members from both the International Telecommunication Union (ITU) and the International Organisation for Standardisation (ISO) have been working together to establish a joint international standard for the compression of continuous-tone (multilevel) still images, both greyscale and colour. This effort has been known as JPEG, the Joint Photographic Experts Group [1, 2]. (The “joint” in JPEG refers to the collaboration between ITU and ISO). Officially, JPEG corresponds to the ISO/IEC international standard 10928-1, digital compression and coding of continuous-tone still images or to the ITU-T Recommendation T.81. The text in both these ISO and ITU-T documents is identical. The process was such that, after evaluating a number of coding schemes, the JPEG members selected a DCT[1]-based method in 1988. From 1988 to 1990, the JPEG group continued its work by simulating, testing and documenting the algorithm. JPEG became a draft international standard (DIS) in 1991 and an international standard (IS) in 1992.

With the continual expansion of multimedia and Internet applications, the needs and requirements of the technologies used, grew and evolved. In March 1997 a new call for contributions were launched for the development of a new standard for the compression of still images, the JPEG2000. This project, JTC[2] 1.29.14 (15444), was intended to create a new image coding system for different types of still images (bi-level, grey-level, colour, multi-component), with different characteristics (natural images, scientific, medical, remote sensing, text, rendered graphics, etc) allowing different imaging models (client/server, real-time transmission, image library archival, limited buffer and bandwidth resources, etc) preferably within a unified system. This coding system should provide low bit-rate operation with rate-distortion and subjective image quality performance superior to existing standards, without sacrificing performance at other points in the rate-distortion spectrum, incorporating at the same time many contemporary features. The standard is intended to compliment and not to replace the current JPEG standards.

The standardisation process, which is co-ordinated by the JTC1/SC29/WG1 of ISO/IEC[3] has already (as of

[1] DCT stands for Discrete Cosine Transform
[2] JTC stands for Joint Technical Committee
[3] SC, WG, IEC stand for Standing Committee, Working Group and International Electrotechnical Commission respectively.

This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th- 12th, pp. 359-366, 2000.

May 2000) produced the Working Draft (WD) [3] and the Committee Draft (CD) [4] documents. DIS is scheduled for July 2000 and the IS for December 2000.

In the present communication the structure of the JPEG2000 standard is presented and performance and complexity comparisons with existing standards, are reported. The paper is organised in the following way: In Section 2 the main areas of application and their requirements are given. The architecture of the standard is described in Section 3, while the multiple-component case is covered in Section 4. The file format aspects, as well as other interesting features of the standard, like region-of-interest coding, error resilience and scalability, are presented in Section 5. Finally, some comparative results are reported in Section 6 of the paper.

2. Applications-Requirements-Features

The JPEG2000 standard provides a set of features that are of vital importance to many high-end and emerging applications, by taking advantage of new technologies. It addresses areas where current standards fail to produce the best quality or performance and provides capabilities to markets that currently do not use compression. The markets and applications better served by the JPEG2000 standard are Internet, colour facsimile, printing, scanning (consumer and pre-press), digital photography, remote sensing, mobile, medical imagery, digital libraries / archives and E-commerce. Each application area imposes some requirements that the standard should fulfil [5]. The features that this standard should possess are the following:

- **Superior low bit-rate performance**: This standard should offer performance superior to the current standards at low bit-rates (e.g. below 0.25 bpp for highly detailed grey-scale images). This significantly improved low bit-rate performance should be achieved without sacrificing performance on the rest of the rate-distortion spectrum. Examples of applications that need this feature include network image transmission and remote sensing. This is the highest priority feature.

- **Continuous-tone and bi-level compression**: It is desired to have a coding standard that is capable of compressing both continuous-tone and bi-level images. If feasible, this standard should strive to achieve this with similar system resources. The system should compress and decompress images with various dynamic ranges (e.g. 1 bit to 16 bit) for each colour component. Examples of applications that can use this feature include compound documents with images and text, medical images with annotation overlays, and graphic and computer generated images with binary and near to binary regions, alpha and transparency planes, and facsimile.

- **Lossless and lossy compression**: It is desired to provide lossless compression naturally in the course of progressive decoding. Examples of applications that can use this feature include medical images, where loss is not always tolerated, image archival applications, where the highest quality is vital for preservation but not necessary for display, network applications that supply devices with different capabilities and resources, and pre-press imagery.

- **Progressive transmission by pixel accuracy and resolution**: Progressive transmission that allows images to be reconstructed with increasing pixel accuracy or spatial resolution is essential for many applications. This feature allows the reconstruction of images with different resolutions and pixel accuracy, as needed or desired, for different target devices. Examples of applications include the World Wide Web, image archival applications and printers.

- **Random codestream access and processing**: Often there are parts of an image that are more important than others. This feature allows user defined Regions-Of-Interest (ROI) in the image to be randomly accessed and/or decompressed with less distortion than the rest of the image. Also, random codestream processing could allow operations such as rotation, translation, filtering, feature extraction and scaling (see also Section 5).

- **Robustness to bit-errors:** It is desirable to consider robustness to bit-errors while designing the codestream. One application where this is important is wireless communication channels. Portions of the codestream may be more important than others in determining decoded image quality. Proper design of the codestream can aid subsequent error correction systems in alleviating catastrophic decoding failures (see also Section 5).

- **Open architecture:** It is desirable to allow open architecture to optimise the system for different image types and applications. With this feature, a decoder is only required to implement the core tool set and a parser that understands the codestream. If necessary, unknown tools are requested by the decoder and sent from the source.

- **Sequential build-up capability (real time coding)**: The standard should be capable of compressing and decompressing images with a single sequential pass. This standard should also be capable of processing an image using component interleave order or non-interleaved order. During compression and decompression, the standard should use context limited to a reasonable number of lines.

3. Architecture of the Standard

The block diagram of the JPEG200 encoder is illustrated in Fig. 1a. It is similar to every transform-based coding scheme. The discrete transform is first

This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th- 12th, pp. 359-366, 2000.

applied on the source image data. The transform coefficients are then quantised and entropy coded, before forming the output codestream (bitstream). The decoder is the reverse of the encoder (Fig. 1b). The codestream is first entropy decoded, dequantised and inverse discrete transformed, thus resulting in the reconstructed image data. It is worth mentioning that, unlike other coding schemes, the JPEG2000 can be both lossy and lossless. This depends on the wavelet transform and the quantisation applied.

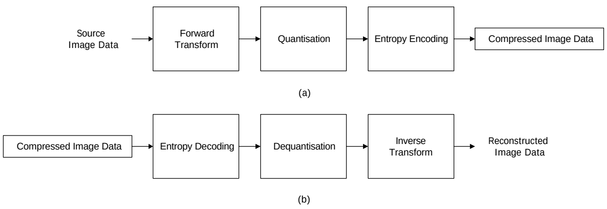

Figure 1. Block diagrams of the JPEG2000 (a) encoder and (b) decoder.

Before proceeding with the details of each block of Fig. 1, it should be mentioned that the standard works on image tiles. The term ‘tiling’ refers to the partition of the original (source) image into rectangular non-overlapping blocks (tiles), which are compressed independently, as though they were entirely distinct images (Fig. 2). This is the strongest form of spatial partitioning, in that all operations, including component mixing (see also Section 4), wavelet transform, quantisation and entropy coding are performed independently on the different tiles of the image. All tiles have exactly the same dimensions, except maybe those, which abut the right and lower boundary of the image. The nominal tile dimensions are exact powers of two. Tiling reduces memory requirements and constitutes one of the methods for the efficient extraction of a region of the image. Prior to computation of the forward discrete wavelet transform (DWT) on each tile, all samples of the image tile component are DC level shifted by subtracting the same quantity (i.e. the component depth) from each sample (Fig. 2).

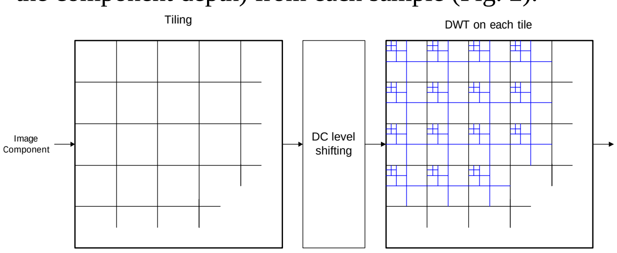

Figure 2. Tiling, DC level shifting and DWT of each image tile component.

3.1. The Wavelet Transform

The tile components are decomposed into different decomposition levels using a wavelet transform. These decomposition levels contain a number of subbands populated with coefficients that describe the horizontal and vertical spatial frequency characteristics of the original tile component planes (Fig. 2). The coefficients provide local frequency information. A decomposition level is related to the next decomposition level by spatial powers of two. To perform the forward DWT the standard uses a 1-D subband decomposition of a 1-D set of samples into low-pass samples, representing a downsampled low-resolution version of the original set, and high-pass samples, representing a downsampled residual version of the original set, needed for the perfect reconstruction of the original set from the low-pass set. Any user supplied wavelet filter banks may be used. The DWT can be *irreversible* or *reversible*. The default *irreversible transform* is implemented by means of the Daubechies 9-tap/7-tap filter [6]. The analysis and the corresponding synthesis filter coefficients are given in Table I. The default *reversible transformation* is implemented by means of the 5-tap/3-tap filter, the coefficients of which are given in Table II [7].

Table I Daubechies 9/7 analysis and synthesis filter coefficients

Analysis Filter Coefficients		
i	Lowpass Filter $h_L(i)$	Highpass Filter $h_H(i)$
0	0.6029490182363579	1.115087052456994
±1	0.2668641184428723	-0.5912717631142470
±2	-0.07822326652898785	-0.05754352622849957
±3	-0.01686411844287495	0.09127176311424948
±4	0.02674875741080976	

Synthesis Filter Coefficients		
i	Lowpass Filter $g_L(i)$	Highpass Filter $g_H(i)$
0	1.115087052456994	0.6029490182363579
±1	0.5912717631142470	-0.2668641184428723
±2	-0.05754352622849957	-0.07822326652898785
±3	-0.09127176311424948	0.01686411844287495
±4		0.02674875741080976

Table II 5/3 analysis and synthesis filter coefficients

Analysis Filter Coefficients		
i	Lowpass Filter $h_L(i)$	Highpass Filter $h_H(i)$
0	6/8	1
±1	2/8	-1/2
±2	-1/8	

Synthesis Filter Coefficients		
i	Lowpass Filter $g_L(i)$	Highpass Filter $g_H(i)$
0	1	6/8
±1	1/2	-2/8
±2		-1/8

The standard supports two filtering modes: a *convolution-based* and a *lifting-based*. For both modes to be implemented, the signal should be first extended periodically as shown in Fig. 3. This *periodic symmetric extension* is used to ensure that for the filtering operations that take place at both boundaries of the signal, one signal sample exists and spatially

This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th- 12th, pp. 359-366, 2000.

corresponds to each coefficient of the filter mask. The number of additional samples required at the boundaries of the signal is therefore filter-length dependent [8].

Figure 3. Periodic symmetric extension of the signal ABCDEFG. (It has been assumed that A>B>…>G for illustration purposes).

Convolution-based filtering consists in performing a series of dot products between the two filter masks and the extended 1-D signal. *Lifting-based filtering* consists of a sequence of very simple filtering operations for which alternately odd sample values of the signal are updated with a weighted sum of even sample values, and even sample values are updated with a weighted sum of odd sample values [4, 9]. For the reversible (lossless) case the results are rounded to integer values. The lifting-based filtering for the 5/3 analysis filter is

achieved by means of eq. (1) below:

$$y(2n+1) = x_{ext}(2n+1) - \left\lceil \frac{x_{ext}(2n) + x_{ext}(2n+2) - 1}{2} \right\rceil \quad (1a)$$

$$y(2n) = x_{ext}(2n) + \left\lfloor \frac{y(2n-1) + y(2n+1) + 2}{4} \right\rfloor \quad (1b)$$

where x_{ext} is the extended input signal, y is the output signal and $\lfloor a \rfloor$, $\lceil a \rceil$ indicate the largest integer not exceeding *a* and the smallest integer not exceeded by *a*, respectively.

3.2. Quantisation

Quantisation is the process by which the coefficients are reduced in precision. This operation is lossy, unless the quantisation step is 1 and the coefficients are integers, as produced by the reversible integer 5/3 wavelet. Each of the transform coefficients $a_b(u,v)$ of the subband b is quantised to the value $q_b(u,v)$ according to the formula:

$$q_b(u,v) = \text{sign}(a_b(u,v)) \left\lfloor \frac{|a_b(u,v)|}{\Delta_b} \right\rfloor \quad (2)$$

The quantisation step $_b$ is represented relative to the dynamic range R_b of subband b, by the exponent $_b$ and mantissa $_b$ as:

$$\Delta_b = 2^{R_b - \varepsilon_b}\left(1 + \frac{\mu_b}{2^{11}}\right) \quad (3)$$

The dynamic range R_b depends on the number of bits used to represent the original image tile component and on the choice of the wavelet transform. All quantised transform coefficients are signed values even when the original components are unsigned. These coefficients are expressed in a sign-magnitude representation prior to coding.

3.3. Entropy Coding

Each subband of the wavelet decomposition is divided into rectangular blocks, called *code-blocks*, which are coded independently using arithmetic coding. These code-blocks are coded at a bit-plane at a time, starting with the most significant bit-plane with a non-zero element to the least significant bit-plane. For each bit-plane in a code-block, a special code-block scan pattern is used for each of three passes. Each coefficient bit in the bit-plane is coded in only one of the three passes. A rate distortion optimisation method is used to allocate a certain number of bits to each block. The recursive probability interval subdivision of Elias coding is the basis for the binary *arithmetic coding* process. With each binary decision, the current probability interval is subdivided into two sub-intervals, and the codestream is modified (if necessary) so that points to the base (the lower bound) of the probability sub-interval assigned to the symbol, which occurred. Since the coding process involves addition of binary fractions rather than concatenation of integer codewords, the more probable binary decisions can often be coded at a cost of much less than one bit per decision [4].

4. MULTIPLE-COMPONENT IMAGES

JPEG2000 supports multiple-component images. Different components need not have the same bit-depths; nor need they have all been signed or unsigned. For reversible systems, the only requirement is that the bit-depth of each output image component must be identical to the bit-depth of the corresponding input image component.

The standard supports two different component transformations, one *irreversible component transformation* (ICT) and one *reversible component transformation* (RCT). It is usual that the input image has three components, the Red, Green and Blue (RGB). Thus in the following we will be using these components without, however, restricting the generality of the presentation. The block diagram of the JPEG2000 multiple-component encoder is shown in Fig. 4. C_1, C_2, C_3 represent in general the colour transformed output components. It is worth mentioning that prior to applying the forward colour transformation, the image component samples are DC level shifted (if needed).

This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th- 12th, pp. 359-366, 2000.

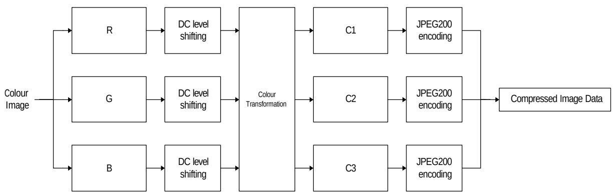

Figure 4. Structure of the JPEG2000 multiple-component encoder.

The ICT may only be used for lossy coding. It can be seen as an approximation of a YC_bC_r transformation of the RGB components. The forward and the inverse ICT transformations are achieved by means of eq. (4).

$$\begin{pmatrix} Y \\ C_b \\ C_r \end{pmatrix} = \begin{pmatrix} 0.299 & 0.587 & 0.114 \\ -0.16875 & -0.33126 & 0.5 \\ 0.5 & -0.41869 & -0.08131 \end{pmatrix} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix} \quad (4a)$$

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} 1.0 & 0 & 1.402 \\ 1.0 & -0.34413 & -0.71414 \\ 1.0 & 1.772 & 0 \end{pmatrix} \cdot \begin{pmatrix} Y \\ C_b \\ C_r \end{pmatrix} \quad (4b)$$

The RCT may be used for lossy or lossless coding. It is a decorrelating transformation, which is applied to the three first components of an image. Three goals are achieved by this transformation, namely, colour decorrelation for efficient compression, reasonable colour space with respect to the Human Visual System for quantisation, and ability of having lossless compression, i.e. exact reconstruction with finite integer precision. For the RGB components, the RCT can be seen as an approximation of a YUV transformation. The forward and inverse RCT is performed by means of eq. (5):

$$\begin{pmatrix} Y_r \\ U_r \\ V_r \end{pmatrix} = \begin{pmatrix} \left\lfloor \frac{R+2G+B}{4} \right\rfloor \\ R-G \\ B-G \end{pmatrix} \quad (5a)$$

$$\begin{pmatrix} G \\ R \\ B \end{pmatrix} = \begin{pmatrix} Y_r - \left\lfloor \frac{U_r+V_r}{4} \right\rfloor \\ U_r + G \\ V_r + G \end{pmatrix} \quad (5b)$$

5. SIGNIFICANT FEATURES OF THE STANDARD

The JPEG2000 standard exhibits a lot of nice features, the most significant being the possibility to define regions of interest in an image, the spatial and SNR scalability, the error resilience and the possibility of intellectual property rights protection. Interestingly enough, all these features are incorporated within a unified algorithm.

Region-of-Interest (ROI): One of the features included in JPEG2000 is the ROI coding. According to this, certain ROI's of the image can be coded with better quality than the rest of the image (background). The ROI scaling-based method used, scales up (DC shifts) the coefficients so that the bits associated with the ROI are placed in higher bit-planes. During the embedded coding process, those bits are placed in the bit-stream before the non-ROI parts of the image. Thus, the ROI will be decoded, or refined, before the rest of the image. Regardless of the scaling, a full decoding of the bit-stream results in a reconstruction of the whole image with the highest fidelity available. If the bit-stream is truncated, or the encoding process is terminated before the whole image is fully encoded, the ROI will have a higher fidelity than the rest of the image [10]. The ROI approach defined in the JPEG2000 Part I is called MAXSHIFT method and allows ROI encoding of arbitrary shaped regions without the need of shape information and shape decoding [10-12].

Scalability: Realising that many applications require images to be simultaneously available for decoding at a variety of resolutions or qualities, this architecture supports scalability. In general, scalable coding of still images means the ability to achieve coding of more than one resolution and/or quality simultaneously. Scalable image coding involves generating a coded representation (bitstream) in a manner which facilitates the derivation of images of more than one resolution and/or quality by scalable decoding. Bitstream scalability is the property of a bitstream that allows decoding of appropriate subsets of a bitstream to generate complete pictures of resolution and/or quality commensurate with the proportion of the bitstream decoded. If a bitstream is truly scalable, decoders of different complexities, from low performance decoders to high performance decoders can coexist. While low performance decoders may decode only small portions of the bitstream producing basic quality, high performance decoders may decode much more and produce significantly higher quality. The most important types of scalability are SNR[4] scalability and spatial scalability [3, 4].

SNR scalability is intended for use in systems with the primary common feature that a minimum of two layers of image quality is necessary. SNR scalability involves generating at least two image layers of same spatial resolution, but different qualities, from a single image source. The lower layer is coded by itself to provide the basic image quality and the enhancement layers are coded to enhance the lower layer. The enhancement layer when added back to the lower layer regenerates a higher quality reproduction of the input image.

[4] SNR stands for Signal to Noise Ratio

This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th- 12th, pp. 359-366, 2000.

Spatial scalability is intended for use in systems with the primary common feature that a minimum of two layers of spatial resolution is necessary. Spatial scalability involves generating at least two spatial resolution layers from a single source such that the lower layer is coded by itself to provide the basic spatial resolution and the enhancement layer employs the spatially interpolated lower layer and carries the full spatial resolution of the input image source.

An additional advantage of spatial and SNR scalability types is their ability to provide resilience to transmission errors, as the most important data of the lower layer can be sent over channel with better error performance, while the less critical enhancement layer data can be sent over a channel with poor error performance.

Both types of scalability are very important for Internet and database access applications, and bandwidth scaling for robust delivery. The SNR and spatial scalability types include the progressive and hierarchical coding modes already defined in the current JPEG, but they are more general.

Error Resilience: Many applications require the delivery of image data over different types of communication channels. Typical wireless communications channels give rise to random and burst bit errors. Internet communications are prone to loss due to traffic congestion. To improve the performance of transmitting compressed images over these error prone channels, error resilient bit stream syntax and tools are included in this standard. The error resilience tools deal with channel errors using the following approaches: data partitioning and resynchronisation, error detection and concealment, and Quality of Service (QoS) transmission based on priority [3, 4].

New File Format with IPR Capabilities: An optional file format (JP2) for the JPEG2000 compressed image data has been defined by the standard. This format has got provisions for both image and metadata, a mechanism to indicate the tonescale or colourspace of the image, a mechanism by which readers may recognise the existence of intellectual property rights (IPR) information in the file and a mechanism by which metadata (including vendor specific information) can be included in the file [4].

6. COMPARATIVE RESULTS

Up to now we have been dealing with the description of the main parts of the JPEG2000 algorithm. In this section we will look at the efficiency of the algorithm in comparison with existing lossless and lossy compression standards.

The rate-distortion behaviour of the lossy (non-reversible) JPEG2000 and the progressive JPEG is depicted in Fig. 5 for a natural image. It is seen that the JPEG2000 significantly outperforms the JPEG scheme. We can easily conclude that for similar PSNR quality, the JPEG2000 compresses almost twice more than JPEG [13-16]. The superiority of the JPEG2000 can be subjectively judged with the help of Fig. 6, where the reconstructed image '*hotel*' (720x576) is shown. Both images were compressed at a rate of 0.125 bpp using the existing JPEG and the upcoming JPEG2000. The degradation of the image in Fig 6a is evident [14, 17].

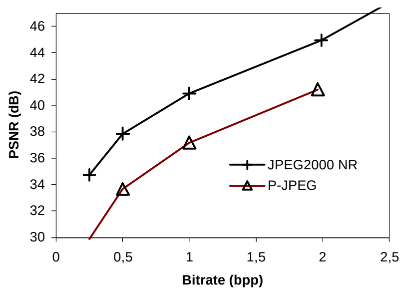

Figure 5. Rate-distortion results for the progressive JPEG2000 versus the progressive JPEG for a natural image.

(a)

(b)

Figure 6. Reconstructed images compressed at 0.125 bpp by means of (a) JPEG and (b) JPEG2000.

This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th- 12th, pp. 359-366, 2000.

One of the interesting and unique features of JPEG2000 is its capability in defining ROI's, which are coded at a better quality than the rest of the image. The regions can be one or more and of any shape and size. In Fig. 7 an example of a circular ROI is given. Experiments have shown that for lossless coding of images, the ROI feature results in an increase in the bitrate of approximately 1 to 8 percent in comparison to lossless coding without using the ROI feature [18].

The lossless compression efficiency of JPEG2000 versus the lossless mode of JPEG [1, 2] and JPEG-LS [19] for a natural and a compound image is reported in Table III. It is seen that JPEG2000 performs equivalently to JPEG-LS in the case of the natural image, with the added benefit of scalability. JPEG-LS, however, is advantageous in the case of the compound image. Taking into account that JPEG-LS is significantly less complex than JPEG2000, it is reasonable to use JPEG-LS for lossless compression. In such a case though, the generality of JPEG2000 will be sacrificed [15].

Figure 7. Reconstructed image in which a ROI of circular shape have been defined.

Table III Lossless compression results (in bpp)

Image	lossless JPEG	JPEG-LS	JPEG2000
Lena (512x512, 24bpp)	14.75	13.56	13.54
Cmpnd1 (512x768, 8bpp)	2.48	1.24	2.12

The comparative results of JPEG, JPEG-LS and JPEG2000 from the functionality point of view are illustrated in Table IV. A plus (minus) sign indicates that the corresponding functionality is supported (not supported). The more the plus signs the greater the support. The parentheses indicate that a separate mode is required. It becomes evident from Table IV, that the JPEG2000 standard offers the richest set of features in a very efficient way and within a unified algorithm [15].

Table IV Summary of functionalities supported by each standard

Compression Algorithm	Lossless	Lossy	Embedded bitstream	ROI	Error resilience	Scalability	Complexity	Random access	Generic
JPEG	(+)	++	-	-	-	(+)	+(+)	+	+
JPEG-LS	++++	+	+	-	-	-	+	-	+
JPEG2000	+++	+++	+++	++	++	++	+++	++	+++

However, all of the above mentioned advantages of JPEG2000 are at the expense of memory and data accesses complexity and execution time. It has been reported in [20] an increase in the memory size by a factor of 40 and an increase in memory accesses by a factor of 1.9 of JPEG2000 versus JPEG. For the decoder, this is mainly due to the wavelet transform, which takes up to 40% of the resources as compared to less than 25% needed by the IDCT for a 1.0 bpp compression rate. As for the execution times, simulation results have shown that the JPEG2000 encoder takes approximately 34 times longer than the JPEG encoder, while the JPEG2000 decoder takes about 8 times longer than the JPEG decoder [21]. It should be stressed though, that these figures refer to non-optimal implementations, which are also platform dependent. Careful optimisation of the algorithm will greatly improve the performance without sacrificing functionality.

7. CONCLUSIONS

JPEG2000 is the new standard for still image compression that is going to be in use by the beginning of next year. It provides a new framework and an integrated toolbox to better address increasing needs for compression and functionalities for still image applications, like internet, colour facsimile, printing, scanning, digital photography, remote sensing, mobile applications, medical imagery, digital library and E-commerce. Lossless and lossy coding, embedded lossy to lossless, progressive by resolution and quality, high compression efficiency, error resilience and lossless colour transformations are some of its characteristics. Comparative results have shown that JPEG2000 is indeed superior to existing still image compression standards. Work is still needed in optimising its implementation performance. The reference software is

This paper appeared in:
Proceedings of the 11th Portuguese Conference on Pattern Recognition (RECPA00D 20; invited paper), Porto, Portugal, May 11th- 12th, pp. 359-366, 2000.

decided to be developed not only in C, but in JAVA as well. The intention is to have a license fee free software for commercial and non-commercial use. The JAVA version can be downloaded from http://jj2000.epfl.ch.

ACKNOWLEDGEMENTS

The authors would like to thank Diego Santa Cruz from EPFL and Mathias Larsson and Joel Askelof from Ericsson, Sweden for their contribution in the simulations.

REFERENCES

[1] G. K. Wallace: "The JPEG Still Picture Compression Standard", IEEE Trans. Consumer Electronics, Vol. 38, No 1, Feb. 1992.

[2] V. Bhaskaran and K. Konstantinides: "Image and Video Compression Standards: Algorithms and Applications", 2nd Ed., Kluwer Academic Publishers, 1997.

[3] ISO/IEC JTC1/SC29/WG1 N1577: "JPEG2000 Part II Working Draft Version 1.0 Pre-Release A", Jan. 26, 2000.

[4] ISO/IEC CD 15444-1: "JPEG2000 Committee Draft v1.0", Dec. 9, 1999.

[5] ISO/IEC JTC1/SC29/WG1 N1271: "JPEG2000 Requirements and Profiles v5.0", March 1999.

[6] M. Antonini, M. Barlaud, P. Mathieu and I. Daubechies: "Image Coding Using the Wavelet Transform", IEEE Trans. Image Proc., pp. 205-220, April 1992.

[7] D. Le Gall and A. Tabatabai: "Subband Coding of Digital Images Using Symmetric Short Kernel Filters and Arithmetic Coding Techniques", Proc. ICASSP, NY, pp. 761-765, 1988.

[8] ISO/IEC JTC1/SC29/WG1 N1595: "USNB Comments on JPEG2000 CD 1.0", Jan. 3, 2000.

[9] A. R. Calderbank, I. Daubechies, W. Sweldens and B.-L. Yeo: "Lossless Image Compression Using Integer to Integer Wavelet Transforms", Proc. ICIP-97, Santa Barbara, USA, Oct. 1997.

[10] D. Nister and C. Christopoulos: "Lossless Region of Interest with Embedded Wavelet Image Coding", Signal Processing, Vol. 78, No 1, pp. 1-17, 1999.

[11] C. Christopoulos**,** J. Askelof and M. Larsson, "Efficient Region of Interest Encoding Techniques in the Upcoming JPEG2000 Still Image Coding Standard", invited paper to IEEE Int. Conf. on Image Processing (ICIP-2000), Special Session on JPEG2000, Vancouver, Canada, Sept. 2000.

[12] C. Christopoulos**,** J. Askelof and M. Larsson, "Efficient Encoding and Reconstruction of Regions of Interest in JPEG2000", to be presented at EUSIPCO-2000, Tampere, Finland, Sep. 2000.

[13] M. Charrier, D. Santa Cruz and M. Larsson: "JPEG2000: The Nest Millennium Compression Standard for Still Images", Proc. ICMCS-99, Florence, Italy, June 1999.

[14] C. A. Christopoulos and A. N. Skodras: "The JPEG 2000 (Tutorial)", Proc. IEEE Int. Conf. Image Processing, ICIP-99, Kobe, Japan, 25-28 Oct 1999.

[15] D. Santa Cruz and T. Ebrahimi: "An Analytical Study of the JPEG2000 Functionalities", to be presented at ICIP-2000, Vancouver, Canada, Sep. 2000.

[16] ISO/IEC JTC1/SC29/WG1 N1583: "Visual Evaluation of the JPEG2000 Color Image Compression Performance", March 2000.

[17] http://etro.vub.ac.be/~chchrist/jpeg2000_contributions.htm

[18] C. A. Christopoulos, J. Askelof and M. Larsson: "Efficient Region of Interest Coding Techniques in the Upcoming JPEG2000 Still Image Coding Standard", to be presented at ICIP-2000, Vancouver, Canada, Sep. 2000.

[19] ISO/IEC JTC1/SC29/WG1 N575: "JPEG-LS (14495) Final CD", July 1997.

[20] L. Nachtergaele, S. Wuytack, B. Vanhoof and J. Bormans: "Study of the VM5.2 Memory Complexity for Cost Efficient Implementation", ISO/IEC JTC1/SC29/WG1 N1469R, Dec. 2, 1999.

[21] D. Pian: "Study of the VM Complexity for Cost Efficient Implementation", ISO/IEC JTC1/SC29/WG1 N1471R, Dec. 3, 1999.