grass valley

WHITE PAPER

Introducing HQX

What are the main features of Grass Valley's Intermediate Codec, and for what projects is it the best choice?

Akira Takemoto, Desktop & Enterprise Solutions

December 2010

The codec you choose to use within your NLE is becoming increasingly critical and confusing as the growth in the type and number of input formats is exceeded only by that of the delivery formats your clients ask for. In this paper, we discuss which workflows are best served by a so-called "Intermediate Codec." We go on to define the key characteristics of Intermediate Codecs and introduce HQX, which is the latest generation in Grass Valley's "HQ" family of codecs.

Intermediate Codec solutions from several vendors are then compared in terms of picture quality, flexibility, and suitability for use with graphics as well as video. We find that in like-for-like tests, HQX offers picture quality that compares favorably with other popular Intermediate Codecs. We also show that HQX has superior multi-generational performance, which when combined with rate and resolution flexibility and built-in alpha channel support, the conclusion that the Intermediate Codec is the best choice for creative artists seeking the best possible image quality.

Contents

Introduction

It seems that the number of different sources and formats that content producers have to use as input material multiplies every day. You may be bringing in RAW 4k content from a RED One, or pulling in HD from a Grass Valley™ LDK 3000. At the same time, your graphics department is providing you with material from After Effects. The format explosion on the input side is mirrored by the growing variety of delivery formats that your clients will ask for, whether that means video files, podcasts, webstreams, or a host of other formats.

As an editor, the value you bring to your client is your creativity, combined with the ability to meet exacting technical production values with on-time delivery. We think that your creativity is maximized when your NLE is responsive, so that experimenting with cool effects and compositions doesn't mean waiting around for renders; and you can put aside worries about whether the layers you're building up are compromising the image data. This all adds up to a good set of reasons to use an Intermediate Codec, such as HQX. So let's start with an overview.

Intermediate Codecs

The performance of an NLE workstation is determined by its image processing power and how quickly content can be moved in and out of storage. Your workstation slows down once either limit is reached. Uncompressed HD content at 155 MB/s[1] will max out the bandwidth of most storage systems[2]. So, were you to use uncompressed HD, the CPU would spend most of its time idling while waiting for the storage to catch up.

Of course, the reason uncompressed video is attractive, in theory, is that it is the best image quality possible. Intermediate Codecs have been developed to solve the problem of storage bandwidth while maintaining picture quality that's visually identical with the uncompressed master. They share the following characteristics:

- Enough compression to relieve the storage bandwidth as a bottleneck – compression ratios tend to be in the 5:1 to 8:1 range. That means a single SATA 1.0 drive can now support four or more channels before maxing out, even after allowing for the overhead of sharing the bandwidth

- Not super-complicated because the CPU is still needed for effects and editing, not just compression

- Transparent image quality – Intermediate Codecs are lossy, but are much less aggressive than acquisition or playout codecs

This achieves a balanced use of the workstation's key resources (CPU and storage bandwidth) without making unacceptable compromises with the pictures.

A typical workflow that uses an Intermediate Codec has three stages:

1) Browse, log, and trim the source material, and convert it from the acquisition format to the Intermediate format

2) The editing session

3) Render the finished piece to the required delivery format(s)

Both the first and the last steps can be time-consuming, so the next section explores the type of projects where using an Intermediate Codec is the right choice.

---

[1] 1920x1080 pixels x 30000/1001 frames per second x (1 + ½ + ½) components x 10 bits per component ÷ 8,000,000 bits per megabyte.

[2] A single SATA hard disk can just about saturate the ~180 MB/s physical layer specified by SATA 1.0.

Intermediate Codecs (cont.)

When to Use an Intermediate Codec

The choice of codec comes down to three factors:

1) Use a "System Codec." It may be possible to use the same codec that was used for acquisition. This has genuine advantages for fast-turn work, especially if both the input and delivery material uses an advanced codec such as AVC-Intra. However, if the acquisition codec uses compromises such as picture scaling, under-sampled chroma, or only 8-bit sampling, it's not so advisable. And, of course, if a mixture of formats were used for the input material, this approach isn't possible.

2) Use "Mixed Formats." You can choose to use mixed formats as-is, and this is a great choice for fast-turn work because you can get started right away. This choice loses its allure for jobs that are more craft oriented and multi-layered because of the issues that will be encountered as multiple generations of the input material are generated by the editing process.

3) The other choice, and the one to which the rest of this paper is devoted, is to use an Intermediate Codec, which delivers these important benefits:

   • All worries about managing the picture quality while building complex edits are eliminated, and this enables you to focus on the creative process, not the technical one.

   • An Intermediate Codec is absolutely the fastest way to get content off your hard drive and onto your screen. This means that waiting for effects to render isn't an issue, and by taking the waiting out of working you can stay "in the moment" and at your most creative throughout an editing session.

EDIUS® 6 from Grass Valley allows you make any and all of these choices so that you can match your workflow to the type of job at hand.

Popular Intermediate Codecs

Table 1 introduces the main Intermediate Codecs in common use today and summarizes their key characteristics.

Codec	Compression Range	Bit Rate (Mb/s)	Bit Depth	Chroma	Alpha Support
HQX	From 25:1 to 2:1	45 to 600	10	4:2:2 YCbCr	Yes, 10-bit
DNxHD	8:1 or 5:1	145 220	8 8 or 10	4:2:2 YCbCr 4:2:2 YCbCr	No No
ProRes	8:1 ProRes422 5:1 ProRes422 (HQ) 7:1 ProRes444	147 220 330	10 10 Up to 12	4:2:2 YCbCr 4:2:2 YCbCr 4:4:4 RGB	No No Yes
CineForm	From 10:1 to 3:1	120 to 400	Up to 12	4:2:2 YCbCr or 4:4:4 RGB or RAW	No

Table 1 – *Summary of Intermediate Codecs.*

CineForm and HQX share the important features of:
• User-definable compression ratios, plus
• Built-in alpha channel support

We'll cover why alpha channel support is useful to the point of necessity later in this paper.

Note that both DNxHD and ProRes422 only offer fixed bit rates, and neither provides alpha support in video color space.

Remember that the purpose of an Intermediate Codec is to emulate uncompressed HD. Uncompressed HD has 10-bit sampling, which is why most of these support at least this, with DNxHD being the exception. The importance of those two bits is covered later.

Now let's briefly review the evolution of HQX.

HQX Evolution

HQX is the latest stage in the evolution of the HQ family of codecs that have been provided as part of the EDIUS NLE.

Features	Released	Resolution	Chroma	Bit Depth	Alpha Support
HQ	2004, EDIUS 3	1440x1080	4:2:2	8	No
HQ with Alpha	2006, EDIUS 4	1920x1080	4:2:2	8	Yes
HQX with Alpha	2010, EDIUS 6	Up to 4k	4:2:2	10	Yes

Table 2 – HQX evolution.

Table 2 highlights the main developments in the HQ family of codecs since 2004. It shows support for steadily increasing resolutions and that provision for alpha support as a fourth component has been built right into the codec and has been a feature of the HQ family since 2006. As we've discussed above, this is one of those features that distinguishes a true Intermediate Codec from those designed more for acquisition or transmission.

In addition to increased resolution, the important upgrade included with HQX is 10-bit video. Let's take a closer look at what benefits those extra bits bring to your editing session.
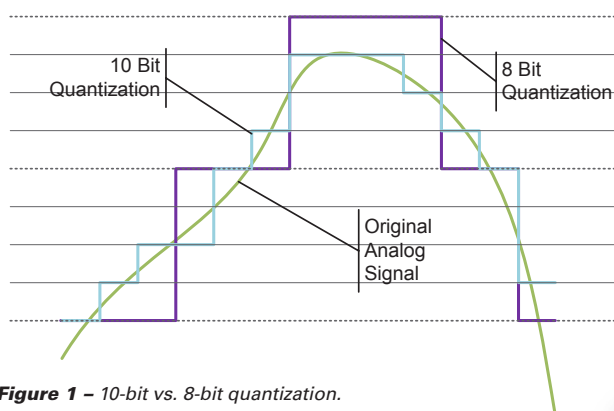
Superior Performance of 10-bit Codecs

When a picture is digitized, errors are introduced by the process of turning a continuously variable signal into a series of numbers that have finite precision. For example, say you have a real-world signal that can range between 0 and 1V. With eight bits of precision, you'd assign the minimum code value (zero) to represent the bottom of the range (0V), and the maximum value (255) to represent the top of the range (1V). With 10 bits of precision, you represent 1V with the maximum code value available with 10 bits, which is 1023. An 8-bit coder can represent 256 distinct levels in the 0-1V range, but a 10-bit one can represent 1024 distinct levels.
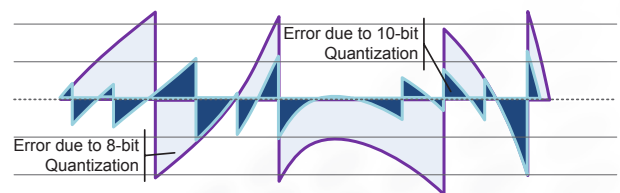
Look at what happens when you encode a real-world signal at 0.489844V in each system:

- With an 8-bit range of 0-255, the ideal coding is 125.4, but an 8-bit coder has to pick the nearest whole number of 125. This causes an overall error of -1.563 mV.

- With 10-bit precision in a range of 0-1023, our code value is 503, rounded down from the ideal value of 503.0753. The error introduced here is just 73.6 µV.

In this example, the quantization error of 8-bit precision is approximately 21 times larger than that for 10-bit precision, which is somewhat of an extreme! On average, 8-bit precision introduces quantization errors that are four times larger than those of 10-bit sampling. This concept is illustrated graphically in Figure 1 and Figure 2.
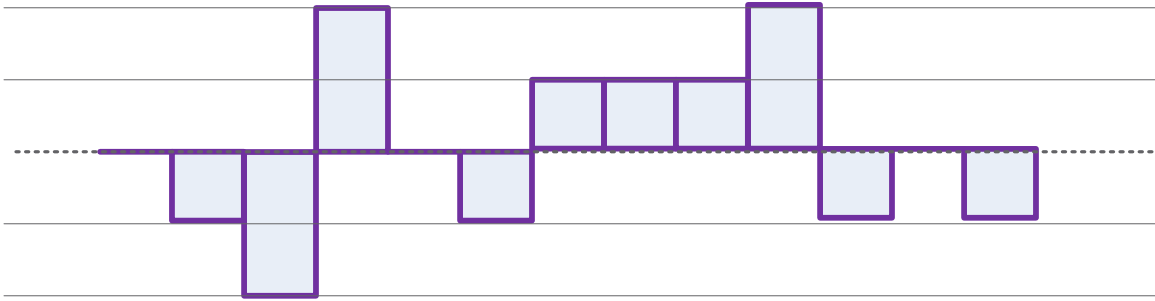

Figure 1 – 10-bit vs. 8-bit quantization.

Figure 2 – Comparison of quantization errors.

Quantization is performed by capture devices such as cameras or field recorders, which are increasingly capable of capturing pictures with at 10-bit precision. The quantization errors (or noise) introduced by the capture device are unavoidable, but only happen once. What does this have to do with the Intermediate Codec used? The answer is whether or not a second, much noisier set of errors is introduced when an 8-bit Intermediate Codec is used to compress a 10-bit signal.

Superior Performance of 10-bit Codecs (cont.)

Figure 3 – Noise added by 8-bit quantization of 10-bit signal.

Figure 3 illustrates the result of this process, and Figure 4 shows a picture side-by-side with noise in just the bottom two bits of each pixel[3].

Figure 4 – A picture and the noise in its bottom two bits.

Quantization adds noise to your picture which is often noticeable as visible banding or contouring, though sometimes a process called "error diffusion" is used to dither the contours so that they are less noticeable, but the total amount of noise is still present. We'll discuss why this is so important after considering another benefit that using 10-bit precision provides: more accurate colors.

Chroma Precision

Because two components are used to code chroma information, 8-bit video can represent, at most, 50,000 different chroma values (code values below 16 and above 240 are not legal, so there are 224, not 256 different values possible for Cb and Cr).

10-bit video can represent 16 times as many: that's over 800,000 different chroma values

One of the important benefits of representing color more accurately is when pulling mattes from green screens (or blue, or red) – you can get much finer tolerances on the colors that are in or out of the matte, which makes your compositions look more convincing.

Now let's look at the structure of the HQX codec, review its capabilities, and see why using 10-bit precision helps the codec deliver the transparent picture quality required of Intermediate Codecs.
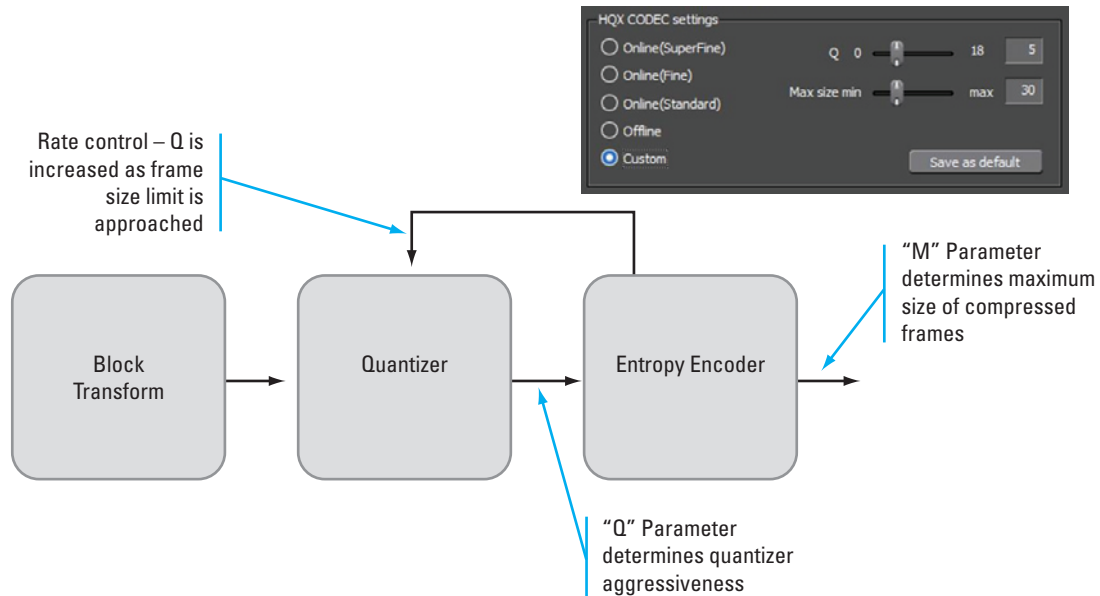
[3]Note: A gain of 40 has been applied to make the noise easily visible.

HQX Architecture and Profiles

HQX has an architecture that's common to most video codecs but does put the user in charge of two key parameters via the dialog box shown in Figure 5. The "M" parameter fixes the maximum bit-rate as a fraction of the uncompressed rate. So, M=10, corresponds to 10:1 compression, M=20 corresponds to 5:1 and so on.

The "Q" parameter governs how aggressively the algorithm compresses the image, with higher values of Q corresponding to more aggressive compression.

Figure 5 – *HQX codec architecture.*

Uncompressed images are first split into small blocks, each of which is transformed. Think of the transform as the equivalent of taking a large amount of paperwork, and sorting it into a well-designed filing system. It doesn't take any picture information away, but by organizing it better, it does reduce the entropy for normal pictures. Of course, for noise, it can't do this.

The Quantizer is where the decisions are made about what information can be discarded. It does this in an intelligent way that preserves the most significant information.

The entropy encoder is a lossless, arithmetic coder, similar to algorithms like WinZip. The feedback loop causes the quantizer to become more aggressive if the maximum frame size could be exceeded.

There are several decisions with any codec architecture that determine the overall efficiency and performance of the codec, and the handling characteristics of the compressed material. Table 3 summarizes the choices that were made for HQX.

HQX Design Characteristics

Parameter	HQX Setting	Benefit
Resolution	Flexible, with support for up to 4k	Work with your content in its native resolution, but without the time penalties of using a resource hungry-acquisition codec.
Chroma Sampling	4:2:2	Matches the chroma sampling of HD-SDI so no chroma information is lost by HQX.
Sampling Precision	10 bits	Lower noise, better compression, avoids banding or contours, finer control on chroma keys, apply gain or level shifts to eight-bit contribution material (e.g., XDCAM) without introducing undesirable artifacts.
Alpha Support	Yes, 10-bit	Graphics are simple to work with once they've been encoded as HQX because access is fast, and high picture quality is maintained.
Bit Rate	User defined as a fraction of the source material's bit rate. Typical values for HD are 145, 220, and 400 MB/s	The user decides the amount of compression that's right for the job. High bit rates are available when near-uncompressed quality is desired.
Bit Rate Variation	VBR, with strict maximum	Better pictures for the same number of bits compared to CBR schemes. Strict maximum avoids network and storage glitches that can be caused by sudden spikes in bit-rate exhibited by some Intermediate Codecs.
Frame Structure	Intra coding only	Easier clip management and editing.

Table 3 – *HQX design characteristics.*

Codec Performance and 10-bit Precision

Although there are many different image and video codec algorithms available today, with their supporters highlighting the technology used, whether it be fractals, wavelets, or DCTs, they all operate on a common set of basic principles.

Principle #1: All images contain information
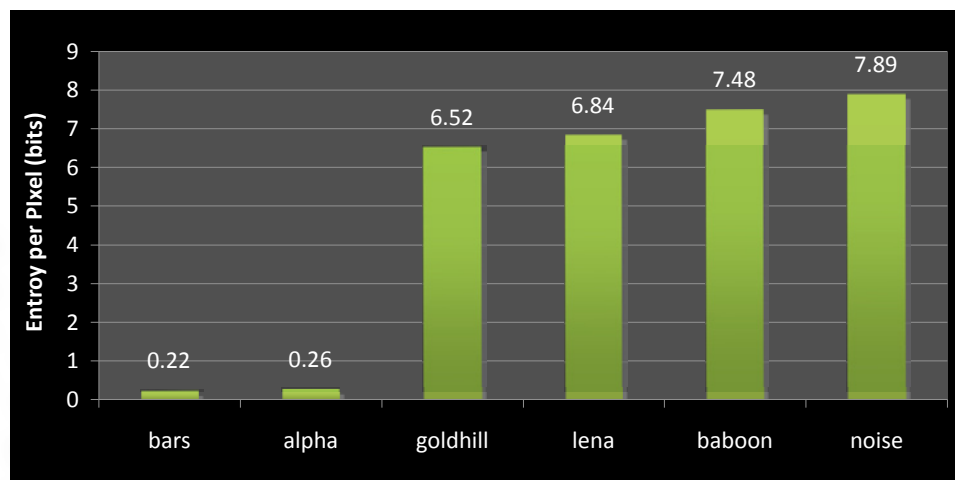
How much information depends on the picture.

- For instance, color bars have very little information because it's almost trivial to predict the value of each pixel based on its neighbors.

- Highly detailed images that contain lots of fine texture contain much more information, because it's not at all trivial to predict the value of any given pixel with accuracy.

Information scientists use the word "entropy" as the measure of the amount of information that an image contains. In layman's terms "unexpectedness" is a good way of understanding entropy. Things that are unlikely and unpredictable contain more information (and entropy) than things that fit an obvious pattern.

An informal way to measure the entropy of an image is to compress it with a lossless algorithm such as WinZip.

Name	Bars	Alpha	Goldhill	Lena	Baboon	Noise
Picture Size (bytes)	1,049,760	307,200	414,720	262,144	65,536	250,000
WinZip File Size (bytes)	9,780	10,077	337,947	224,019	61,299	246,572
Estimated Entropy per Pixel (bits)	0.22	0.26	6.51	6.83	7.48	7.89

Table 4 – *Entropy content of various images.*

Figure 6 – *Entropy content of various images.*

Figure 6 and Table 4 show the amount of entropy in each picture expressed as "bits per pixel." And sure enough, it shows that color bars and typical alpha channel content contain hardly any entropy, but that normal pictures have quite a lot, and noise is nothing but entropy.

Codec Performance and 10-bit Precision (cont.)

Principle #2: Random noise is all entropy

In lay terms, lossless codecs like WinZip work by separating the entropy in a file from the "padding" (i.e., two extra null bits inserted at the end of an 8-bit word to become a 10-bit word). The entropy is saved into the compressed file, and the "padding" can be discarded. When de-compressed, the padding is put back to exactly reconstruct the original data.

When a signal is random noise, it is all entropy, so there is no "padding" to discard. By the way, that's why you can't compress files that have already been compressed—for example zipping a JPEG rarely gains you any disk space—because the entropy of the JPEG file has been maximized to the best ability of the JPEG algorithm.

Principle #3: Lossy codecs need to discard entropy as well as "padding"

A lossless codec is doing well if it achieves a compression ratio of 2:1 (i.e., the compressed file is half the size of the uncompressed one). The compression ratios used by Intermediate Codecs are higher than this, with ratios of 5:1 and 8:1 in common use. This means that the Intermediate Codec must discard entropy with the "padding," which makes it a "lossy" codec.

Of course, the more entropy the original image contains, the more of it will have to be discarded during the encoding process.

This is why there is nothing to be gained ignoring 20% of your image data—because doing so adds enough quantization noise to counteract any sought-for benefit such as better picture quality or more coding efficiency. The somewhat counter-intuitive result is that, for the same image quality, codecs using 10-bit precision can achieve the same bit-rates as codecs that only use 8-bit precision on 10-bit input data. This finding is demonstrated in Table 5, which shows that the long-run picture quality for codecs with 10-bit precision (HQX and ProRes422) is very close to DNxHD, which is an 8-bit codec in the 8:1 class of compression profiles.

So, as the extra two bits effectively come for "free," why not take advantage of their benefits?

Picture Quality Comparisons

The normal way of making objective comparisons between different video codecs is via their peak signal-to-noise ratio (PSNR) performance for some well-known reference sequences. Although the PSNR does not model the way in which the human visual system perceives noise created by image compression, it is both objective and straightforward to compute (see sidebar). Differences between codecs are only significant if the difference in PSNR exceeds approximately 2 dB. It's important that the same reference content is used when comparing codecs because codec performance is a function of the uncompressed image.

The European Broadcasting Union makes available a set of uncompressed reference sequences that are commonly used to evaluate codec performance.
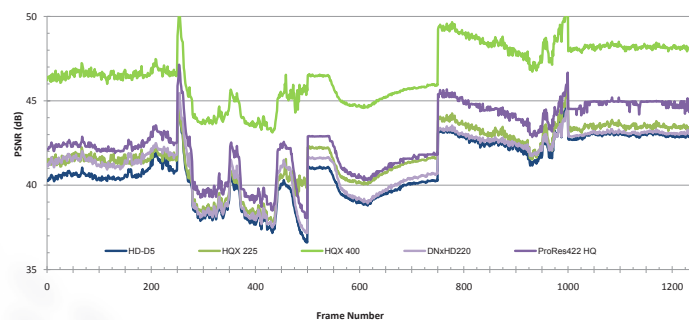
The EBU sequences, comprising 1250 frames of video, were used to evaluate two key performance measures for a number of Intermediate Codecs.

Single Generation Performance

The first performance metric is to calculate the significance of the errors that are introduced by a single encode/decode generation for each codec.

We've plotted the PSNR for each frame in the sequence for Intermediate Codecs in these three classes of increasing compression ratio: 3:1, 5:1, and 8:1. Remember bigger numbers are better—the line at the top of the graph shows the codec with the best picture quality on a frame-by-frame basis.

To provide a long-run quality measure, the average PSNR for each codec over all 1250 frames has also been computed. The results are presented in Table 5.

The acronym "PSNR" stands for "Peak Signal to Noise Ratio"

It's a handy, single-figure measurement of how different a picture that's been through a codec is compared to the uncompressed original.
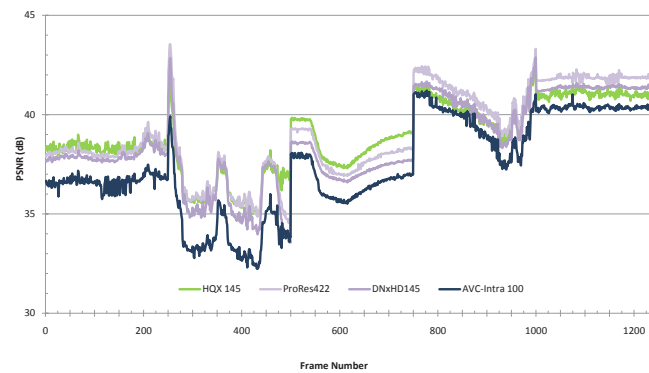
Its calculation requires a few simple steps:

1) Compress and decompress a reference picture to produce a resultant picture.
2) Subtract the resultant from the original; the non-zero values in the resultant are errors that were introduced by the codec.
3) Square all the resultant pixels and add them together, then divide by the total number of pixels to get a figure for the Mean Square Error.
4) Form the Ratio by dividing the Mean Square Error into the Peak value a pixel can represent (SMPTE 292 defines "peak white" as code value 940, but provides headroom up to 1019).
5) Express the ratio in dB.

$$MSE = \frac{1}{N \cdot M} \cdot \sum_{i=1}^{M} \sum_{j=1}^{N} (ref_{i,j} - res_{i,j})^2$$

$$PSNR = 10 \cdot \log_{10} \left(\frac{Peak^2}{MSE} \right)$$

Figure 7 – PSNR comparison of 5:1 class codecs, and HQX at 3:1.

Figure 7 presents the PSNR results for our test sequence of 1250 reference images. The trace with the best PSNR is HQX at 400 Mb/s with the "Q" parameter set to zero and the "M" parameter set to 35. As you can see, it is much better than the other codecs in this trial because the bit rate is significantly higher than the others.

The traces for the 5:1 class show that HQX at 220 Mb/s is comparable to the other Intermediate Codecs, and noticeably better than D5.

If we reduce the average bit rate, we encounter the class of codecs using an 8:1 ratio. The PSNR results for which are shown in Figure 8.

Picture Quality Comparisons (cont.)

Figure 8 – PSNR comparison of 8:1 class codecs.

Here we see that HQX is the best codec for some sequences, and always very close to the ProRes422 and DNxHD codecs in this compression class.
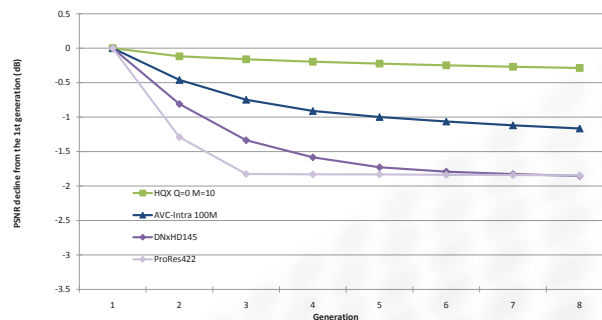
Codec	8:1 Class		5:1 Class		3:1 Class	
	Mean bit rate	Mean PSNR	Mean bit rate	Mean PSNR	Mean bit rate	Mean PSNR
HQX	145	38.9 dB	224	41.8 dB	400	46.4
DNxHD	145	38.6 dB	220	41.4 dB		
ProRes422	147	39.1 dB	220	42.7 dB		
AVC-Intra	100	37.4 dB				
HD-D5			220	41.0 dB		

Table 5 – *Summary of picture quality comparisons.*

The long-run performance of each codec is summarized in Table 5. It shows that HQX provides comparable quality to DNxHD and ProRes422 over the long run, at comparable bit rates. It also shows that HQX allows you to turn the quality dial all the way up should your project merit it.

Multi-generation Performance

The second key performance metric for Intermediate Codecs, in particular, is how they behave when a picture is compressed and decompressed several times. If you build up your edits in layers, then it's likely that you are creating multiple generations of the original input material, and this causes compression errors to accumulate across generations. The performance differences between the Intermediate Codecs are more marked here as illustrated in Figure 9.

Figure 9 – Multi-generation performance of intermediate codecs.

Figure 9 shows that HQX does the best job of maintaining picture quality over multiple generations, and that DNxHD and ProRes422 lose picture quality quickly—even over the first two or three generations.

Conclusion – HQX Maximizes Your Creativity

We have shown that it is important to be able to match the codec you use to the work you do. For fast-turn operations, a system codec, or using mixed formats is probably the best choice; but when image quality and creativity are paramount, an Intermediate Codec should be selected because it will deliver these important benefits:

- Input sources are handled transparently and consistently, even those including alpha channels such as graphics

- Technical concerns over picture quality are separated from the creative process

- An intermediate codec is the fastest way to get great-looking material between disk and screen

In terms of picture quality, HQX offers excellent performance when compared like-for-like with the other major Intermediate Codecs that are available from NLE manufacturers, including ProRes422, and DNxHD.

HQX offers superior performance in key areas when compared to other Intermediate Codecs:

- By providing direct support for alpha channels, it's fast and easy to work with graphics while remaining in the video color space

- By minimizing the degradation caused by multiple generations of encoding and decoding, the picture quality of your finished product will be better than most other popular Intermediate Codecs

- By putting you, the user, in charge of the bit-rate and quantization approach used, HQX allows you to offer your clients picture quality that is a clear 10 dB better than your competitors

So, when faced with a complex assortment of input sources, a demanding creative brief, and a tight deadline, the speed and transparency benefits of Intermediate Codecs in general, put together with HQX's superior performance in particular, form a combination that adds up to increased creativity for you, and increased value to your client.

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