

# JPEG 2000 Image Compression

By Christine Bako [christine.bako@analog.com]

## INTRODUCTION

The JPEG (Joint Photographic Experts Group) 2000 standard, finalized in 2001, defines a new image-coding scheme using state-of-the-art compression techniques based on wavelet technology. Its architecture is useful for many diverse applications, including Internet image distribution, security systems, digital photography, and medical imaging.

A lot of confusion exists as to what JPEG 2000 is and how it compares with other compression standards such as MPEG (Moving-Picture Experts Group) -2, MPEG-4, and the earlier JPEG. With brief comparisons to other compression standards, this article is primarily intended to highlight some of the often misunderstood and rarely mentioned potential-become-actual benefits of JPEG 2000.

## Applications

### CCTV Security

When transmitting or storing picture information, *compression* must be employed to maintain picture resolution while making best use of limited channel bandwidth. Compression is defined as *lossless* if full recovery of the original is available from the channel without any loss of information; otherwise, it is *lossy*. Standards are required to ensure interoperability. JPEG 2000 is the only standard compression scheme that provides for both lossless and lossy compression. As such, it lends itself to applications that require high-quality images despite limitations on storage or transmission bandwidths.

An important feature of systems based on JPEG 2000 is the ability to extract a variety of resolutions, components, areas of interest, and compression ratios from a single JPEG 2000 code stream. This is not possible with any other compression standard because the image size, bit rate, and quality must be specified on the encode side and can not be determined or changed on the decode side.

For example, a closed-circuit TV (CCTV) security system can make use of this feature by sending a single JPEG 2000 code stream over a low bandwidth network. High-resolution images can be stored on a hard-disk drive (HDD), while several lower-resolution images are displayed on monitors. The operator on the *receive* side can decide what information to extract from the single code stream sent.

JPEG 2000 is *frame accurate*, in that every single frame of the input is contained in the compressed format. MPEG systems, on the other hand, reduce the amount of data through *temporal compression* (which does not encode each frame as a complete image), so MPEG compression is not *frame-accurate*. For this reason, legal issues restrict the use of MPEG compression in some security applications. To get around this problem, security system and equipment providers have had to develop their own compression schemes—or use the highly inefficient *motion JPEG (M-JPEG)* compression standard—in order to provide a compressed stream that contains every single field of the original. They can now use JPEG 2000 for new designs.

### Internet Image Distribution

*Progressive coding*, another feature of the JPEG 2000 standard, means that the bit stream can be coded in such a way as to contain less-detailed information at the beginning of the stream and more detailed information as the stream progresses. This makes it ideal for Internet/network applications—especially with large images and low bandwidths—as the image can be seen instantly on the decoding side, even with low-speed networks or image databases. The lower subbands are shown first, and more detail is added as time progresses. The picture thus becomes sharper and more detailed over time, and the entire image does not have to be downloaded before it can be seen.

With the low-quality image instantly available, the user at the receiving end can decide whether to view the picture in its fully decoded version, or to pass it by and scan the next picture instead. Clients can view images at different resolutions or quality levels [*compression rates*] making them suitable for any transmission bandwidth, connection speed, or display device. In addition, JPEG 2000 coding provides the option to zoom in or out on a particular area of the image—or to display a particular region of the image at a different resolution or compression rate.

### High Definition

At extreme compression levels, JPEG 2000 video starts to blur, but is still quite viewable. MPEG or JPEG artifacts are much more disturbing to the eye, with the picture visibly broken down into small blocks at high compression ratios. The high image quality at medium-to-high bit rates and contents that contain a lot of motion, lack of block artifacts, and high efficiency make JPEG 2000 ideal for high-definition (HD) applications, such as digital cinema, HD recording systems, and HD camera equipment.

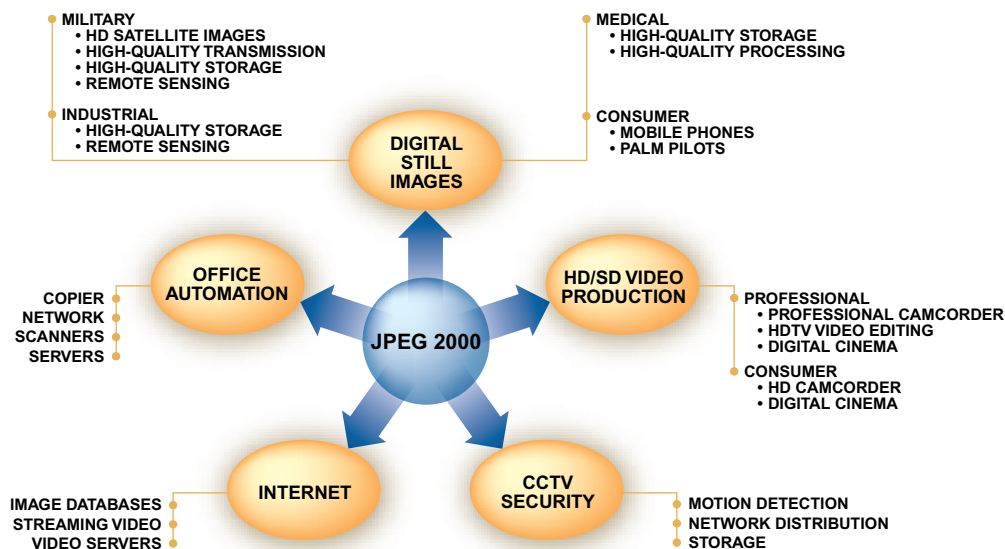


Figure 1. JPEG 2000 applications.

Many applications require exact bit-rate control, which only JPEG 2000 can provide. Exact bit-rate control is possible because an entire frame or field is transformed at once; it is then broken down into bit streams or code blocks that can be processed independently with the techniques described below. In systems using DCT, quantization is the only technique used, and this makes exact bit-rate control difficult. In order to control bit rate in DCT systems, the information must be repeatedly re-processed and re-quantized. The rate-control algorithm used in JPEG 2000 truncates each bit stream to meet a specific target bit rate, adjusting the truncation and re-quantization of each code block's data as required. In addition to programming the target bit rate, the standard allows the user to specify a particular quality metric. In this case, the target bit rate will vary to meet the specified quality factor, as long as the performance does not fall below a specific *peak signal-to-noise ratio*. The PSNR is an indication of picture quality comparable to *perceived picture quality*.

### JPEG 2000 Code Stream

A given input image or part of the image [*tile*] is sent to a set of wavelet filters, which transform the pixel information into wavelet coefficients, which are then grouped into several *subbands* [the use of *wavelets* in encoding was first explained in Analog Dialogue 30-2 (1996)]. Each subband contains wavelet

coefficients that describe a specific horizontal and vertical spatial frequency range of the entire original image. This means that lower-frequency, less-detailed information is contained in the first transform level, while more-detailed, higher-frequency information is contained in higher transform levels. For simplicity, only two levels of transform are shown here. The first transform level results in subbands LH1, HH1, HL1, and LL1. Only subband LL1 is passed on for further filtering, generating the next transform level and creating subbands LH2, HH2, HL2, and LL2.

Equally sized code blocks, which are essentially bit streams of data, are generated within each subband. This break-down is necessary for coefficient modeling and coding, and is done on a code-block-by-code-block basis. In essence, the actual compression is achieved by truncating and/or re-quantizing the bit streams contained in each code block. These bit streams are then optimally truncated using a technique known as *post-compression-rate-control* (PCRC).

Code blocks can be accessed independently. Their bit streams are coded with three coding passes per bit plane. This process, called context modeling, is used to assign information about the importance of each individual coefficient bit. The code blocks can then be grouped according to their significance. On the decoding side it is then possible to extract information according to its significance, allowing the most significant information to be seen first.

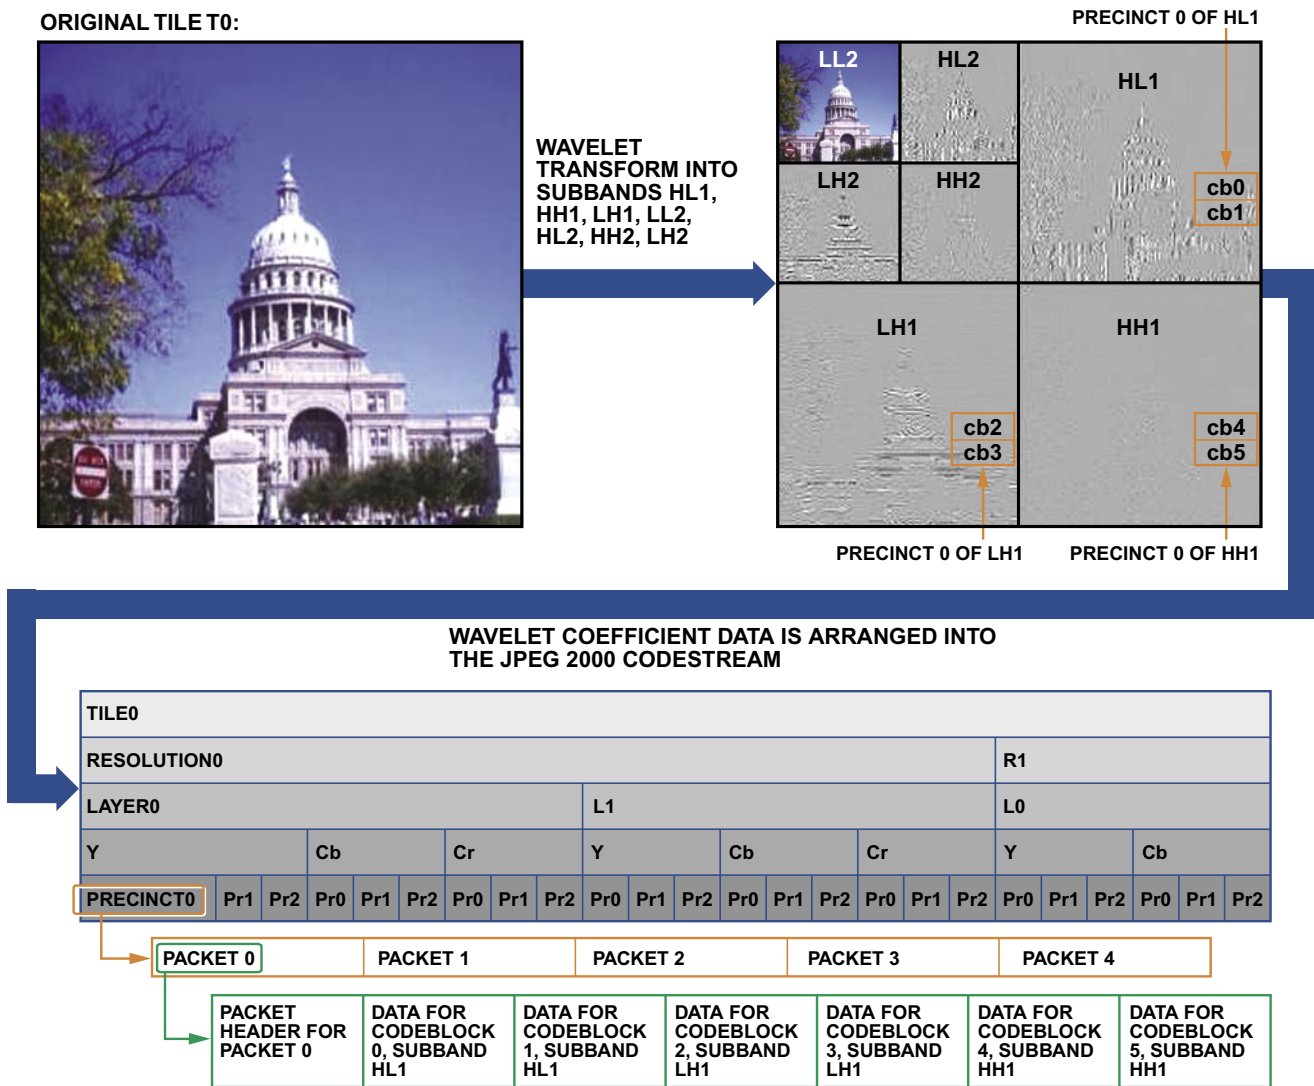


Figure 2. ENCODE—image over wavelet transform into subbands and resolutions.

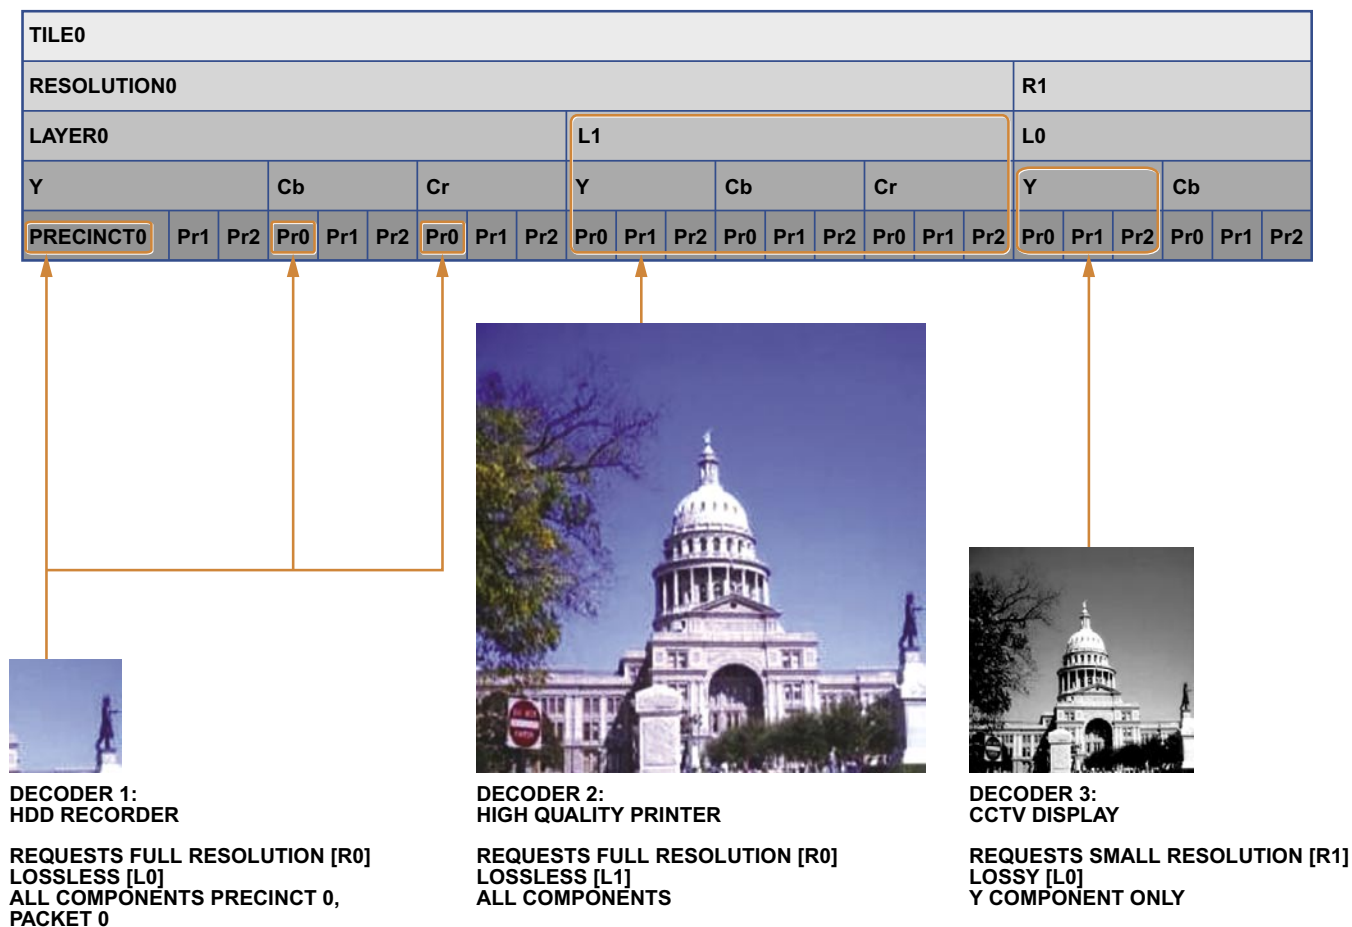


Figure 3. DECODE—one JPEG 2000 stream is received by several decoders.

JPEG 2000 can contain a user-defined number of layers, which are defined by PCRC and context modeling. Each layer stands for a particular compression rate, where the compression rate is achieved from the quantization-, rate-distortion-, and context modeling processes. Layer 0, for example, contains bit streams from the lossy WT transform—that are heavily truncated, contain no coding passes, and thus provide the highest compression rate and the lowest quality. Layer 16 can then contain bit streams that are less truncated and use a higher number of coding passes, thus providing low compression and high quality.

Tiles or images are further partitioned into *precincts*. Precincts contain a number of code blocks, and are used to facilitate access to a specific area within an image in order to process this area in a different way, or to decode only a specific area of an image. The JPEG 2000 bit stream is generated by arranging code blocks or precincts into an array of packets with the lower subbands coming first.

The JPEG 2000 stream starts with a main header containing information such as: uncompressed image size, tile size, number of components, bit depth of components, coding style, transform levels, progression order, number of layers, code block size, wavelet filter type, quantization level, etc. The entire image data, grouped in code blocks of LL, HL, LH, and HH subbands, follows the header. Data is *not* contained in the

header information. Also, a *table of contents* can be stored on the encode side, and allows a decoder to call up a certain resolution on demand, without first having to decode or download the entire JPEG 2000 code stream.

#### DCT versus WT

JPEG 2000 uses the *wavelet transform* (WT) to reduce the amount of information contained in a picture, while MPEG and JPEG systems use the *discrete cosine transform* (DCT). It is true that the WT requires more processing power than the DCT, but MPEG systems require more than just the DCT. The DCT, or any type of *Fourier* transform, expresses the signal in terms of frequency and amplitude—but only at a single instant in time. The WT transforms a signal into frequency and amplitude over time, and is therefore more efficient. The figures on the following page demonstrate this.

To obtain the same amount of information as with one WT pass, the DCT must be used for every frequency; and each of these frequencies must be transformed at each time instant for each  $8 \times 8$  pixel block. In addition, MPEG systems use inter-frame compression [motion estimation] in order to reduce the amount of data further for motion estimation. This requires storage of at least two entire fields in external memory. The computation-intensive motion estimation process requires a very powerful processor. Temporal compression can be used in JPEG 2000 systems, but it is not inherent in the JPEG 2000 standard.




Figure 4. Input signal 1 containing frequencies A, B, C, and D.




Figure 5. Input signal 2 containing frequencies A, B, C, and D.




Figure 6. Wavelet transform of signal 1.




Figure 7. Wavelet transform of signal 2.

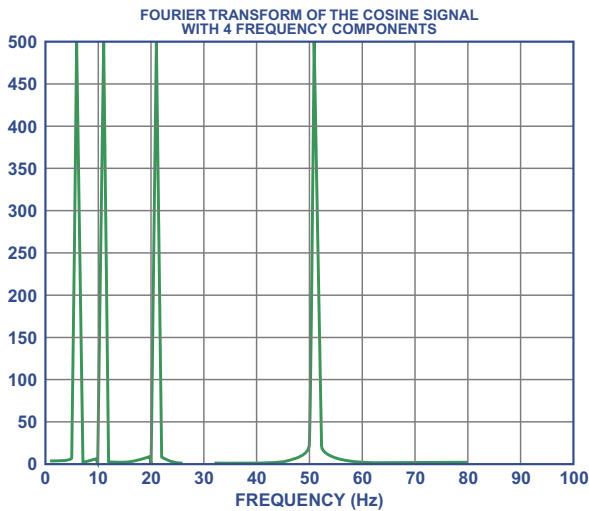


Figure 8. Fourier transform of signal 1.

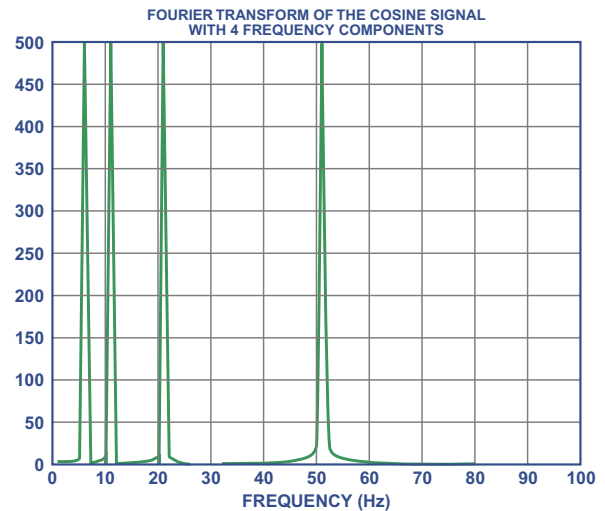


Figure 9. Fourier transform of signal 2.




Figure 10. ADV202 block diagram.

### JPEG 2000's Advantages Over Other Compression Standards

All MPEG standards are complex and computation intensive. This translates into extensive processing latency and memory requirements in *standard-definition* (SD) applications. These factors become even more of a problem when *high-definition* (HD) formats are considered, and JPEG 2000 becomes even more desirable. Another strength of JPEG 2000 is the standard itself, which allows immense flexibility and control in many different applications. There is also much versatility regarding formats: JPEG 2000 supports anything from 8-bits per sample to an unlimited amount of bits per sample, whereas MPEG only supports 8-bit data.

JPEG 2000 continues to gain popularity, even though MPEG-2 is the established standard for DVD and broadcast applications. JPEG 2000 is also very popular in HD applications that require high-quality storage or transmission of HD images over wireless or other links.

### The ADV202

Since the early 1990s, Analog Devices has invested heavily in wavelet-compression R&D. We were the first to introduce a wavelet-compression hardware solution in 1996 with the ADV601. Now ADI's newest wavelet codec, the ADV202, released in July 2004, is so far the only dedicated JPEG 2000 IC on the market. A complete single-chip JPEG 2000 compression/decompression IC, the ADV202 works with high-definition video, standard-definition video, and still images. It supports all features of the ISO/IEC15444-1 [JPEG 2000] image-compression standard [except Maxshift ROI]. Its patented SURF™ (spatial ultra-efficient recursive filtering) technology enables low-power, low-cost wavelet-based compression. Containing a dedicated wavelet transform engine, three entropy coders, a RISC processor, and on-board memory systems, the ADV202

provides a glueless interface to common video standards such as ITU.R.BT656, SMPTE274M, or SMPTE296M. It can create a fully compliant JPEG 2000 code stream [.j2c, .jp2]. It can also provide raw code-block and attribute data, allowing the host processor to have complete control over the generation- and compression processes.

Even though digital signal-processor (DSP) performance has improved significantly, a DSP would have to perform 20 billion instructions per second to match the performance of the ADV202 in a standard-definition encode application. Effectively serving as accelerators, the ADV202's three dedicated on-chip entropy coders are responsible for the high throughput rate.

### CONCLUSION—THE OUTLOOK FOR JPEG 2000

A major advantage of using a JPEG 2000 hardware solution is lower latency than any other compression scheme, a factor which is of particular importance in medical applications, for example.

Several major manufacturers of video or broadcast equipment have implemented JPEG 2000 into such future HD products as real-time encoding and decoding systems and video servers.

The Digital Cinema Initiative (DCI) has recently announced that it will use JPEG 2000 as the compression method in the delivery of digital motion pictures. The ADV202 has already found its way into many designs in the CCTV/security market in *video-over-network* applications. For more information, please visit <http://www.itscj.ipsj.or.jp/sc29/29w02901.pdf>

Because of its flexibility and image-compression quality, the ADV202—operating under JPEG 2000—could find its way into virtually every design that uses image or video compression. ▀