

WHITEPAPER

JPEG 2000 in the HD Broadcast Production Workflow

John R. Naylor and Dave Bancroft

February 2010

JPEG 2000 is an ideal acquisition codec for serious HD productions, because it allows the full picture to be captured in full-color at economical bit rates.

Contents

Why is Compression Needed?	1
Capture	2
Edit	2
Play-out	2
Emission	2
Acquisition Codec Requirements	3
JPEG 2000 Compared with Other Codecs	4
JPEG 2000 Standardization Status and Applications	5
Scalability of JPEG 2000	6
JPEG 2000 and the Infinity Product Range	7
JPEG 2000 and NLEs	7
Conclusion	8

Why is Compression Needed?

The technically ideal production workflow wouldn't use compression at all. Unfortunately, "technically ideal" uncompressed HD would need a 155 MB/s stream to represent it.

Just to put that into human terms, a bookshelf full of books that is one meter long could be represented by 100 MB of data. So now imagine that bookshelf growing in length at a rate slightly faster than a walking pace (over 5 kph or 3 mph) and it's easy to realize that 155 MB/s is too much data to handle.

After considering costs of purchasing, and managing this volume of books (and shelves!) the technology that can shrink the space requirements by a factor of around twenty becomes very attractive—as long as it doesn't seriously affect the quality of the pictures; and here we get to the serious point of video compression: it involves making compromises.

Sometimes, the tradeoff is between picture quality and the amount of storage needed to hold the compressed images. Sometimes it is between the codec's computational complexity and the bit-

rate of the compressed video stream. And at other times it can be between the ability to address and manipulate the images in the video stream and the complexity of the stream itself.

The right compromise is determined by what you need to do with or to the video stream—which is another way of saying that it's determined by the stage in the broadcast production workflow in which you find yourself. So let's define the stages in this chain of different operations, each with a common requirement for great-looking pictures, but differing priorities in how they are used.

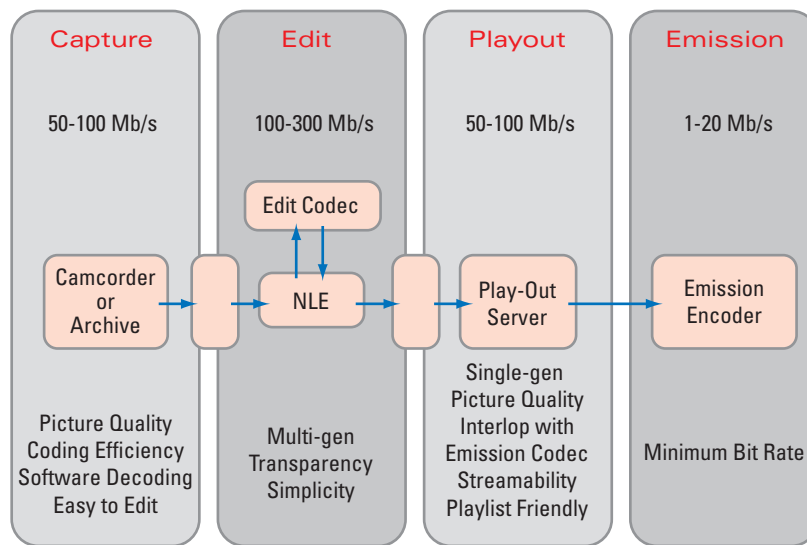


Figure 1 – Production workflow separated into four stages, namely: capture, edit, playout and emission; and summary of the codec priorities in each of these stages.

$$1920 \times 1080 \times 10 \text{ bits} \times (1 \times Y + \frac{1}{2} \times Cr + \frac{1}{2} \times Cb) \times 30 \text{ fps} + B \approx 155 \text{ mb/s}$$

Capture

“GIGO” or garbage-in/garbage-out is especially true in video production and this places special emphasis on the codec used at its very start.

While upgrades to the stages downstream of capture are always technically possible in the future, they will only be worthwhile if the input material is of sufficiently high quality. By noting today’s recordings are tomorrow’s archives, the conclusion that

the capture stage is not a place where picture quality should be traded-off lightly is impossible to avoid. Yet that has been common practice until recently.

SMPTE 292 HD-SDI encodes HD video at resolution of 1920 x 1080, using 10 bits per sample and 4:2:2 color sampling. Using this as the “gold standard,” some of the compromises used at the capture stage are as follows:

- Sub-sampled raster: instead of capturing a full 1920 x 1080 raster, some systems only capture 1440 x 1080, thereby discarding 25% of the available information.
- Sub-sampled colors: instead of capturing a full 4:2:2 color-space, some systems only capture 4:2:0 or worse, thereby discarding another 25% of the available information.
- Truncated samples: for example, using only 8 bits per sample – another 20% of the signal gone forever. And this is before any compression has even occurred!

Systems that use legacy codecs such as MPEG-2 and DVCPRO have typically made at least one and sometimes all three of these compromises which can accumulate to an overall information loss of 55% before encoding. The reason for this is simple: these encoders

simply couldn’t produce good-looking compressed pictures at reasonable bit-rates unless most of the HD signal was discarded.

Systems that use legacy codecs such as MPEG-2 and DVCPRO have typically made at least one and sometimes all three of these compromises which can accumulate to an overall information loss of 55% before encoding.

The new generation of advanced codecs such as JPEG 2000 and H.264/AVC changes these rules because higher compression ratios can be achieved without degrading the picture quality. However, “advanced codec” can also mean “computationally complex codec,” which is an acceptable trade-off for the capture stage, but not necessarily good for the next, edit stage.

Edit

Different requirements are important inside the NLE and today are mainly met by so-called “intermediate codecs” such as HQ, CineForm, ProRes 422, and DNxHD that are characterized by having high transparency, but relatively poor coding efficiency.

A picture can go through multiple encode/decode cycles in creating an effect, so the codec used needs to be very transparent with excellent multi-generation performance. However, hard disk and RAID storage is so fast and economical today that bandwidth isn’t too much of a concern on modern hardware platforms.

A special requirement of intermediate codecs is their ability to support multiple simultaneous decodes. Consider the fairly common example of putting four different video streams into an effect

simultaneously, maybe showing each in its own quarter of the screen—this requires four decodes and one encode to happen at the same time on a single workstation. Real-time NLE performance on practical hardware only comes with codec simplicity.

Play-out

It’s normal for NLEs to render or conform to a good playout codec. Typically, only one encoding is done and the requirements focus is a balance between transparency and efficiency.

Multi-purpose playout is an emerging requirement in this stage of the workflow. This is where HD, SD, and mobile versions of an asset can be produced from the same file at the same time. More on this later.

Emission

Very low bit rates with high picture quality matter most here, so we’re back to an advanced codec such as H.264 with a long GOP necessary to stay within the limited bit-rate budget.

So having established the importance of the capture stage to the performance of the overall system, and outlined the main priorities of codecs used here, let’s take a closer look at the requirements of a capture or acquisition codec.

Acquisition Codec Requirements

An acquisition codec has at least nine bosses (Figure 2). And some of them are giving different orders—for example some codecs get high coding efficiency by using a long GOP structure, but this increases the latency—or how much time it takes between pictures going into the encoder and coming out of the other side.

That’s why designing a great acquisition codec is such a puzzle.

The open, non-proprietary JPEG 2000 standard solves this puzzle in the following ways:

- Using an I-Frame only approach eliminates the complexity issues that can make editing difficult. It also minimizes the latency
- Because of the advanced and sophisticated compression algorithm, it’s still efficient, even without using long GOP
- Uniquely, it also provides scalable resolution and picture quality—more on that later
- It uses wavelets instead of DCTs (discrete cosine transforms) which eliminates the blocking artifacts we’ve grown so used to with other codecs that quickly degrade multi-generation performance
- It performs really well in multi-generation tests
- And, as implemented in the Infinity product range, it captures all the colors and the entire picture raster

So, JPEG 2000 is an excellent acquisition codec, but it isn’t the only one—so how does it compare with the alternatives?

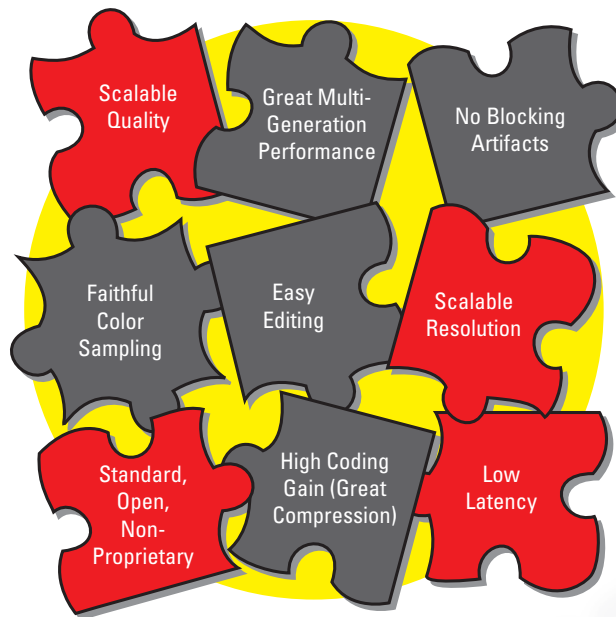


Figure 2 – Acquisition codec requirements.

JPEG 2000 Compared with Other Codecs

Figure 3 shows a side-by-side comparison of codecs that are commonly used for HD acquisition.

We’ve already talked about how I-frame only coding reduces latency—which can be important if you’ve got an anchor talking with a journalist live in the field, as some GOPs are over half a second in length.

There’s a high-quality I-frame at each end of a GOP. These are used as the basis for the in-between pictures. The in-betweens tend to be noisier because they’re working at much higher com-

pression ratios than the I-Frame pictures and can even lead to the picture quality “pumping” every GOP.

One of the criticisms often leveled at JPEG 2000 is that its sophistication means you can’t decode it in real-time on a laptop. Well that isn’t true, especially once the inherent scalability is exploited in a sensible way such as in the EDIUS® and Aurora™ NLEs—more on that later.

Perhaps the biggest difference between JPEG 2000 and the other compression schemes is that when the bit-rate

is lower than ideal, its performance degrades gracefully and compressed pictures get softer, instead of blocky. The reason for this is fundamental: only JPEG 2000 uses wavelets, all the rest use DCTs. This is important because all those block edges are actually really hard to re-encode which can put you in a downward picture quality spiral further along the production chain. In contrast, softer pictures are actually easier to encode, which eliminates the potential for further downstream damage to your production values.

Features	MPEG-2 I-Frame	MPEG-2 Long GOP	MPEG-4 AVC	JPEG 2000	JPEG 2000 Benefits
Intra-coding Only	Yes	No	No	Yes	Easy editing. Avoid introducing noise caused by GOP pumping, GOP alignment
Low Latency	Yes	No	No	Yes	Avoid delays in live news interviews. Reduce AV sync issues
Efficient Coding	No	Yes	Yes	Yes	High-quality pictures in manageable size files
Scalable Decoding	No	No	Yes	Yes	Real-time browse even on a laptop. Embedded proxy eliminates mismatch
Main Artifacts	Blocks	Blocks and motion aliasing	Blocks and motion aliasing	Softness	Downstream encoders use bits to encode your pictures instead of artifacts
Multi-Gen Performance	OK	Poor if GOPs misaligned	Poor with deblocking	Good	Confidence that the emission picture quality is as high as it can be
Open Standard	Yes	Yes	Yes	Yes	Interoperability. Choose from multiple vendors
No/Low Licensing	No	No	No	Free	Lower costs, richer toolsets
10-bit, 4:2:2 Color	No	No	Sometimes	Yes	Accurate green/blue screen work today. Perfect pictures from your archive tomorrow

Figure 3 – HD codec comparison.

JPEG 2000 Standardization Status and Applications

The process to standardize JPEG 2000 started in 1996 when Ricoh, a member of the JPEG committee, proposed a program to make major improvements to the existing JPEG standard for still pictures.

Some 24 suggestions, evaluations, and four years of committee work later, the main parts (1 & 2) became standards under the mandates of the International Standards Organization and International Electro-technical Commission (ISO/IEC). This happened in 2000, which is why standard 15444 is known as JPEG 2000.

From the outset, there were plans for enhancements to the standard—is actually comprised of 12 parts, numbered 1 to 13; (there is no part 7) and the majority of these have been formalized since 2003 and 2006.

The remaining parts, having to do with volumetric imaging and security, are still being worked on.

The benefit of standards is that they encourage companies to develop technologies, products, and interoperability methods—some examples are on Figure 4. Perhaps the most famous is Digital Cinema Initiatives LLC (DCI), which undertook a serious investigation into codec technologies before settling on JPEG 2000 for their specification, which is designed to last at least as long as the previous 70 years of the SMPTE specification for 35 mm film!

You'll see that Grass Valley™ isn't the only camera manufacturer using JPEG 2000 – with Ikegami and RED also finding applications for it.

JPEG 2000 also finds a number of applications outside the broadcast industry:

- It's used for medical imagery where faithful compression is an absolute requirement to avoid potential misdiagnosis lawsuits due to compression artifacts
- When you see pictures get gradually sharper on your internet browser, that's JPEG 2000

Having such a wide range of applications and large installed base is a good thing. For example, the benefits of meeting military requirements have often flowed into those of broadcasting. This is similar to the way billions of dollars invested in IT storage technology can lead to better and more affordable solutions that, left to itself, the broadcast industry would not have been able to make.

Now let's examine one of the unique characteristics of JPEG 2000, its scalability, and the benefits that it brings to broadcast applications.

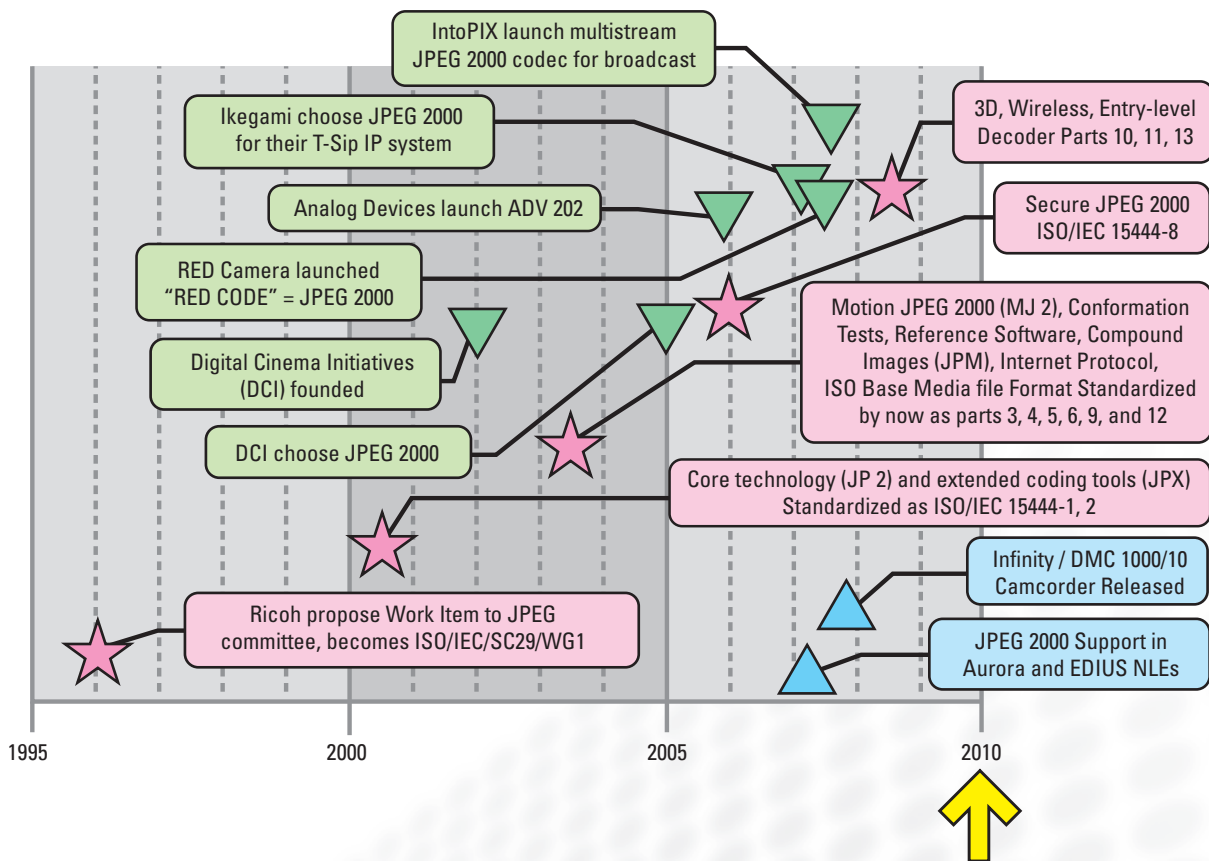


Figure 4 – JPEG 2000 applications.

Scalability of JPEG 2000

JPEG 2000 allows the information that makes up a full-resolution image to be ordered in three useful ways:

- By resolution, so you can decode a high quality quarter size picture first, followed by a high-quality, half-size picture, and finally the high-quality, full-resolution image
- By quality, so you get a full-resolution, but fuzzy image first that gets progressively sharper as more data is decoded
- By region of interest, so you decode part of the picture at full-resolution and full-quality first, followed by the rest of it

The Infinity™, Aurora, and EDIUS products exploit the first of these two types of scalable decoding as illustrated in Figure 5.

The stacked bar on the left shows how the encoded bit-stream is layered. This allows a decoder to use just the parts that are needed for the given application. The 6 Mb/s layer is used as an “embedded proxy.”

On the right shows the resolutions supported by each layer:

- For HD, decoding layer 1 gives a quarter-resolution image, followed by a half-resolution image with layer 2, and the full-resolution delivered with layer 3

- For SD, decoding layer 1 gives a half-resolution image, layer 2 gives a satisfactory full-resolution image, that is sharpened to full-quality by decoding layer 3

This is useful because it allows you to work with full HD material at full speed on a laptop—just one of the promises of scalability that has been realized today.

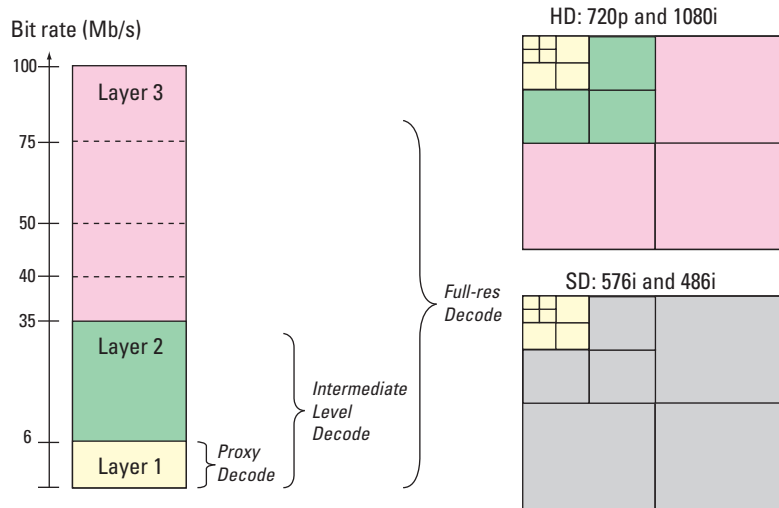


Figure 5 – JPEG 2000 decoding options.

JPEG 2000 and the Infinity Product Range

Figure 7 shows the formats that the Infinity product range supports today and the various JPEG 2000 profiles that support them.

SD is still important which is why our products support it. The HD profiles support two main applications:

- EFP, where quality is at a premium because assets captured today are likely to be monetized over a long period and need to have good archival characteristics. We recommend the 100 Mb/s profile for this work
- ENG is more about speed and pictures for today, so fewer bits per second are supported by the 75 and 50 Mb/s profiles, although the 50 Mb/s profile is only provided to allow users to trade off picture quality against time to send over a network

The two media types supported by the Infinity product range allow for plenty of affordable storage, from a minimum of 15 minutes of the highest profile on CompactFlash, to over an hour of the lowest HD profile on REV PRO™ disks.

JPEG 2000 and NLEs

What about NLE support for JPEG 2000? If you believe what some people say, it can't be done.

Well, there are two editing systems that prove otherwise: Aurora for fast-turn work, and EDIUS for craft applications.

Both of these support fast and easy content browsing, including jog, shuttle, and real-time playback of the embedded proxy that's in every Infinity JPEG 2000 file. Even on a laptop!

On a normal turnkey Aurora 6.3 or later workstation, JPEG 2000 can be transcoded directly to the timeline or to an edit bin.

Transcoding directly to the timeline means you select only those parts of the original clips you need so that no time is wasted transcoding material you don't want.

The fast-turn Aurora NLE can then push the finished program to a play-out Grass Valley K2 server that can start playing to air even before the transfer has completed.

The agile codec support in EDIUS means you can use JPEG 2000 directly on the timeline—with no need to transcode and you can mix and match with material recorded with virtually any other codec. Also, EDIUS makes smart use of the quality and resolution layering to keep the editing experience both fluid and pristine.

The craft-focused EDIUS NLE can then conform content into any format that's useful—so the target can be DVD, a playout server, or any number of media file types such as Windows Media. The more recent EDIUS releases also support export direct from timeline to high-definition Blue-ray.

Additionally, there are third-party tools available to enable the use of JPEG 2000 with other edit systems such as Apple Final Cut Pro.

Raster	System Hz	Bit Rated Mb/s	Recommended Use	Trec 35 Gb Minutes	Trec 16 Gb Minutes
1920x1080i	50	100	EFP	>30	>15
1280x720p	59.94	75	ENG	>45	>20
		50	"Longplay" mode	>60	>30
720x576i	50	50	These bit rates are good for any purpose	>60	>30
720x486i	59.94	40		>75	>40
		30		>100	>50

Figure 7 – Infinity JPEG 2000 formats.

Conclusion

In this paper we have shown that JPEG 2000 is an ideal acquisition codec for serious HD productions, because it allows the full picture to be captured in full-color at economical bit rates.

Even at the highest bit rates, its scalable characteristics make it very usable on today's laptop and workstation edit systems, with a limited number of full bit rate real-time video tracks. As computing hardware improves, the number of real-time layers can only increase.

For high-end applications, transcoding to an intermediate codec is also an efficient workflow. JPEG 2000's unobjectionable compression artifacts help ensure that this is a clean as well as a quick operation when bit rate is at a premium.

In play-out applications, JPEG 2000's scalability holds the promise of supporting "create once for multiple purposes" type operations.

JPEG 2000 is not suitable for emission applications, but it is a perfect, future-proofed archive format.

SALES

Local and regional sales contacts can be found by visiting
www.grassvalley.com/sales

SUPPORT

Local and regional support contacts can be found by visiting
www.grassvalley.com/support