(a) Original 1920x1080 RGB image resampled for this display to 500x281

Quality Analysis Comparing CineForm's Visually Perfect™ HD Codec versus Native Camera Formats

Testing Methodology

The following comparative quality analysis uses visual inspection, difference images, and PSNR (peak signal-to-noise ratio) as measures of image quality for comparing CineForm's 10-bit Visually Perfect™ codec to DVCPRO HD. PSNR analysis uses a standard mathematical model to measure an objective difference between two images. It is commonly used in the development and analysis of compression algorithms, and for comparing visual quality between different compression systems. For the results shown shown in the table, PSNR was calculated on the uncompressed luma channel of the YUV signal. For testing the DVCPRO HD codec the data was sent across HD-SDI for compression in the AJ-HD130DCP DVCPRO HD deck. For testing CineForm's 10-bit Visually Perfect codec the uncompressed YUV source was presented to the production version of CineForm's software codec that ships with Prospect HD. No other compression was introduced during testing of either codec. The above image (a) is a reduced version of the original 1920x1080 RGB image used for performing the PSNR analysis and difference tests.

Spatial Analysis

Image (b) below represents a windowed region from the first capture of the source image (a) using CineForm's Visually Perfect 10-bit codec. Images (c) and (d) represent the same windowed region after rendering through 5 and 15 generations respectively. Each of the images are 250x300 pixel regions extracted losslessly (PNG format) from the full-size image. Upon close inspection it is clear from the rendered CineForm images that there is no visually-discernable change in quality even between the originally-captured image (b) and the 15th-generation render (d). Hence the term "Visually Perfect".

CineForm Images - Captured and Rendered Through 15 Generations

(b) Windowed region of image originally captured using CineForm's codec.

(c) Windowed region from 5th-generation render using CineForm's codec.

(d) Windowed region from 15th-generation render using CineForm's codec.

DVCPRO HD Images - Captured and Rendered Through 15 Generations

(e) Windowed region of image originally captured using DVCPRO HD codec.

(f) Windowed region from 5th-generation render using DVCPRO HD codec.

(g) Windowed region from 15th-generation render using DVCPRO HD codec.

An identical analysis was performed using the DVCPRO HD codec. Images (e) - (g) show the originally captured image, the 5th-generation rendered version, and the 15th-generation rendered version respectively. At first glance it is difficult to notice a difference. To perform a closer inspection - that's okay, stick your eyes a few inches from the screen for a better look - look to the end of the overlaid white line on the left, where the lowest part of the building moulding touches the wall. Notice in the first-generation captured DVCPRO HD image (e) that substantial detail has been lost compared to the equivalent CineForm image (b). Further, compare (e) and (d) (the 15th generation CineForm render). Notice that the 15th generation CineForm render has substantially more detail than the originally captured DVCPRO HD image. If you now compare the originally-captured CineForm image (b) with the 15th generation (d), there no discernable loss of detail.

CineForm Difference Images Through 15 Generations

(h) Difference image between the captured CineForm image (b) and the original source image (a) for the windowed region.

(i) Difference image between the 5th-generation CineForm rendered image (c) and the original source image (a) for the windowed region.

(j) Difference image between the 15th-generation CineForm rendered image (d) and the original source image (a) for the windowed region.

DVCPRO HD Difference Images Through 15 Generations

(k) Difference image between the captured DVCPRO HD image (e) and the original source image (a) for the windowed region.

(l) Difference image between the 5th-generation DVCPRO HD rendered image (f) and the original source image (a) for the windowed region.

(m) Difference image between the 15th-generation DVCPRO HD rendered image (g) and the original source image (a) for the windowed region.

Difference Analysis

Another useful technique to analyze losses introduced during multi-generation rendering is to create difference images at each generation of the render and essentially perform a subtraction of the rendered image from a reference image. If the pixels are identical the pixel difference will be zero, but if pixels change during rendering the difference becomes non-zero. Non-zero difference pixels represent a loss of information compared to the original. In our examples below, "0" difference pixels are scaled up to 127 (middle grey), which is a luminance region where the human eye is more sensitive.

In performing the difference analysis on both codecs we have used the source image (a) as the reference in all cases. Images (h) - (j) show the differences for CineForm's Visually Perfect codec compared to the original image. Notice that you can see very minor differences in grey level representing minor deviations from the original image. When you compare (h) and (j) there is no discernable difference, demonstrating that CineForm's Visually Perfect codec holds up very well during multi-generation rendering.

In performing the same analysis on the DVCPRO HD codec, notice when comparing the difference between first-generation capture (k) and the original image (a) that substantial detail has been lost. This loss represents the "capture loss", and is due largely to the lower spatial resolution of the DVCPRO HD format. Also notice in the (k) - (m) difference images that visual loss continues to increase (more white areas) through successive rendering steps.

Note: To enhance the visibility of the non-zero differences, we have increased the contrast for all six difference images in an identical manner.

PSNR Analysis

When we analyze the difference data objectively, we can determine the PSNR (Peak Signal-to-Noise) measurement for each codec in decibels (dB). PSNR calculations use difference data for each codec like that shown in images (h) to (m) as part of the calculation. Before we discuss the data, it is useful to know that the human eye does not have enough sensitivity to detect changes in visual data for PSNR measurements above approximately 50dB, although this may vary in a minor way for each person. This is why the term "Visually Lossless" is often used.

In the Table below, the blue-circle plot indicates the PSNR performance for CineForm's Visually Perfect 10-bit codec, which even after 15 generations of rendering, maintains a 63dB PSNR measurement, staying well into the visually lossless region. The yellow-square plot represents the performance for DVCPRO HD, which dips well below 50dB upon multiple renders for which visual loss is discernable.

It is important when choosing a post-production format to be sure the format offers PSNR performance that remains well within the visually lossless region through multi-generation processing as does CineForm's Visually Perfect 10-bit codec.

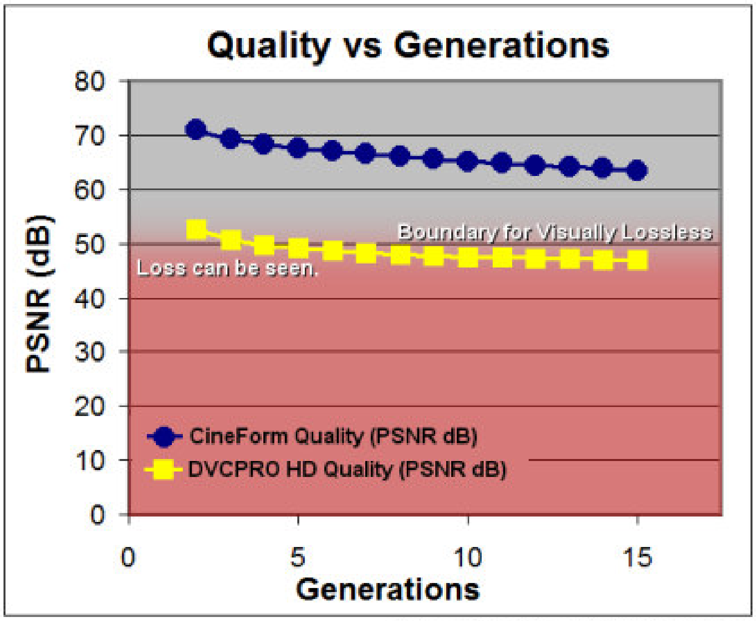

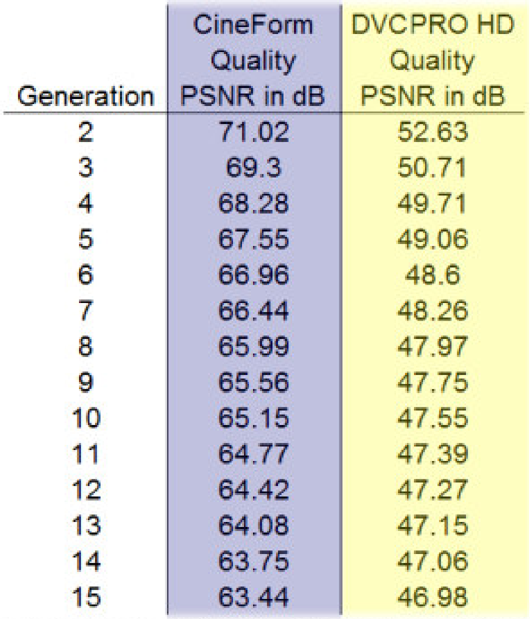

Generation	CineForm Quality PSNR in dB	DVCPRO HD Quality PSNR in dB
2	71.02	52.63
3	69.3	50.71
4	68.28	49.71
5	67.55	49.06
6	66.96	48.6
7	66.44	48.26
8	65.99	47.97
9	65.56	47.75
10	65.15	47.55
11	64.77	47.39
12	64.42	47.27
13	64.08	47.15
14	63.75	47.06
15	63.44	46.98

* The higher the PSNR (dB) number the higher the quality (lower loss between generations).

Final Word - Why does CineForm's Visually Perfect Codec Substantially Outperform DVCPRO HD?

First of all, the DVCPRO HD codec always resamples 1920x1080 source material down to 1280x1080. That is a reduction in horizontal data by 33% compared to the CineForm codec which always preserves full source resolution. Second, DVCPRO HD is only an 8-bit codec, whereas CineForm's codec operates at either 8-bit or 10-bit; these tests were performed with CineForm's 10-bit codec. Visual quality for multiple generations of rendering simply cannot stand up well using an 8-bit codec. Third, DVCPRO HD is a fixed bitrate codec, regardless of how complex a scene is. CineForm's Visually-Perfect codec is a variable bitrate codec that maintains constant quality regardless of scene complexity.

Finally, for those more technically minded, the DVCPRO HD codec compresses image data using a variation of a DCT (discrete cosine transform) algorithm. DCT compression divides images up into small (usually 8 pixel by 8 pixel) regions known as "blocks". Each block is transformed and compressed separately from all other blocks in the image. Under stressful (highly complex) images, these blocks begin to break down, yielding "block artifacts" familiar to many people. The transform structure employed by CineForm's Visually-Perfect codec uses a Wavelet transform that never breaks the image into small regions, eliminating the possibility of block artifacts that often accompany DCT compression algorithms.

In summary, CineForm's Visually-Perfect codec provides the visual performance to stand up to the rigor required throughout the post-production process.

PSNR Quality Analysis Methodology

The quality analysis comparison performed between CineForm's 10-bit Visually Perfect codec and the DVCPRO HD codec was performed at CineForm's facilities by CineForm employees. The analysis uses visual inspection, image differencing, and PSNR (peak signal-to-noise ratio) calculations as measures of image quality. PSNR analysis uses a standard mathematical model to measure an objective difference between two images. It is commonly used in the development and analysis of compression algorithms, and for comparing visual quality between different compression systems. PSNR is calculated by the following formula:

PSNR = 20*log10 (255 / RMSE) // For 8bit inputs.

where RMSE is the square root of the mean squared error for the entire image.

Pseudo code (for those who wish to perform their own PSNR analysis):

```
For (each pixel)
{
    Difference  = Pixel from Image A – Pixel from Image B
    SummedError = SummedError + Difference * Difference
}
MeanSquaredError = SummedError / Number of Pixels
RMSE = sqrtt (MeanSquaredError)
PSNR = 20*log10 (255 / RMSE)
```

For the Quality results shown in the table, PSNR was calculated on the uncompressed luma (Y) channel of the YUV signal. Luma is the best channel for showing errors the human eye is most likely to perceive. Also, by analyzing the luma channel directly we avoid any PSNR error that may occur in YUV to RGB conversion; therefore we get a more accurate measure of the codec’s quality performance.

NOTE: The Luma channel for both sources (CFHD 10-bit and DVCPRO HD 8-bit) was scaled up to a 16-bit value for analysis (i.e. the data was normalized – requiring 255 in the formulas to be 65535.) If you perform your own analysis, be weary that most off-the-shelf PSNR tools will only consider 8-bit sources.

PSNR was calculated using the first-captured image as the reference image, not the original 1920x1080 source image. The reason for this is the DVCPRO HD codec always resamples 1920x1080 frames to 1280x1080, substantially reducing (by 1/3) its spatial resolution. Consequently, using the original image as the reference for PSNR analysis would have further reduced the measured PSNR of the DVCPRO HD codec. Consequently the PSNR analysis uses the first-captured image as the reference for both codecs.

To perform each step in the multi-generation testing, a 1920x1080 uncompressed RGB frame was converted to 10-bit YUV 4:2:2 and presented to each codec. For testing the DVCPRO HD codec the data was sent across HD-SDI for compression in the AJ-HD130DCP DVCPRO HD deck. For testing CineForm's CFHD codec the uncompressed YUV source was presented to the production version of CineForm's software codec that ships with Prospect HD. This effectively simulates a DVCPRO HD camera acquiring its first image, or the CineForm codec acquiring its first image from an uncompressed camera head. By performing the PSNR analysis in this way we are testing the performance of each codec for multi-generational characteristics – i.e. we are testing how well it preserves the capture quality through a series of rendering stages. No additional compression was used in the signal processing paths of either codec.

The PSNR graph on the Quality Analysis page shows the boundary generally considered the dividing point between "visually lossless" and "visible". This boundary is subjective, as some people will see errors that others will not. In practice, it becomes very difficult to see differences for PSNR greater than 50dB between the source and the compressed image.

The difference images were calculated by comparing the source RGB image to the resulting compressed images. (Remember, the PSNR analysis used the first-captured image, not the original image). This is why both codecs show some loss in their first generation. Some small differences are due to the RGB-to-YUV conversion (common to most compressed video formats), however most of the error shown with the DVCPRO HD codec indicates the lower resolution nature of this compressed format. The difference images where created by extracting uncompressed frames from the timeline (in RGB.) These frames were loading into an image-processing tool (GIMP) where the source and resulting compressed images were "differenced" (not subtracted which can hide errors.) These difference images were lifted to a middle grey by setting the brightness to 127, a luminance region where the eye is more sensitive. In order to enhance the visibility of codec-induced errors, the contrast for all difference images was lifted about 400%.