Apple ProRes

White Paper
December 2013

Contents

Introduction

Apple ProRes is one of the most popular codecs in professional post-production. The Apple ProRes family of video codecs has made it both possible and affordable to edit full-frame, 10-bit, 4:2:2 high-definition (HD), 2K, 4K, and 5K video sources with multistream performance in Final Cut Pro X. This white paper provides in-depth information about all five members of the Apple ProRes family, including technical specifications and performance metrics.

Authorized Apple ProRes Implementations

Apple ProRes is a codec technology developed for high-quality, high-performance editing in Final Cut Pro X. Apple has licensed Apple ProRes to select companies for use in specific products and workflows.

In some instances, unauthorized codec implementations have been used in third-party software and hardware products. Using any unauthorized implementation (like the FFmpeg and derivative implementations) may lead to decoding errors, performance degradation, incompatibility, and instability.

For a list of all authorized Apple ProRes licensees and developers, and for licensing information, go to support.apple.com/kb/HT5959. If you are using or considering purchasing a product that encodes or decodes Apple ProRes, and that product is not on this list, please contact Apple at ProRes@apple.com.

Apple ProRes Family Overview

All members of the Apple ProRes family provide an unparalleled combination of multistream, real-time editing performance coupled with impressive image quality at reduced storage rates. Additionally, all five codecs preserve standard-definition (SD), HD, 2K, 4K, and 5K frame sizes at full resolution.

As a variable bit rate (VBR) codec technology, Apple ProRes uses fewer bits on simple frames that would not benefit from encoding at a higher data rate. All Apple ProRes codecs are frame-independent (or "intra-frame") codecs, meaning that each frame is encoded and decoded independently of any other frame. This technique provides the greatest editing performance and flexibility.

- **Apple ProRes 4444:** The Apple ProRes 4444 codec preserves motion image sequences originating in either 4:4:4 RGB or 4:4:4 $Y'C_BC_R$ color spaces. With a remarkably low data rate (as compared to uncompressed 4:4:4 HD), Apple ProRes 4444 supports 12-bit pixel depth with an optional, mathematically lossless alpha channel for true 4:4:4:4 support. Apple ProRes 4444 preserves visual quality at the same high level as Apple ProRes 422 (HQ), but for 4:4:4 image sources, which can carry the highest possible color detail.

  Apple ProRes 4444 is ideal for next-generation video and the most pristine computer graphics sources for cinema-quality compositing. For projects that previously required the Animation codec, Apple ProRes 4444 is a modern replacement with real-time playback support in Final Cut Pro X. Because color detail is preserved, Apple ProRes 4444 is ideal for color grading with 4:4:4 image sources.

- **Apple ProRes 422 (HQ):** Apple ProRes 422 (HQ) offers visually lossless preservation of the highest-quality professional HD video that a (single-link) HD-SDI signal can carry. For this reason, Apple ProRes 422 (HQ) has enjoyed widespread adoption across the video post-production industry. This codec supports full-width, 4:2:2 video sources at 10-bit pixel depths, while retaining its visually lossless characteristic through many generations of decoding and re-encoding.

  Apple ProRes 422 (HQ) can be used both as an intermediate codec to accelerate workflows for complex, compressed video sources and as an affordable, high-performance alternative to uncompressed 4:2:2 video.

- **Apple ProRes 422:** Apple ProRes 422 offers nearly all the benefits of its big brother, Apple ProRes 422 (HQ), but at a significantly lower data rate. It provides visually lossless coding performance for the same full-width, 10-bit, 4:2:2 sequences as Apple ProRes 422 (HQ) with even better multistream real-time editing performance.
- **Apple ProRes 422 (LT):** Like Apple ProRes 422 (HQ) and Apple ProRes 422, the Apple ProRes 422 (LT) codec supports full-width, 10-bit video sequences, but at a target data rate even lower than that of its siblings. Apple ProRes 422 (LT) weighs in at 100 Mbps or less, depending on the particular video format. Apple ProRes 422 (LT) balances incredible image quality with small file sizes, and is perfect for digital broadcast environments where storage capacity and bandwidth are often at a premium.

  Apple ProRes 422 (LT) is ideal for live multicamera and on-location productions where large amounts of footage are acquired to disk.
- **Apple ProRes 422 (Proxy):** The lowest-bandwidth member of the Apple ProRes family is Apple ProRes 422 (Proxy). This codec maintains HD data rates below 36 Mbps, yet like its higher-rate Apple ProRes 422 siblings, it supports full-frame, 10-bit, 4:2:2 video.

  Apple ProRes 422 (Proxy) is intended for draft-mode or preview uses where low data rates are required, yet full-resolution video is desired. It is the ideal format for use in offline editing workflows on notebook computers or other systems where storage is at a premium and you have a large quantity of source media.

  Today's HD workflows demand offline video formats that support native frame size and aspect ratio. Apple ProRes 422 (Proxy) supports full 1920 x 1080 and 1280 x 720 resolutions, enabling full HD resolution while editing, and accurate representation of Final Cut Pro X motion effects from the creative stages through finishing.

Properties of Digital Images

The technical properties of digital images correspond to different aspects of image quality. For example, high-resolution HD images can carry more detail than their lower-resolution SD counterparts. 10-bit images can carry finer gradations of color, thereby avoiding the banding artifacts that can occur in 8-bit images.

The role of a codec is to preserve image quality as much as possible at a particular reduced data rate, while delivering the fastest encoding and decoding speed. The Apple ProRes family supports the three key properties of digital images that contribute to image quality—*frame size, chroma sampling,* and *sample bit depth*—while offering industry-leading performance and quality at each supported data rate. In order to appreciate the benefits of the Apple ProRes family as a whole and to choose which family members to use in various post-production workflows, it is important to understand these three properties.

Frame Size (Full-Width Versus Partial-Width)

Many video camcorders encode and store video frames at less than the full HD widths of 1920 pixels or 1280 pixels, for 1080-line or 720-line HD formats, respectively. When such formats are displayed, they are upsampled horizontally to full HD widths, but they cannot carry the amount of detail possible with full-width HD formats.

All Apple ProRes family members can encode full-width HD video sources (sometimes called "full-raster" video sources) to preserve the maximum possible detail an HD signal can carry. Apple ProRes codecs can also encode partial-width HD sources if desired, thereby averting potential quality and performance degradation that results from upscaling partial-width formats prior to encoding.

Chroma Sampling

Color images require three channels of information. In computer graphics, a pixel's color is typically defined by R, G, and B values. In traditional digital video, a pixel is represented by Y', C_B, and C_R values, where Y' is the "luma" or grayscale value and C_B and C_R contain the "chroma" or color-difference information. Because the eye is less sensitive to fine chroma detail, it is possible to average together and encode fewer C_B and C_R samples with little visible quality loss for casual viewing. This technique, known as *chroma subsampling,* has been used widely to reduce the data rate of video signals. However, excessive chroma subsampling can degrade quality for compositing, color correction, and other image-processing operations. The Apple ProRes family handles today's popular chroma formats as follows:

- *4:4:4* is the highest-quality format for preserving chroma detail. In 4:4:4 image sources, there is no subsampling, or averaging, of chroma information. There are three unique samples, either Y', C_B, and C_R or R, G, and B, for every pixel location. Apple ProRes 4444 fully supports 4:4:4 image sources, from either RGB or $Y'C_BC_R$ color spaces. The fourth "4" means that Apple ProRes 4444 can also carry a unique alpha-channel sample for every pixel location. Apple ProRes 4444 is intended to support 4:4:4:4 RGB+Alpha sources exported from computer graphics applications such as Motion, as well as 4:4:4 video sources from high-end devices such as dual-link HDCAM-SR.
- *4:2:2* is considered a high-quality, professional video format in which the chroma values of $Y'C_BC_R$ images are averaged together such that there is one C_B and one C_R sample, or one "C_B/C_R chroma pair," for each Y' (luma) sample. This minimal chroma subsampling has traditionally been considered adequate for high-quality compositing and color correction, although better results can be achieved using 4:4:4 sources. 4:2:2 sources are generated by many popular higher-end video camcorder formats, including DVCPRO HD, AVC-Intra/100, and XDCAM HD422/50. All Apple ProRes 422 family members fully support the chroma resolution inherent in 4:2:2 video formats.
- *4:2:0* and *4:1:1* have the least chroma resolution of the formats mentioned here, with just one C_B/C_R chroma pair for every four luma samples. These formats are used in a variety of consumer and professional video camcorders. Depending on the quality of a camera's imaging system, 4:2:0 and 4:1:1 formats can provide excellent viewing quality. However, in compositing workflows it can be difficult to avoid visible artifacts around the edges of a composited element. HD 4:2:0 formats include HDV, XDCAM HD, and AVC-Intra/50. 4:1:1 is used in DV. All Apple ProRes 422 formats can support 4:2:0 or 4:1:1 sources if the chroma is upsampled to 4:2:2 prior to encoding.

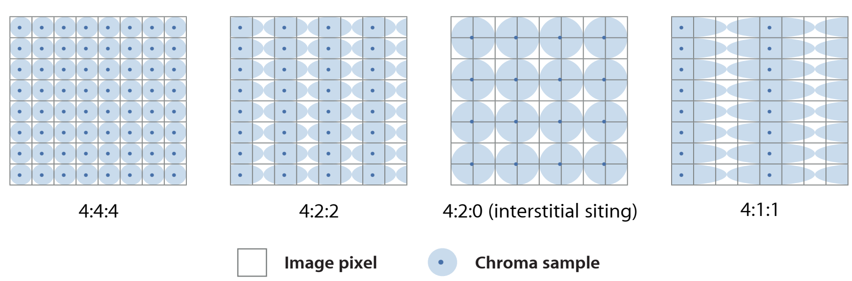

Sample Bit Depth

The number of bits used to represent each Y', C_B, or C_R (or R, G, or B) image sample determines the number of possible colors that can be represented at each pixel location. Sample bit depth also determines the smoothness of subtle color shading that can be represented across an image gradient, such as a sunset sky, without visible quantization or "banding" artifacts.

Traditionally, digital images have been limited to 8-bit samples. In recent years the number of professional devices and acquisition techniques supporting 10-bit and even 12-bit image samples has increased. 10-bit imagery is now often found in 4:2:2 video sources with professional digital (SDI, HD-SDI or even HDMI) outputs. 4:2:2 video sources rarely exceed 10 bits, but a growing number of 4:4:4 image sources claim 12-bit resolution, though with sensor-derived images the least significant one or two bits may have more noise than signal. 4:4:4 sources include high-end film scanners and film-like digital cameras and can include high-end computer graphics.

Apple ProRes 4444 supports image sources up to 12 bits and preserves alpha sample depths up to 16 bits. All Apple ProRes 422 codecs support up to 10-bit image sources, though the best 10-bit quality will be obtained with the higher bit rate family members—Apple ProRes 422 and Apple ProRes 422 (HQ). (Note: Like Apple ProRes 4444, all Apple ProRes 422 codecs can in fact accept image samples even greater than 10 bits, although such high bit depths are rarely found among 4:2:2 or 4:2:0 video sources.)

Properties of Apple ProRes Codecs

Every image or video codec can be characterized by how well it behaves in three critical dimensions: compression, quality, and complexity. *Compression* means data reduction, or how many bits are required compared to the original image. For image sequences or video streams, compression means data rate, expressed in bits/sec for transmission or bytes/hour for storage. *Quality* describes how closely a compressed image resembles the original. "Fidelity" would therefore be a more accurate term, but "quality" is the term widely used. *Complexity* relates to how many arithmetic operations must be computed to compress or decompress an image frame or sequence. For software codec implementations, the lower the complexity, the greater the number of video streams that can be decoded simultaneously in real time, resulting in higher performance within post-production applications.

Every image or video codec design must make a tradeoff among these three properties. Because codecs used within professional camcorders or for professional video editing must maintain high visual quality, the tradeoff amounts to one of data rate versus performance. For example, AVCHD camcorders can produce H.264 video streams with excellent image quality at low data rates. However, the complexity of the H.264 codec is very high, resulting in lower performance for real-time video editing with multiple video streams and effects. In comparison, Apple ProRes features excellent image quality as well as low complexity, which results in better performance for real-time video editing.

The following sections describe how the various Apple ProRes codecs behave and compare to one another in terms of these three important codec properties: *data rate, quality,* and *performance*.

Data Rate

The Apple ProRes family spans a broad range of data rates to support a variety of workflow and application purposes. This section describes how Apple ProRes data rates compare to each other and to uncompressed video. The section also illustrates how frame size and frame rate affect Apple ProRes data rates. Finally, the text includes information on the variable bit rate (VBR) nature of the Apple ProRes codec family.

The bar chart below shows how the data rate of each Apple ProRes format compares to uncompressed, full-width (1920 x 1080), 4:4:4 12-bit and 4:2:2 10-bit image sequences at 29.97 frames/sec. The chart shows that even the two highest-quality Apple ProRes formats—Apple ProRes 4444 and Apple ProRes 422 (HQ)—offer significantly lower data rates than their uncompressed counterparts.

Data Rates - Uncompressed and Apple ProRes at 1920 x 1080, 29.97 fps

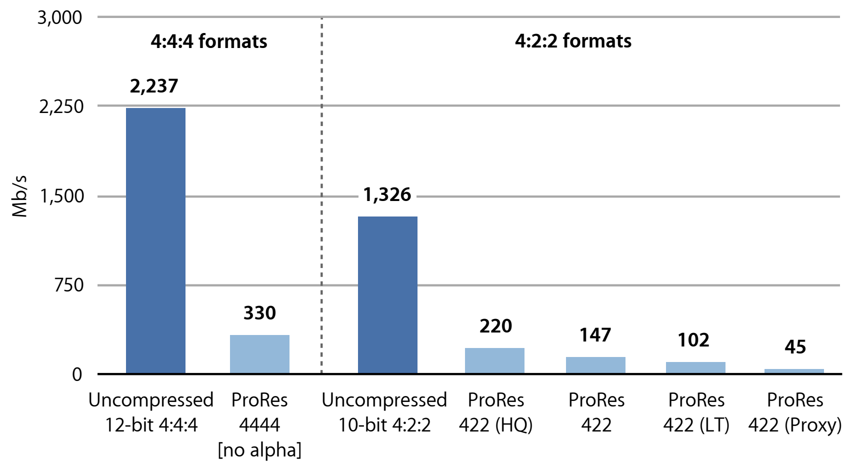

The data rates shown in the bar chart above are for "full-width" (1920 x 1080) HD frames at 29.97 frames/sec. The Apple ProRes family also supports the 720p HD format at its full width (1280 x 720). In addition to full-width HD formats, Apple ProRes codecs support three different "partial-width" HD video formats used as the recording resolutions in a number of popular HD camcorders: 1280 x 1080, 1440 x 1080, and 960 x 720.

The data rate of an Apple ProRes format is determined primarily by three key factors: Apple ProRes codec type, encoded frame size, and frame rate. The chart below shows some examples of how varying any one of these factors changes an Apple ProRes format's data rate. A table of data rates for a number of Apple ProRes formats supported for real-time editing in Final Cut Pro X can be found in the appendix.

Data Rates - Apple ProRes 422 (LT) versus Apple ProRes 422

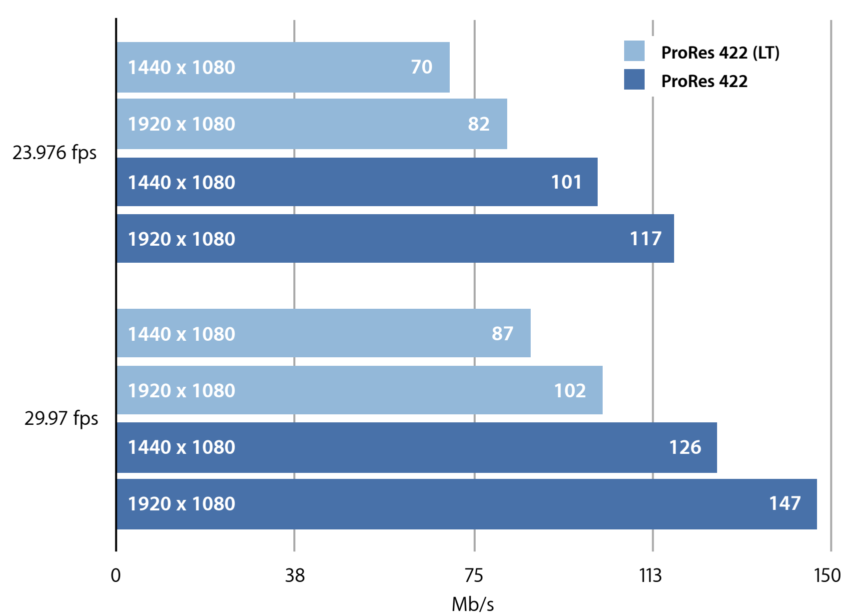

Apple ProRes is a variable bit rate (VBR) video codec. This means that the number of bits used to encode each frame within a stream is not constant, but varies from one frame to the next. For a given video frame size and a given Apple ProRes codec type, the Apple ProRes encoder aims to achieve a "target" number of bits per frame. Multiplying this number by the frames per second of the video format being encoded results in the target data rate for a specific Apple ProRes format.

Although Apple ProRes is a VBR codec, the variability is usually small. The actual data rate is usually close to the target data rate. For a given Apple ProRes format, there is also a maximum number of bits per frame that is never exceeded. This maximum is approximately 10 percent more than the target number of bits per frame. The graph below plots the actual number of bits used per frame in an example Apple ProRes video sequence.

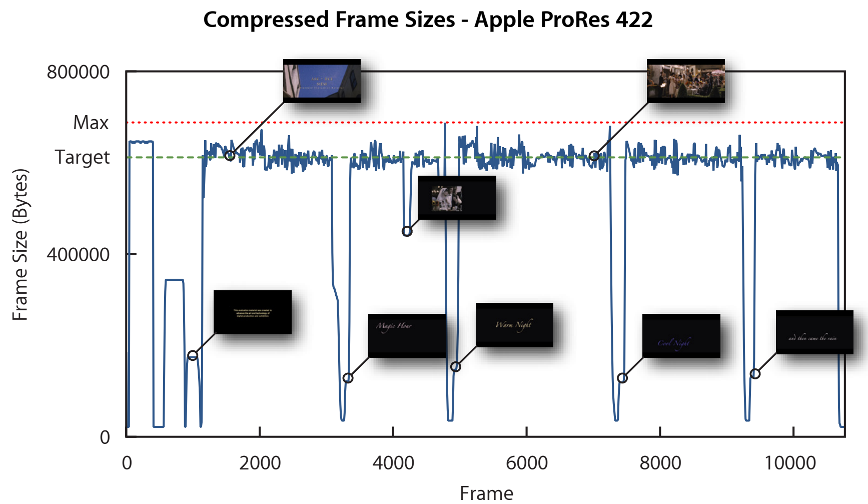

Sequence depicted is ASC/DCI Standard Evaluation Material (StEM) Mini-Movie at 1920 x 1080.

Note that for this particular sequence of over 10,000 frames, only one frame uses the maximum number of bits and most frames are clustered within a few percent of the target. However, many frames use significantly fewer bits than the target. This is because Apple ProRes encoders only add bits to a frame if doing so will produce a better match to the original image. Beyond a certain point, simple image frames, such as an all-black frame with a few words of text, incur no quality benefit if more bits are added. Apple ProRes encoders do not waste bits on any frame if adding more will not improve the fidelity.

Quality

While the ability to produce high-quality output is a key attribute of image and video codecs, it is quality preservation—or fidelity—that is the actual goal of a codec. Imagery often goes through many stages of processing prior to Apple ProRes encoding, and these stages may add visible flaws, or "artifacts," to the images. If an image sequence has visible artifacts to begin with, Apple ProRes will perfectly preserve these artifacts, which can make viewers mistakenly think such flaws are caused by the Apple ProRes codec itself. The goal of every Apple ProRes family member is to perfectly preserve the quality of the original image source, be it good or bad.

The quality-preserving capability of the various Apple ProRes codecs can be expressed in both quantitative and qualitative terms. In the field of image and video compression, the most widely used quantitative measure of image fidelity is peak signal-to-noise ratio (PSNR). PSNR is a measure of how closely a compressed image (after being decompressed) matches the original image handed to the encoder. The higher the PSNR value, the more closely the encoded image matches the original. The graph below plots the PSNR value for each image frame in a test sequence for three different codecs: Apple ProRes 422 (HQ), Avid DNxHD, and Panasonic D5.

PSNR Comparison - Apple ProRes, DNxHD, and D5

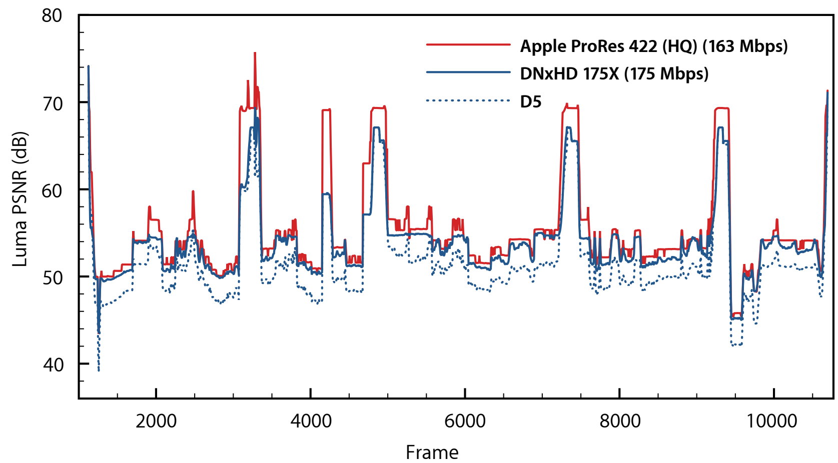

Measured using ASC/DCI Standard Evaluation Material (StEM) Mini-Movie at 1920 x 1080.

The next graph shows the same sequence plotted for each Apple ProRes 422 codec. As the graph shows, there is a difference in PSNR between one family member and the next. These differences correspond to the comparative data rates of the Apple ProRes 422 codecs. PSNR for Apple ProRes 422 (HQ) is 15-20 dB higher than that for Apple ProRes 422 (Proxy), but the Apple ProRes 422 (HQ) stream has nearly five times the data rate of the Apple ProRes 422 (Proxy) stream. The benefit of higher fidelity comes at the cost of larger file sizes, so it is important to select the Apple ProRes family member according to your workflow requirements.

PSNR Comparison - Apple ProRes 422 Family

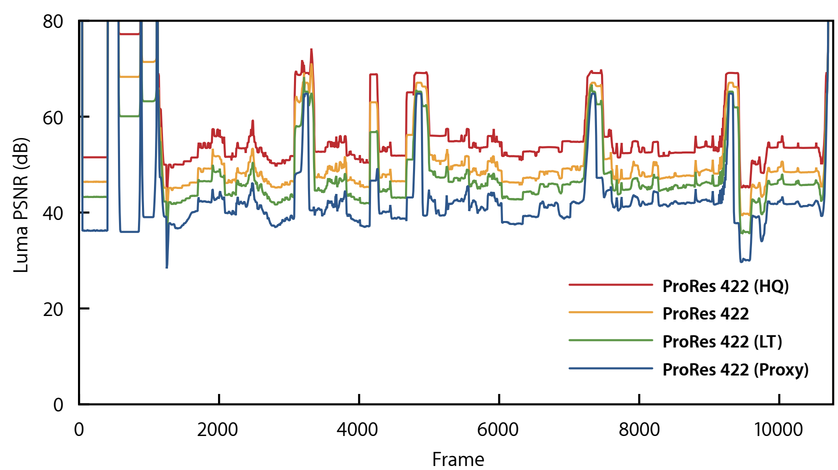

Measured using ASC/DCI Standard Evaluation Material (StEM) Mini-Movie at 1920 x 1080.

In addition to indicating visual fidelity, the difference in PSNR values also denotes headroom. For example, if you were to view the original sequence used in the graph above, and then view the Apple ProRes 422 (HQ) and Apple ProRes 422 encoded versions of the same stream, all three would look visually identical. However, the higher PSNR value for Apple ProRes 422 (HQ) indicates greater quality headroom. This increased headroom means that an image sequence can be decoded and re-encoded over multiple generations and still look visually identical to the original, as shown in the graph below:

Multigeneration PSNR

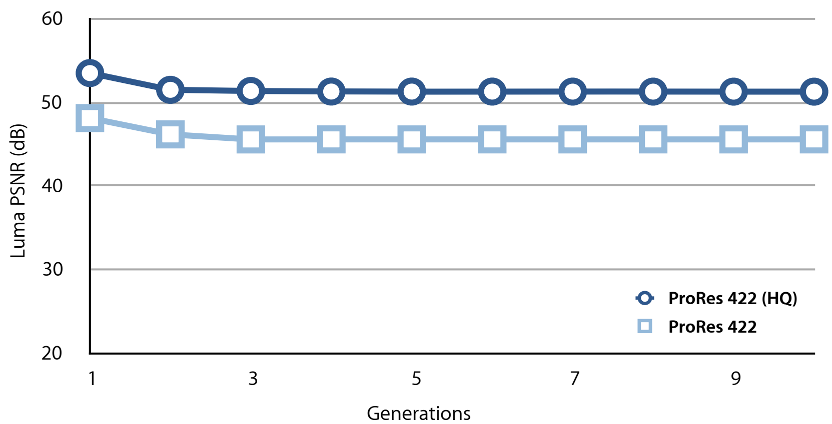

Because PSNR is not a perfect measure of compressed image fidelity—there is no particular PSNR number that can absolutely guarantee that a compressed image will have no visible difference from the original—it is useful to have some qualitative description of expected image quality for each Apple ProRes codec type. Note that in the table below, the qualitative description for Apple ProRes 4444 (without an alpha channel) is identical to that for Apple ProRes 422 (HQ). This is because Apple ProRes 4444, though its target bit rate is 50 percent higher than that of Apple ProRes 422 (HQ), uses extra bits to encode the greater number of chroma samples in 4:4:4 at the same high quality headroom ensured by Apple ProRes 422 (HQ) for 4:2:2 sources.

Apple ProRes Codec	Visible differences (1st gen.)	Quality headroom
ProRes 4444	Virtually never	Very high, excellent for multi-gen. finishing
ProRes 422 (HQ)	Virtually never	Very high, excellent for multi-gen. finishing
ProRes 422	Very rare	High, very good for most multi-gen. workflows
ProRes 422 (LT)	Rare	Good for some multi-gen. workflows
ProRes 422 (Proxy)	Subtle for high-detail images	OK, intended for first-gen. viewing and editing

Performance

The Apple ProRes family of codecs is designed for speed, and high speed of both encoding and decoding is essential to avoid workflow bottlenecks.

Fast decoding is especially critical for multistream, real-time editing in Final Cut Pro X. The Apple ProRes codec family performs exceptionally well in this regard. For each Apple ProRes codec type, the following charts show the number of full-width streams that can be edited simultaneously in real time on a Mac Pro and a MacBook Pro computer. (In each chart, the number of possible 4096 x 2160 4K streams is shown in light blue. The number of possible 1920 x 1080 HD streams is shown in dark blue.) In practice of course, you may not often need to edit five, six, or more streams simultaneously, but these charts give an idea of how much processing time will be available for real-time titling, effects, and so on, when just one, two, or three streams are being used.

Mac Pro – Final Cut Pro X Multicam Streams*

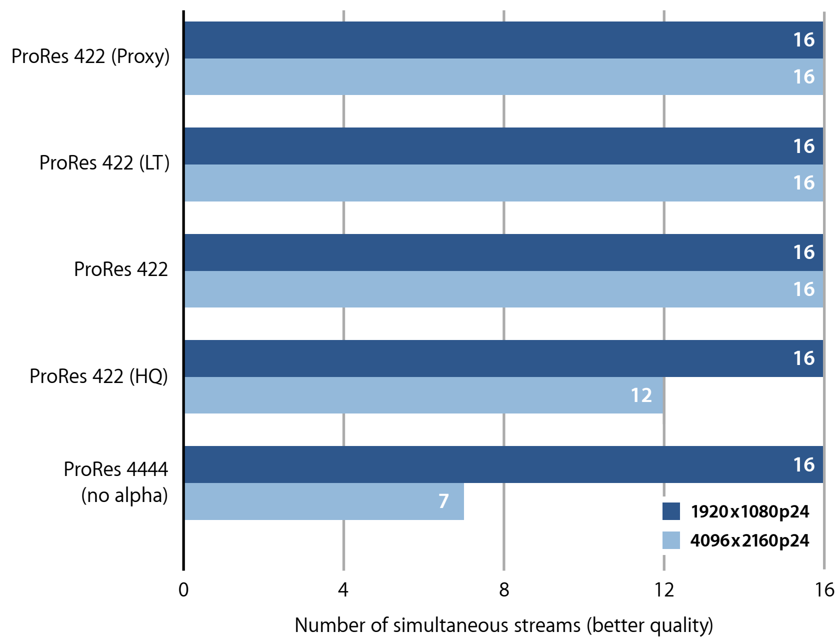

*The Final Cut Pro X Multicam feature allows you to view up to 16 angles simultaneously while switching or cutting angles in real time.

Testing conducted by Apple in September 2013 using a prerelease version of OS X Mavericks v10.9 and a prerelease 2.7GHz 12-core Intel Xeon-based Mac Pro configured with 250GB of flash storage and 16GB of RAM. Testing conducted with a prerelease version of Final Cut Pro X using 1-minute, 26-second 1920x1080p24 and 4096x2160p24 ASC-DCI Standard Evaluation Material Multicam clips for each content type. Mac Pro continuously monitors system thermal and power conditions, and may adjust processor speed as needed to maintain optimal system operation. Performance may vary depending on system configuration and content.

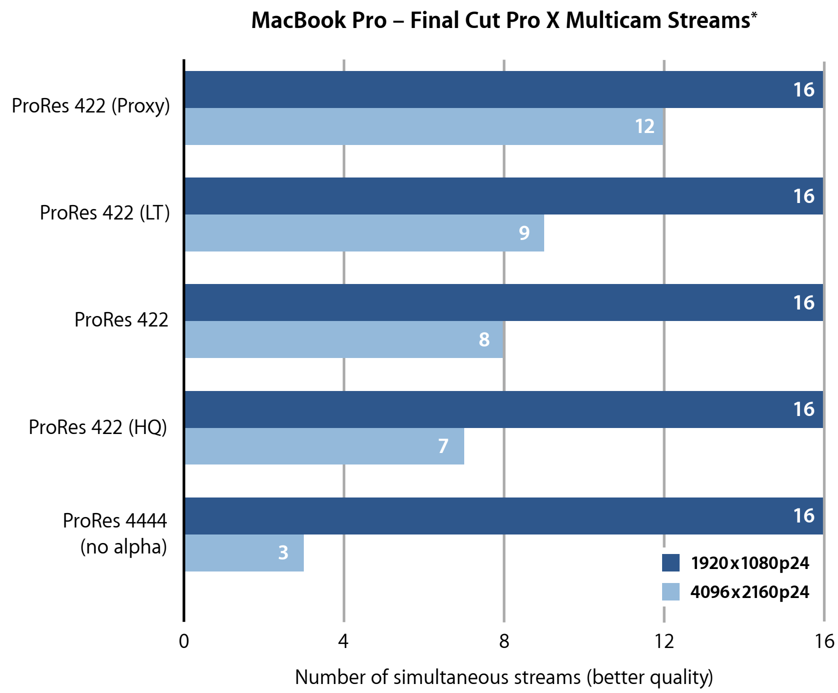

*The Final Cut Pro X Multicam feature allows you to view up to 16 angles simultaneously while switching or cutting angles in real time.

Testing conducted by Apple in September 2013 using a prerelease version of OS X Mavericks v10.9 and a prerelease 2.6GHz quad-core Intel Core i7-based 15-inch MacBook Pro with Retina display configured with 250GB of flash storage and 16GB of RAM. Testing conducted with a prerelease version of Final Cut Pro X using 1-minute, 26-second 1920x1080p24 and 4096x2160p24 ASC-DCI Standard Evaluation Material Multicam clips for each content type. MacBook Pro continuously monitors system thermal and power conditions, and may adjust processor speed as needed to maintain optimal system operation. Performance may vary depending on system configuration and content.

Today's Mac notebook and desktop machines rely on multicore processing, so the speed of a fast editing decoder must scale up—meaning that decoding time per frame should decrease—as the number of processing cores increases. Many industry codec implementations "hit the wall" and do not realize further performance gains as more processors are added, but Apple ProRes codecs continue to get faster as more cores are added, as the following chart shows.

Multiprocessor Scaling – Apple ProRes 422 (HQ) at 1920 x 1080

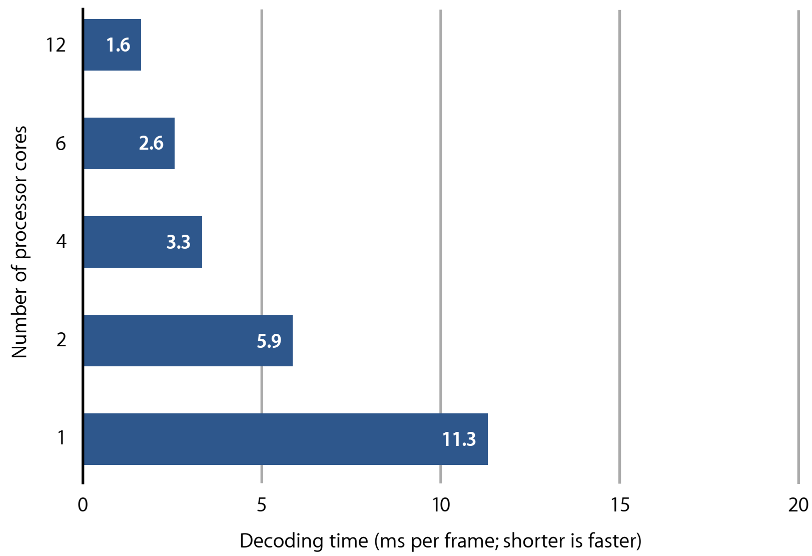

Testing conducted by Apple in September 2013 using a prerelease version of OS X Mavericks v10.9 and a prerelease Mac Pro 2.7GHz 12-core Intel Xeon with 16GB of RAM. Performance may vary depending on system configuration, content, and performance measurement tool use.

Apple ProRes decoders are designed to work especially well as high-quality, high-performance editing codecs for Final Cut Pro X. Not only are they fast for decoding video at full frame size and quality, but they are even faster at decoding frames at "half-size" frame (1/2 height and 1/2 width). Especially for high-resolution formats like HD and 2K, half-size images provide plenty of onscreen detail for making editing decisions.

The chart below shows that half-size decoding is substantially faster than already-fast full-size decoding, especially for the higher-quality Apple ProRes codecs. The faster decoding speed means more CPU time is available for decoding more streams or more real-time effects.

Reduced Resolution Decoding Speed at 1920 x 1080

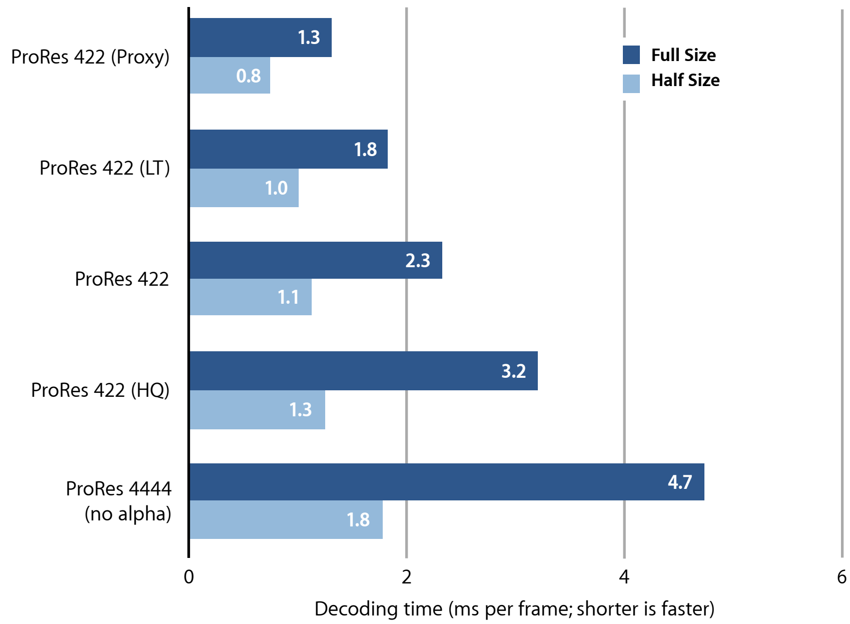

Testing conducted by Apple in September 2013 using a prerelease version of OS X Mavericks v10.9 and a prerelease 2.6GHz quad-core Intel Core i7-based 15-inch MacBook Pro with Retina display configured with 250GB of flash storage and 16GB of RAM. MacBook Pro continuously monitors system thermal and power conditions, and may adjust processor speed as needed to maintain optimal system operation. Performance may vary depending on system configuration, content, and performance measurement tool use.

Although fast decoding speed is the primary factor in real-time editing performance, fast encoding speed is also important for key steps in post-production workflows. Like Apple ProRes decoders, the Apple ProRes family of encoders have all been built as efficient software implementations, and fast encoding is achieved through efficient use of multicore processors. Fast encoding speed is essential for some steps and important in virtually all others.

For real-time capture and Apple ProRes encoding of baseband video signals (either analog or digital SD or HD signal sources), Apple ProRes software encoders must be fast enough to keep up with the incoming real-time video frames. An appropriate video capture card must be used for this purpose, but otherwise no specialized encoding hardware is required to achieve real-time capture of baseband video to Apple ProRes formats.

For file-based transcoding of video files that have been encoded with other (non-Apple ProRes) video codecs, transcoding to Apple ProRes entails both decoding of the starting technique and re-encoding to Apple ProRes. The minimum total transcoding time will therefore be the sum of the time required to decode the file and the time required to re-encode it to Apple ProRes. For certain video codec formats known to be highly complex and therefore relatively slow to decode, such as JPEG-2000 and the REDCODE® RAW (R3D) native codec format, the overall transcoding time will be dominated by the decoding time. Still, fast Apple ProRes encoding helps make the total transcoding time faster.

Fast encoding and decoding also benefits rendering and exporting. Rendering effects, as part of a creative process or the final step before output, is basically a decode of the source media and a re-encode to the chosen final output format. During the rendering process, all of the decoding, blending, and compositing steps must be precomputed before encoding to the compressed format defined in your Final Cut Pro X project. Although you can choose any Apple ProRes codec as a rendering format—from Apple ProRes 422 (LT) to Apple ProRes 4444—and change it at any time during post-production, Final Cut Pro X defaults to rendering in Apple ProRes 422.

When rendering to Apple ProRes, the total rendering time is determined by the speed of both the decoding and encoding steps, which can be significantly quicker compared to other, more complex and slower codecs. The speed advantage of Apple ProRes is also beneficial when exporting a file at the end of a project. If you need to deliver to the web, DVD, or Blu-ray disc, you can speed up the export process by choosing to edit in Apple ProRes instead of other professional formats, including uncompressed.

Apple ProRes 4444 Alpha Channel Support

In addition to supporting $Y'C_BC_R$ or RGB 4:4:4 pixel data, the Apple ProRes 4444 codec type supports an optional alpha channel. The sampling nomenclature for such $Y'C_BC_RA$ or RGBA images is 4:4:4:4, to indicate that for each pixel location, there is an alpha—or A—value in addition to the three $Y'C_BC_R$ or RGB values. An alpha value specifies the proportion of its associated RGB or $Y'C_BC_R$ pixel that should be blended with the pixel at the corresponding location of a background image, creating the effect of varying transparency for use in compositing workflows. Unlike $Y'C_BC_R$ or RGB pixel values, alpha values do not represent samples of a real-world image, or even samples of a computer-generated image, both of which are intended for human viewing.

Alpha values are essentially numeric data that specify how to blend, or composite, a foreground image into a background image. For this reason, Apple ProRes 4444 encodes alpha values exactly rather than approximately. This kind of exact encoding is called "lossless" (or sometimes "mathematically lossless") compression. It uses different encoding techniques from those used by the Apple ProRes codec family for RGB or $Y'C_BC_R$ pixel values, where approximate encoding is acceptable as long as differences from the original are not visible to the viewer and do not affect processing. The Apple ProRes 4444 codec losslessly encodes alpha channel values of any bit depth up to and including 16 bits.

In summary, the Apple ProRes 4444 codec can be considered "visually lossless" for encoding the $Y'C_BC_R$ or RGB pixel values intended for viewing, but "mathematically lossless" for encoding the alpha values that specify compositing. As a result, the degree of quality or fidelity is never a question for Apple ProRes 4444 alpha channels because the decoded data always matches the original perfectly.

With any kind of lossless compression, the data rate varies according to the amount of image detail being encoded. This is true of Apple ProRes 4444 lossless alpha channel compression as well. However, in practice alpha channels typically contain just the information related to object outlines, so the optional alpha channel typically adds just a few percent to the overall Apple ProRes 4444 data rate. For this reason, the presence of an alpha channel in an Apple ProRes 4444 stream typically reduces decoding and encoding performance by only about 10 percent or less.

Appendix

Target Data Rates

Dimensions	Frame Rate	ProRes 422 (Proxy)		ProRes 422 (LT)		ProRes 422		ProRes 422 (HQ)		ProRes 4444 (no alpha)	
		Mb/s	GB/hr	Mb/s	GB/hr	Mb/s	GB/hr	Mb/s	GB/hr	Mb/s	GB/hr
720 x 486	24p	10	4	23	10	34	15	50	23	75	34
	60i, 30p	12	5	29	13	42	19	63	28	94	42
720 x 576	50i, 25p	12	6	28	13	41	18	61	28	92	41
960 x 720	24p	15	7	35	16	50	23	75	34	113	51
	25p	16	7	36	16	52	24	79	35	118	53
	30p	19	9	44	20	63	28	94	42	141	64
	50p	32	14	73	33	105	47	157	71	236	106
	60p	38	17	87	39	126	57	189	85	283	127
1280 x 720	24p	18	8	41	18	59	26	88	40	132	59
	25p	19	9	42	19	61	28	92	41	138	62
	30p	23	10	51	23	73	33	110	49	165	74
	50p	38	17	84	38	122	55	184	83	275	124
	60p	45	20	101	46	147	66	220	99	330	148
1280 x 1080	24p	31	14	70	31	101	45	151	68	226	102
	60i, 30p	38	17	87	39	126	57	189	85	283	127
1440 x 1080	24p	31	14	70	31	101	45	151	68	226	102
	50i, 25p	32	14	73	33	105	47	157	71	236	106
	60i, 30p	38	17	87	39	126	57	189	85	283	127
1920 x 1080	24p	36	16	82	37	117	53	176	79	264	119
	50i, 25p	38	17	85	38	122	55	184	83	275	124
	60i, 30p	45	20	102	46	147	66	220	99	330	148
	50p	76	34	170	77	245	110	367	165	551	248
	60p	91	41	204	92	293	132	440	198	660	297

Target Data Rates (continued)

Dimensions	Frame Rate	ProRes 422 (Proxy)		ProRes 422 (LT)		ProRes 422		ProRes 422 (HQ)		ProRes 4444 (no alpha)	
		Mb/s	**GB/hr**	**Mb/s**	**GB/hr**	**Mb/s**	**GB/hr**	**Mb/s**	**GB/hr**	**Mb/s**	**GB/hr**
2048 x 1080	24p	41	19	93	42	134	60	201	91	302	136
	25p	43	19	97	44	140	63	210	94	315	142
	30p	52	23	116	52	168	75	251	113	377	170
	50p	86	39	194	87	280	126	419	189	629	283
	60p	103	46	232	104	335	151	503	226	754	339
2048 x 1556	24p	56	25	126	57	181	81	272	122	407	183
	25p	58	26	131	59	189	85	283	127	425	191
	30p	70	31	157	71	226	102	340	153	509	339
	50p	117	52	262	118	377	170	567	255	850	382
	60p	140	63	314	141	452	203	679	306	1019	458
3840 x 2160	24p	145	65	328	148	471	212	707	318	1061	477
	25p	151	68	342	154	492	221	737	332	1106	498
	30p	182	82	410	185	589	265	884	398	1326	597
	50p	303	136	684	308	983	442	1475	664	2212	995
	60p	363	163	821	369	1178	530	1768	795	2652	1193
4096 x 2160	24p	155	70	350	157	503	226	754	339	1131	509
	25p	162	73	365	164	524	236	786	354	1180	531
	30p	194	87	437	197	629	283	943	424	1414	636
	50p	323	145	730	328	1049	472	1573	708	2359	1062
	60p	388	174	875	394	1257	566	1886	848	2828	1273
5120 x 2160	24p	194	87	437	197	629	283	943	424	1414	636
	25p	202	91	456	205	655	295	983	442	1475	664
	30p	243	109	546	246	786	354	1178	530	1768	795
	50p	405	182	912	410	1311	590	1966	885	2949	1327
	60p	485	218	1093	492	1571	707	2357	1061	3535	1591

Glossary

alpha channel: An additional channel of information that may optionally be included with RGB and $Y'C_BC_R$ images. If included with an RGB image, for each R, G, and B value that defines a pixel, there is an "A" value that specifies how the RGB pixel should be blended with a background image. Typically, one extreme value of A indicates 100% transparency and the other extreme value indicates 100% opacity. Values in between the extremes indicate the degree of opacity.

Apple ProRes format: An Apple ProRes-encoded bitstream, typically in the form of a .mov file, for which the Apple ProRes codec type and video format are specified. For example, an "Apple ProRes 422 (HQ) 1920 x 1080i 29.97 format."

codec: Abbreviation for *compressor/decompressor*. A general term referring to both encoder and decoder.

decoder: An algorithm or processing system that takes a compressed bitstream as input and delivers a sequence of images or video frames as output. For Apple ProRes, this term refers to a QuickTime decompressor component that converts an Apple ProRes-encoded .mov file to a sequence of images, for further processing or display.

encoder: An algorithm or processing system that takes uncompressed images as input and delivers a compressed bitstream as output. For Apple ProRes, this term refers to a QuickTime compressor component that generates an Apple ProRes-encoded .mov file.

image sequence: An ordered set of image frames that, when displayed at a specified frame rate, is perceived by the viewer as a real-time motion image sequence. If not referred to as "video," an image sequence is often a set of RGB images (with an optional alpha channel), such as the DPX, TIFF, and OpenEXR file formats.

lossless: A type of codec for which putting an image frame through encoding followed by decoding results in an image that is mathematically guaranteed to have exactly the same pixel values as the original.

video: An image sequence for which the image frames typically use the $Y'C_BC_R$ color space and subsampled chroma channels, usually with one of the following patterns: 4:2:2, 4:2:0, or 4:1:1.

video format: A video sequence for which the frame height, frame width, and frame rate are all specified. For example, a "1920 x 1080i 29.97 video format."

visually lossless: A type of codec for which putting an image frame through encoding followed by decoding results in an image that is not mathematically lossless, but is visually indistinguishable from the original when viewed alongside the original on identical displays.

December 2013