

Compression Primer

Introduction

This primer is for anyone learning how to prepare high-quality compressed video files. It is for web and video professionals who are new to compression, and for people who want to improve their skills. Compression is often performed at the last minute, so feel free to skip to the section covering the format and codecs you care about.

TABLE OF CONTENTS

| | |
|----|---|
| 1 | Introduction |
| 2 | Getting started with compression |
| 3 | Understanding the fundamentals of compression |
| 4 | Types of compression |
| 5 | Frame types and compression |
| 6 | Color modes |
| 6 | Color depths |
| 7 | Color sampling |
| 8 | Audio compression |
| 9 | Compression efficiency |
| 9 | Authoring video for future compression |
| 12 | Capturing video: choosing formats and connections |
| 13 | Choosing the right capture codec |
| 14 | Exporting video |
| 15 | Understanding preprocessing |
| 17 | Cropping |
| 17 | Scaling |
| 18 | Noise reduction |
| 19 | Image processing |
| 20 | Audio processing |
| 21 | Encoding |
| 25 | Delivery methods for video |
| 27 | Choosing a compression format |
| 28 | Exporting to Windows Media format |
| 29 | Windows Media players |
| 30 | Windows Media encoding modes |
| 32 | Windows Media video codecs |
| 33 | Windows Media audio codecs |
| 34 | Exporting to RealMedia format |
| 34 | RealMedia video codecs |
| 35 | RealMedia audio codecs |
| 36 | Exporting to QuickTime format |
| 37 | QuickTime video codecs |
| 42 | QuickTime audio codecs |

Getting started with compression

The best starting point for learning about how to achieve great results from compression is to understand why and how it's done. While the mathematical bases of compression are very complex, the basic principles are quite approachable. The main reason you need to compress video is because uncompressed video is far too large to transport effectively.

| Bitstream type | Data rate (Kilobits per second) | Time on one 700 MB CD-ROM |
|---|---------------------------------|-----------------------------------|
| Uncompressed high definition (1920 x 1080 29.97 fps) | 745750 | 7.5 seconds |
| Uncompressed standard definition (720 x 486 29.97 fps) | 167794 | 33 seconds |
| DV25 (miniDV, DVCAM, DVCPRO) | 25000 | 3 min, 44 seconds |
| Typical DVD | 5000 | 18 min, 40 seconds |
| Video CD | 1167 | 80 minutes |
| Broadband web video | 100-2000 | (500 Kbps) 3 hours, 8 minutes |
| Modem web video | 18-48 | (32 Kbps) 48 hours, 37 minutes |

The ratios between data rates used for authoring video and for distributing compressed video are huge. For example, converting from digital video (DV) to DVD typically entails a 5:1 compression, and converting an uncompressed source to a modem-compatible video uses about 9000:1 compression. That leaves only about 0.1% of the original data left. It's almost amazing that video on the web looks and sounds as good as it does!

In addition, different types of compressed video vary in bit rate. For example, in web video, there is about a 100:1 ratio between the low end of modem rates to the high end of broadband rates.

Compression and DVD formats

More Americans watch compressed video on DVD than on any other media. (Digital broadcasts over cable and satellite are close runners up.)

Modern DVD-R burners can put 4.7 GB of data on a disc (called a DVD-5). Manufactured DVDs can use a dual-layer format that can hold 8.54 GB (called DVD-9) of data, which burners do not support yet. Even with a DVD-5, it's possible to put a high-quality two-hour movie on a disc.

The kind of DVD that plays in a DVD player is more accurately known as DVD-Video. There is also DVD-ROM, which is a computer-formatted disc that requires a computer to play it—in effect, it's a large capacity CD-ROM. There are also hybrid DVD-Video/DVD-ROM discs, which play in normal DVD players, but also include information only available for computer playback.

DVD-Video discs are mastered in either NTSC (the format used in the U.S. and Japan) or PAL (the format used in Europe and Asia). A set-top player (such as your home DVD player) for one format typically won't play discs in the other format; however, computers can typically play discs of both types. DVDs can also be region coded, so they only play in DVD players or software intended for a certain region. Region coding is an optional feature.

- 45 Exporting to MPEG-1 format
- 46 MPEG-1 video and audio codec
- 46 Exporting to DVD and MPEG-2 formats
- 47 MPEG-2 audio codecs
- 48 Exporting to MPEG-4 format
- 50 MPEG-4 audio codecs
- 51 Exporting to AVI format
- 51 AVI video codecs
- 52 AVI audio codecs

Compression efficiency

The top two rows of uncompressed video are substantially larger in size and shorter in duration than the compressed video in the last five rows.

Compression and VideoCD formats

VideoCD is an older disc-based format. It has a much lower resolution and data rate than DVD, and it provides lower quality. However, it uses standard CD media, so VideoCD discs can play in a vast variety of devices, including most personal computers. A manufactured VideoCD can play in any DVD player.

Like DVD, VideoCD discs can be either NTSC or PAL. VideoCD is very popular in PAL countries, especially in Asia. Also like DVD, software players support both formats, and many set-top players only support one.

Compression and CD-ROM formats

Starting in the early 1990s, CD-ROM video was a popular medium. While the web has largely replaced CD-ROM as a medium for interactive content, people still use CD-ROM video for many applications, like computer games and kiosks. Compared to web video, CD-ROM offers much higher bandwidth than broadband connections and is capable of DVD quality on fast computers.

Compression and Internet formats

Video on the web is the hottest area for ongoing development of compressed video technologies. The main reason is that the bandwidth available on the web is much less than on discs, so it is critical to make every bit count. Another factor is that it's possible to distribute updated player software to computers, which isn't possible with consumer electronic devices, like a DVD.

Web video can use either the real-time streaming mode or progressive download mode. Real-time streaming files play in real time but require specialized streaming server software and are limited in quality by the available bandwidth. Progressive download files come off standard web servers and may take a long time to download, but they may provide higher and more reliable video and audio quality. Understanding the differences between these modes, and the strengths and weaknesses of each, is a major theme of this Primer.

Compression and mobile devices

Mobile devices, such as personal digital assistants (PDAs) and modern cell phones create a new area for compressed video. These devices have small screens and slow processors, even compared to a laptop, but the ability to play video back on a device with a weight measured in ounces (not pounds) can have a lot of value for some markets.

Most formats for compressed video have players for some mobile devices, including MPEG-4, Microsoft® Windows Media, and RealMedia. Dedicated mobile formats, like Kinoma, are also available.

In many ways, video on mobile devices, especially cell phones, is reminiscent of the early days of video on desktop computers. The technology is still evolving, but it's amazing to see it work. Today, there aren't any universal formats or even settings for mobile devices. Instead, content has to be tuned to a small class of devices or even a particular model. But when that works, the results can be great. Mobile video is rapidly maturing and should become easier to compress at a higher quality over the next few years.

Understanding the fundamentals of compression

Compression tools and formats hide the mathematical complexity of compression from the user, but to achieve the best results, you need to understand the fundamentals of why and how compression tools and formats operate.

How we see and hear

Lacking video adapters and audio ports in our heads, we don't directly perceive video or audio, compressed or not. Instead, we use our senses of sight and hearing, which are powerful, but have limitations that make compression possible.

The eye is often described as a camera, but the human visual system includes both the eye and the brain and is massively more complicated and intertwined than a simple CCD-to-tape device. The whole sensory system is tuned according to the evolutionary environment of our ancestors. Thus, humans are skilled at noticing motion, like a saber-toothed tiger jumping out of a jungle, or the difference between ripe and unripe fruit by color. But we aren't skilled at seeing high motion and detail at the same time, or high motion and color at the same time. In the same way, hearing is powerful, but it has its limitations. We may easily hear a sound in a quiet environment but not hear it in a loud one. Human ears are best attuned to the frequencies of human speech and are less aware of higher or lower tones. Compression technologies take advantage of the strengths and weaknesses of human senses to spend the available bits on the information that humans perceive more acutely.

How compression works

Fundamentally, compression works by summarizing the sequences of images and sounds as efficiently as possible. Files encoded for delivery don't provide an identical representation of the source back to the user—the data rates would be far too high. Instead, the encoded files describe the important details of the content as tersely as possible while still providing a reproduction that, to the human sensory system, provides the experience of the original as accurately as possible.

A matched pair of a compressor and a decompressor, better known as a *codec*, performs the compression. It's important that the pair be matched because the decoder (decompressor) needs to understand the encoder's (compressor's) summary of the data.

Types of compression

Video codecs can use spatial compression, temporal compression, or a combination of both.

Spatial compression

Spatial compression affects only a single frame at a time, also called *intraframe compression*. Most authoring codecs only use spatial compression, so there aren't any interdependencies among frames. This type of compression makes random access and editing easy.

Spatial compression finds redundancy within a frame—either in areas of flat or similar colors or areas of a frame that look like other areas of the frame.

Many codecs are based around a technique called Discrete Cosine Transformation (DCT). These codecs deal very well with smooth gradations in an image, but they can have trouble encoding random details or sharp edges at lower bit rates. DCT, or similar techniques trying to encode overly complex images at too low of a data rate, cause two of the most common kinds of artifacts: ringing and blocking. One common area of trouble is small text with sharp edges. Some codecs use very different base technologies, like wavelet, and therefore deal better with sharp edges.

Temporal compression

Temporal compression adds the ability to use other frames as a reference for the current frame. In most video, a majority of the frames are similar enough to the frame previous to them that temporal compression is applicable.

Temporal compression gets a lot of research attention and has the most room for future improvements.

Frame types and compression

You can reveal many fundamental concepts of compression by looking at the kinds of frames in a file. Depending on the codec and mode, one, two, or three kinds of frames may be in a file.

About Keyframes (I-frames)

A Keyframe, or I-frame (Intraframe), is a self-contained image. An I-frame isn't based on any other frames in the file. The first frame of a movie is always an I-frame.

For delivery codecs, most frames aren't I-frames because they don't offer as much compression efficiency as other frame types, which are able to reuse elements that are repeated in multiple frames.

I-frames are important for random access because they're the only frames that can be decoded all by themselves. Skipping to a frame that isn't an I-frame means that the frames the I-frame is based on must also be decoded, and the frames those P-frames refer to must be decoded, and so on. As the frequency of the I-frames increases, the number of frames that needs to be decoded decreases.

Most authoring codecs, like Motion JPEG and DV, are I-frame only. Editing is easy with I-frames because you can reassemble frames without regard for the other frames they might be based on.

About Delta frames (P-frames)

A Delta frame, or P-frame (Predictive frame) in MPEG, isn't self-contained like a keyframe. Instead, the frame is based on the frame before it, and only changes from the previous frame are encoded. So, if two frames are mathematically identical, and the second one is a P-frame, it will require only a few bits to encode. But if the frames are completely different, a P-frame will be as large as an I-frame (and most codecs automatically make the frame an I-frame). For typical content, a P-frame might be only 10-20% the size of the same frame at the same quality as an I-frame.

Codecs vary in how well they encode different kinds of changes among frames. For example, fades up and down are very difficult for some codecs, but some newer codecs have code built in to handle fades.

About Bidirectional frames (B-frames)

A Bidirectional, or B-frame, is mainly used in MPEG formats. A B-frame can use image elements from the frames on either side of it, which improves compression efficiency. However, B-frames cannot use image elements from other B-frames. This restriction avoids frames being based on frames that are based on the first frame—an impossible recursive loop.

display. Some codecs supported even lower modes, like 4 or 16 colors, providing a very simple compression scheme by reducing the number of bits required to store the image. The 8-bit color depth doesn't work well for video source because even the simplest scenes use more than 256 colors. The only significant format that still requires 8-bit is GIF.

- **5 bits per channel** (often called *Thousands of Colors* or *High Color*) After the 8-bit era, the most common consumer bit depth was 5 bits per channel. On Macintosh computers, this was exactly 5 bits per red, green, and blue, with one bit reserved for an alpha channel. Microsoft Windows® gave an extra bit to green, which improved luma quality, but left no room for an alpha channel. The 5-bit color mode was used for some ancient video-compression codecs. Today, it is mainly used for screen recording applications where it may provide better compression and performance than 8 bits per channel without any obvious degradation.
- **8 bits per channel** (also called *Millions* or *True Color*) All modern video-compression codecs use this mode. It is sometimes called 24-bit or 32-bit mode. The 32-bit mode is still 8 bits per channel but with an added alpha channel.
- **10 bits per channel** This mode is used in high-end Y'CbCr authoring systems, like Digital Betacam, D1, and some analog capture cards. It provides 1024 gradations instead of 256, which reduces banding, especially with multiple generations of effects. It isn't used in any delivery formats, although some authoring codecs support it. The 10-bit Y'CbCr mode provides four times more accuracy than 8 bit when converting to or from 16-bit-per-channel RGB in After Effects.
- **16 bits per channel** Some high-end applications use this mode to preserve more color detail. This mode is especially useful when doing multiple generations or when converting from and then back to Y'CbCr. After Effects 6 supports 16-bit-per-channel processing for RGB. Video that is saved at 16 bits per channel is never directly delivered to users.

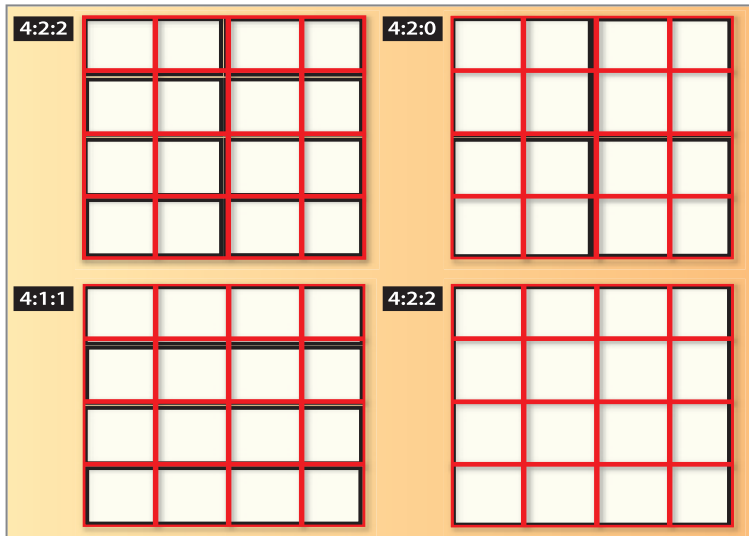
Color sampling

Y'CbCr formats support color subsampling. Color sampling records fewer pixels of chroma information than luma information. The terminology for color sampling is confusing because it uses an X:Y:Z format. The first value is the number of luma samples that are described—this is almost always four. The second value is the number of chroma samples per luma sample on the first, third, fifth, and succeeding odd lines. So if this value is two, there are two chroma samples per four luma samples horizontally. The last value is the number of chroma samples per luma sample on the second, fourth, sixth, and succeeding even lines. If this number is zero, the chroma sample from the first line is used.

- **4:4:4** Chroma is sampled every pixel with no subsampling. Y'CbCr delivery codecs don't use this format because it uses a lot of bits and doesn't provide any additional apparent quality. In authoring applications, it is sometimes internal. RGB is always 4:4:4.
- **4:2:2** This format is used for high-end video authoring, in formats like DV50, Digital Betacam, D1, and Motion JPEG. A few delivery codecs like H.263 use it as well.
- **4:2:0** The vast majority of delivery codecs use this format. There is one chroma sample per 2 x 2 block of pixels. This chroma sample looks nearly as good as 4:4:4, but at half the bits per frame. It is also used for PAL DV (but not PAL DVCPRO).
- **4:1:1** NTSC DV and PAL DVCPRO use this format. It contains the same number of bits per frame as 4:2:0 but with one sample per 4 x 1 block of pixels. This format is better for inter-

laced content because each field gets its own sample; however, it can result in visible chroma artifacts in areas of high saturation, which can be a problem with bluescreen or motion graphics source.

- **YUV-9** This format has one chroma sample per 4 x 4 block of pixels. This format results in noticeably degraded quality, so it isn't used in modern codecs. Legacy codecs that use it include the Indeo family and the original version of Sorenson Video. (Sorenson Video 3 uses 4:2:0.)



Color subsampling
The most common types of color subsampling.

Audio compression

While audio may use only a small fraction of the bits in a compressed file, it is responsible for half of the experience. We shouldn't neglect its importance.

Sampling rate

The most fundamental factor in audio is sampling rate, which is measured in Hertz (Hz), or cycles per second. Audio CDs use 44100 Hz, or 44.1 KHz, which provides excellent reproduction.

| Sample rate | Used | Sufficient for |
|-------------|------------------------------|--------------------------|
| 8 KHz | Phone system | Speech |
| 11 KHz | Old multimedia standard | Better speech |
| 22.050 KHz | Also old multimedia standard | Recognizable music |
| 32 KHz | Common in broadband video | Minimum for decent music |
| 44.1 KHz | Audio CD | Music |
| 48 KHz | DVD, DAT | Music |
| 96 KHz | High-end audio recording | High fidelity mastering |

Sample rate chart
Typical sample rates for common uses.

Bit depth

Bit depth determines the precision of audio volume at any given sample—it has a massive effect on audio quality.

- **8 bit** Originally, multimedia audio used 8 bit, which means there was a measurement of volume from 0-256 for each sample. This bit depth wasn't nearly enough for decent quality, which led to the harsh quality of early CD-ROM titles. There is no reason to distribute audio as 8 bit anymore. Modern 16-bit codecs provide better quality at a smaller file size than 8-bit codecs were ever able to.

-
-
- **16 bit** This bit depth is the current standard for audio distribution. Most modern codecs are 16 bit, and audio CDs are 16 bit as well.
 - **20 bit** This bit depth is used for professional audio recording to give a better dynamic range than 16 bit. Codecs that support 24-bit audio (described later) can effectively take advantage of 20-bit sources.
 - **24 bit** This bit depth is rapidly becoming the standard for professional audio recording because it provides more sonic detail than 16 bit, meaning you can use many more effects and layers without introducing errors. Few delivery codecs now support 24 bit, most prominently Microsoft WMA Professional does. The difference between 16 bit and 24 bit won't be audible on a typical home computer speaker system, but it could begin to matter as people integrate personal computers into home theaters.

Channels

Most audio from computers is either mono with a single channel or stereo with two channels. Low bit rates generally use mono because it requires fewer bits. The industry is evolving to use multichannel (5.1 and 7.1) audio support as well. Multichannel support is included today with Microsoft WMA9 Pro and RealNetworks RealAudio 10. And of course, DVD has always supported multichannel through Dolby Digital.

Perceptual audio compression

Reducing the sample rate and number of channels can reduce the data rate; however, decent compression must go well beyond that. Most modern codecs use perceptual encoding techniques, that is, an internal model of the human auditory system determines the important parts of an audio source and spends the available bits on those parts. At low data rates, the audio can still sound quite altered from the source. The number of bits required for “good enough” quality has been dropping precipitously in recent years.

Compression efficiency

So, how does someone compare codecs? The fundamental measurement is compression efficiency, which is defined by how many bits are required to reach a particular quality level. When improving compression efficiency, you can choose to lower the data rate while keeping the quality the same, increase the quality while keeping the data rate the same, or pick a point in between.

You can't measure compression efficiency with a single number. Relative performance varies greatly and depends on the source and the data rate. Typically, the higher the data rate, the smaller the quality differences. It's with aggressive data rates and difficult content that compression efficiency makes a difference.

Authoring video for future compression

One key to great compression is creating content that compresses well. Sometimes it isn't possible to create this type of content—legacy content is legacy content. But when creating content for compression, there are a variety of techniques that you can apply to produce better results. You can apply these techniques during preproduction, production, or post-production.

Preproduction

In preproduction, you plan out the various elements of a project. It's best to introduce the idea of compression early in a project, so all members of the team can consider the implications.

Production

During production, you can make adjustments to the video, such as changing the shutter speed, camera motion, and background, and you can make the video as easy to compress as possible.

- **Interlaced versus progressive scan** For content shot with a video camera, one fundamental decision is whether to use interlaced or progressive scan video. Historically, video was always interlaced. In an interlaced frame, the odd and even lines, called *fields*, contain images that are captured half the frame duration apart. Each video frame is made up of two fields.

Why do fields exist in video? The first television sets used a single electron gun to shoot a steady beam of electrons at the inside surface of the picture tube. These electrons would excite the phosphorous material that coated the inside of the picture tube and make it glow, giving us an image. The gun would move from one side of the screen to another, illuminating a single line known as a scan line. It would then drop down a line and move across once more. It would do this for the 486 lines that comprised the image. The problem was that by the time the gun would get about half way down the screen, the lines at the top were already beginning to fade out.

The solution was to split the 486 lines into two fields called *odd* and *even* (consisting of odd numbered lines and even numbered lines). Now the electron gun could skip every other line and draw a complete field from top to bottom. It would then start from the top and fill in the missing lines. The result was that the eye would see a complete image because only half would be fading away at any one time. This “interlacing” of the odd and even lines became known as *interlace scan* or *field scan* mode.

Motion cameras and video tape recorders were built to record video in the same way. As a result, the second field shows an image that is exactly one-sixtieth of a second later in time in NTSC and one-fiftieth of a second later in PAL (NTSC=30 frames per second, each frame=two fields, each field=1/60th of a second; PAL=25 frames per second, each frame=two fields, each field=1/50th of a second). If motion occurs in the video, then the image in one field will be slightly off from the image in the other. You can see this effect if you pause your VCR at home. Many VCRs will pause on a frame and continue to alternate between the two fields that compose the frame. This causes a jittering effect called *field jitter*.

This difference between fields adds complexity to the compression process. A good video compression tool must read and understand fields, compress them independently, and allow for them to be played back as fields in the compressed format. Compression for the web generally does not show fields because computer monitors are not field based. Advanced technology for computers allows the entire screen to be drawn in a single pass, usually with more than one electron gun for tube-based displays. Flat panel monitors do not use electron guns, and they generate the entire screen instantly. This process is known as *progressive scan*.

Interlace video does not look good on progressive scan monitors. Compression techniques for the web usually change the video from interlace scan to progressive scan. The most basic method for this change is to throw away one field and duplicate the other to create the full frame. This process reduces the amount of data you start with by 50%, which can adversely affect the final quality. The more complex method is for the compressor to read both fields and *interpolate* (blend) them into a single frame by changing the pixels to smooth out the timing difference between the fields. This more complex method allows you to still use all the original data from the video and produce higher quality streaming media.

-
-
- **Camera motion** The amount and type of camera motion has a huge effect on the compressibility of content. Camera motion changes every pixel of every frame. Lack of camera motion is the simplest, of course. But modern codecs can encode camera motion, like dollies, quite well. Tilts and pans are more difficult. Depending on the codec, rotating the camera (a Dutch) can be trivial or extremely costly in compression efficiency. The most difficult motion for any codec is hand-held camera motion with random jostling.

For compression, it's much easier to handle changes of view by editing rather than camera motion, especially at lower data rates.

- **Shutter speed** This speed controls how long the aperture is open in the camera, and hence, how long light hits the film or the CCD. Slower shutter speeds let through more light, and hence less grain, which helps compression. Slower shutter speeds also cause moving parts of the image to become blurry, which makes them easier to compress.

The default shutter speed for most cameras is half that of the frame rate: 1/60 of a second for NTSC, 1/50 for PAL, and 1/48 for film. The defaults provide good results for most video meant for compression. The exception to those defaults is action and sports videography, where videographers avoid motion blur to make slow-motion images crisp and clear. Be aware of these exceptions in compression tests.

- **Backgrounds** A difficult motion for codecs is foliage rustling in the breeze. Foliage is green, which makes up the majority of luma and therefore gets allocated a lot of bits in the compression. Leaves have hard, irregularly shaped edges, and details that require a lot of bits to encode well. Additionally, the motion of foliage in the breeze is largely random, with leaves constantly overlapping and revealing, all motions that are very difficult for motion estimation to handle.
- **Depth of field** A useful and subtle technique for controlling the complexity of images is depth of field. By managing the depth of field, you can keep unimportant objects in the background or close foreground out of focus while clearly showing the actual subject of the shot. Depth of field is very common in cinematography and photography, so it looks natural to the audience.

If there is not enough spatial separation among objects, backing up the camera and using a telephoto lens narrows the region in focus.

Post-production

Post-production techniques can have as much effect on compressibility as production techniques. Because post-production techniques occur later in the process, it is often possible to adjust them based on the results of test compressions, making fine-tuning much easier.

- **Static versus moving images** As in production, superfluous moving images aren't appropriate for content that you plan to aggressively compress. They are not that important of an issue for DVD, but they are for web and CD-ROM.
- **Motion blur** In the same way that a real-world shutter speed introduces motion blur that helps compression, you can use the virtual shutter in rendering applications. For effects that support motion blur, render them with a shutter speed that matches that of the source. This speed might be expressed as a duration, or, as in After Effects, as the angle of a virtual shutter. The default is 180°, which corresponds to a shutter speed half that of the frame rate.
- **Anti-aliasing** This rendering technique softens the edges of rendered elements, letting a pixel at the edge of an area be an intermediate value. Smoothing out transitions among parts of the image improves compressibility and appearance.

Anti-aliasing may be its own control, or it may be part of an overall group of rendering settings. For example, in After Effects, anti-aliasing is automatically turned on when encoding in Best mode.



Anti-aliasing

These samples show an Adobe Illustrator® file, exported to a Photoshop file (PSD) aliased and anti-aliased, and then converted to JPEG in Adobe ImageReady®. The files were converted to the same size, and, as you can see, the anti-aliased image has significantly fewer artifacts.

- **Progressive rendering** When working with interlaced source video, it may be appropriate to render motion graphics as progressive scan when making the compression master. By rendering out the project in progressive scan mode, you can deinterlace the interlaced source, but the graphical elements remain a perfect progressive.

Capturing video: choosing formats and connections

Capturing video is how content gets from an external source like a VTR or camcorder into the computer. The type of video, or format, and to a lesser extent the connection type you select to capture content, can dramatically affect the quality of the video on the computer. It's always desirable to start with the highest possible quality video and audio signal. In almost every case, compression doesn't make video look better than the original source. The higher the quality of the source material used, the higher the quality of the final compression results. The inverse is also true, so it is helpful to build a basic understanding of the video and audio formats available and how they compare in quality.

Today's editing software, such as Adobe Premiere Pro, and many compression software tools, can capture content in a variety of ways. Video can be captured from analog or digital sources. Both have a variety of different video and audio formats available and various physical connection types that can be used. In most cases, you'll need to install a video capture card with the proper physical connectors in your computer to perform the captures.

To understand the video formats, it's first necessary to understand the video signal components. In the beginning of video, we dealt only with brightness, or luminance (luma) signals. The variations from very low brightness all the way to complete brightness gave us black and

white television. Color was added by generating two channels of color, or chrominance (chroma) values for the signal. To accommodate the millions of black and white television sets still in use, a method for adding color was invented that allowed a black and white set to ignore the colors and still show the luminance, or black and white, picture. Since broadcast television is transmitted over antennas, the luminance and two color channels are sent together as a single signal. The various video formats used today were created using this technology.

A list of the common video signal formats follows, along with the popular types of cables and connectors for transmitting them.

- **Analog formats** Analog video uses fluctuating voltage and frequency signals to transmit image data. It ranges greatly in quality but is inherently susceptible to visual interference from electronic noise, like power lines, and suffers attenuation, or the loss of signal power and quality over long cable runs (usually 300 meters or more). Equipment that supports analog signals is generally more affordable and analog video formats are still the most commonly used because of the long-term use of legacy equipment that is still active today.

Because computers only use digital data, the analog source must be converted to a digital copy to be captured. When it is sent back to tape, it is converted back to an analog signal. This conversion, known as encoding and decoding, can lead to minor errors (quantization errors) being introduced because of the imprecise relationship between analog and digital signals. By using the right format, good equipment, cables, and connections, these quantization errors can be minimized.

- **Composite** A composite video signal carries the luminance and the chrominance channels in a single signal. The biggest problem with composite video is that luma information can leak into chroma, and vice versa. This leakage can lead to noise, known as chroma crawl, that can be amplified by the video capture process and reduce the overall quality. Composite video is generally understood to be the lowest quality video signal that you can use and is more appropriate for video delivery (broadcast) than as a source for content creation (video editing, DVD authoring, or encoding for the web). Common sources of composite video are VHS (1/2 inch tape) VCRs and camcorders, broadcast television signals, and older U-Matic 3/4-inch professional VTRs.

The connectors common to composite video are RCA or Cinch in the consumer space and Bayonet Fitting Connector (BNC), a twist-locking connector common to professional and broadcast video equipment. As the quality of the RCA or BNC cable that is used increases, the noise and attenuation decreases.

- **S-Video** The introduction of the S-Video format was a major advance in quality for consumer and professional video editors. S-Video splits the luma and chroma channels into two separate wire pairs, which eliminates generally half of the composite noise artifacts. The two chroma channels are still combined and still interfere with each other, leaving some noise in the signal. Many professional and consumer-grade video equipment manufacturers, including Hi8 camcorders and S-VHS VCRs, provide an S-Video connector. While no longer considered true professional grade, it is a big improvement from composite and can be considered a good video capture minimum standard for editing and compression.

S-Video, or Y/C (Y for Luma and C for chroma) video signals mostly use a specific four-pin connector on S-Video cables. These cables are intended for short distances (usually 100 feet or less). Some types of professional and broadcast equipment use two discreet BNC cables, one

Analog to DV converters

There are boxes that can perform a live conversion from analog video formats to DV25 for capture. They usually offer composite, S-Video, and unbalanced analog audio conversion to the DV format for storage on the computer. The 25-Mbps data rate of DV is sufficient for content that camcorders produce, but it can be troublesome for converting sources with a lot of analog noise on lower-end versions of the converters.

for luma and one for chroma, to carry S-Video signals. This setup provides a connection that can be run farther and that can be locked into place to prevent the cable from falling out and losing the signal.

- **Component analog** Component analog, also known as YUV (Y for luminance, U for one chroma channel, and V for the other chroma channel) was the professional and broadcast standard for many years and it is still widely used today. With component analog, the luminance signal and the two color signals are all transmitted over their own dedicated cables. Because all three components or channels of the video signal are transmitted independently, the quality of the signal is quite high. Noise is very low and the colors in the video are richer and more precise. In broadcast and professional environments, component analog has mostly yielded to digital video formats. Component analog has become very popular in the consumer market, though, as the preferred format for connecting home DVD players to new television sets, owing to the pristine nature of the signal compared to S-Video and Composite formats.

In professional environments, the component analog signals were carried by three individual BNC cables because of the long run and locking nature of this connection. In the new home DVD players, it is common to see RCA or Cinch connectors used that are color coded red, green, and blue to make it easy to connect the DVD player correctly.

- **Unbalanced audio** Analog audio is basically a single voltage and frequency-based analog signal, no matter what the connector type is. The connector type can greatly affect the quality and generally determines the use of the audio

There are three basic kinds of unbalanced audio connectors. The most familiar to consumers is traditional unbalanced audio that uses RCA or Cinch jacks (normally red for right and white or black for left). The RCA connectors use a single two-wire connection for the left and right channels, which requires either a single RCA cable for each channel or a single “stereo pair” cable. The left and right channel cables are manufactured as a single cable with two RCA connectors on each end.

Other single-cable unbalanced connectors exist. The One-Quarter Inch connector is common to high-end headphones and is standard on most home stereo equipment. Similar but smaller is the One-Eighth Inch connector that is standard for portable audio devices and computer sound cards. Because these are not amplified connections, they are also called line level signals. In the correct circumstances, unbalanced audio can work well, but it is susceptible to interference from other electronic sources’ grounded outlet loops and isn’t appropriate for many professional applications.

- **Balanced audio** The professional standard for analog is balanced audio. Unlike the two wires of unbalanced audio, balanced audio uses three wires, one for the positive signal, one for the negative signal, and one for an isolated ground. Using three wires dramatically improves resistance to interference and ground loop noise. Balanced audio normally uses the relatively large XLR, or three-pin, locking connector. Once again, longer cable runs and a locking connector form the basic standard for a professional connection.

Digital formats

Digital video records the original video signal on tape using digital data similar to that used by computers. All brightness and color data is stored in this very accurate computer-style method and as such suffers no problem of interference or noise between the channels of the video signal. Digital signals are therefore cleaner and crisper than analog sources. The digital data format

allows digital video to be “copied” from tape to a computer hard disk without the conversion process that analog video must go through when it is captured.

Many digital formats have emerged in the video industry. Many of these formats are related to compression and not acquisition or filming with a digital camera. This section covers only the most popular acquisition formats of digital video.

- **DV25 (DV)** This affordable digital video format is most commonly known by the physical connection standard it uses, IEEE-1394 or FireWire. It is actually a digital video format that uses only a fraction of the data rate and speed available in the IEEE-1394 connection standard. This format revolutionized the video editing industry, becoming the fastest growing format for new editors and the de facto standard for news, owing to the small size and relatively low cost of the DV camcorders. The video format is compressed 5:1 during acquisition or recording on the camcorders in order to reduce the data rate of the signal and is only slightly higher quality than the S-Video analog format. Higher end DV camcorders can produce a DV signal that is almost as good as component analog video at the same DV data rate because of better recording technology.

The DV format is digital. It includes a series of computer-style 1s and 0s and no analog signals. This format enables the video data to be “copied” from the DV tape directly to a computer equipped with a FireWire connection. The data on the computer is exactly the same data as on the tape, so the result is a perfect copy. Many nonlinear video editors, like Premiere Pro, can read and edit this format natively. The finished video is virtually the same as the original source footage when it is final. Few to no processing errors are introduced.

The physical connector for the DV format is based on an IEEE standards committee design known as IEEE-1394 (also known as FireWire and i-Link). As a digital format, both the audio and the video signals can be carried on a single, inexpensive FireWire or 1394 cable. In addition to video and audio signals, this cable also carries device control signals, allowing the computer to control the DV camcorder automatically. Originally, only DV25 was used on FireWire, a 25 megabit (Mbit) data rate format. Increasingly, other digital formats are being supported, like DV50 (50 Mbit DVCPPro and DVCam) in the professional and broadcast markets. Newer 100 Mbit formats like DVCPProHD are becoming popular as high-definition digital video formats.

Adobe Premiere Pro supports capturing DV25 and DV50 over any standard FireWire port common to most new computers. Inexpensive FireWire add-in cards can be purchased and added to most computers and laptop computers. Both four-pin and six-pin DV connectors are available, and they are compatible with each other. DV camcorders usually only feature the four-pin connector as it is a smaller connector and only the four pins are needed for data transmission. The extra two pins on the six-pin connector are for power commonly used, for example, by FireWire hard disks.

- **SDI (Serial Digital Interface)** SDI is the high-end professional format for digital signals. The SDI format, technically known as CCIR-601 or “601”, is superior in quality to component analog video and has become the standard in professional and broadcast television production environments. Because it is digital and uncompressed, the format can be edited and copied many times with no detectable loss of quality. SDI was originally meant for standard definition, but a high-definition version, HD-SDI, is growing in popularity as HD becomes more common. (Please see the Adobe HD Primer.) SDI is typically only supported in high-end gear and can be quite expensive, especially for SDI video monitors.

As a digital format, the video, and sometimes the audio, can be carried in a single BNC cable.

-
-
- **Digital audio** There are many different digital audio formats. Consumer grade digital audio is carried on electronic (such as SPDIF) or optical (such as square TOSlink or round Miniplug) connection cables. DV carries an embedded digital audio signal with the digital video signal. The most common professional grade digital audio format is Audio Engineering Society/European Broadcasting Union (AES/EBU) digital audio format. A stereo channel of AES/EBU can be transmitted over single BNC or single XLR connection cables. SDI video equipment can sometimes carry embedded AES/EBU digital audio, but it is more common for a separate AES/EBU or balanced audio stereo pair to be used with the SDI video signal.

Choosing the right capture codec

When capturing content, you must choose an appropriate codec. The choice depends on the capture hardware you're using, the format you're capturing from, and the disc storage you need. There are three basic types of codecs to choose from: DV, Motion JPEG, and uncompressed.

Understanding DV codecs

The DV25 format is ideal when capturing from a DV25 deck. Over the FireWire connector, you can use DV25 to copy the native bits from the tape onto a hard drive, making a perfect copy. If you use a different format at a higher data rate, the resulting video will still be slightly lossy.

Using a DV bridge to convert analog source to DV25 may work, but the quality lessens. The data rate of DV25 isn't quite enough for optimal encoding from source with analog noise. Also, the 4:1:1 color space also entails a loss in quality.

Understanding Motion JPEG codecs

Motion JPEG has been the leading format for high-end video-editing systems for some time. The proprietary codecs supported by many vendors of capture cards, like the Targa and Pinnacle codecs, are variants of Motion JPEG and share its properties.

Motion JPEG is 4:2:2, and most implementations allow you to specify a data rate or quality level. Some implementations can even do a mathematically lossless Motion JPEG. Typically, data rates of 50 Mbps or greater are considered broadcast quality.

Understanding uncompressed codecs

More and more cards are supporting true 4:2:2 uncompressed video. The data rates are large, and the quality isn't any better than a lossless compression. One significant advantage of some uncompressed codecs is support for 10-bit-per-channel captures, which can help the quality from analog and especially SDI sources. High-end tools, like After Effects 6 Professional, are able to process video in more than 8 bit per channel, improving the quality. Some hardware vendors also include support for 10-bit rendering in Adobe Premiere Pro. Because uncompressed codecs have such high data rates, you should use a fast drive system, such as RAID.

Note: Don't confuse uncompressed 4:2:2 with uncompressed RGB codecs, often called *None*.

Exporting video

Capturing from tape is required for many projects, but for some projects, all content is created electronically. In other projects, a lot of creative work is done in editing and compositing on the computer. In those cases, going back out to tape is an unneeded distraction. Many tools, like After Effects and Adobe Premiere Pro, can export to a compressed file from the timeline, which works well for a single file. However, if features in other encoding tools are required, like volume encoding to multiple output formats, it is useful to export an intermediate file from the timeline and use that for further processing.

Exporting as file

A file exported for later processing is called an *intermediate file*. The goal of creating an intermediate file is to provide high quality in the most useful format without introducing any compression artifacts.

Choosing whether to preprocess or to avoid rendering

The first question that you should ask is when to apply preprocessing (described in the next section): during the rendering phase or later. Preprocessing during rendering generally results in a smaller file but adds more rendering time. It also means you can't use an encoding tool for special operations not available in the exporting application, like inverse telecine.

Generally, when you avoid rendering, you should use the same codec and file format that you used in the source video. If you don't preprocess, you don't need to uncompress and then recompress unmodified frames—instead, you copy them from the source into the output. Different tools support this process in difficult ways—in Adobe Premiere Pro, you do this by turning off the Recompress option in the Video pane of the Export Movie Settings dialog box. For this process to work, the video must truly be unchanged with the same frame rate, frame size, pixel aspect ratio, and so on. The simplest way to ensure that no changes were made to the video is to use the same project preset for export that was used for import, which is the default behavior in Adobe Premiere Pro.

There is one exception to exporting with the source codec: when the source is DV25 and contains a lot of motion graphics or compositing. Because these synthetic elements have a lot of precise color, exporting to a 4:2:2 codec, like Motion JPEG, instead of NTSC DV25's 4:1:1 provides twice the color detail.

When preprocessing, choose the highest frame rate and largest frame size that you will use. If targeting web codecs, use a square pixel aspect ratio.

Choosing a file format for rendering

You can choose from a variety of file formats for rendering. The most common are QuickTime and AVI, both offer good-quality intermediate codecs. You can also use MPEG-2 at high data rates. Modern tools like After Effects 6 and Adobe Premiere Pro can use Windows Media 9 for exporting. The key is to make sure that the encoding tool supports the file format.

Choosing a codec for exporting

When exporting an intermediate file, pick a codec and a data rate that doesn't introduce any compression artifacts. When you are not preprocessing, it's typical to use the source codec or Motion JPEG. When you use Motion JPEG, use it in single field when you export as progressive scan or interlaced when you export as interlaced.

For AVI files on Windows, the Huffuyv codec is quite popular to use.

Understanding preprocessing

Preprocessing is the middle step of video compression and involves transforming the decoded frame of the source video into the optimized frame given to the codec. The goal of preprocessing is to transform the video into the final format.

Why preprocessing matters

Preprocessing is one of the subtler steps of compression, but it is often the most time consuming and the most difficult. It's also very important for quality, and getting it right can dramatically


improve the quality of the final video. A well preprocessed video can look better than the same source poorly preprocessed, but can be encoded at twice the data rate.

Preprocessing is critical for web video, but generally it is not needed when going to video formats like MPEG-2 for DVD or SVCD. Because those formats are played back as video, instead of on a computer screen, you can leave them as is.

Also, you can create rendered graphics natively for computer playback without any preprocessing.

Deinterlacing video

The most basic and critical form of preprocessing for the web is deinterlacing. Most video is produced and edited as interlaced (although, an increasing amount is done as progressive scan). But almost all web video is encoded as progressive scan. If the source frame of video is left with interlaced lines intact, the output shows the stacked lines. The lines look bad, but worse, those lines have sharp edges and are difficult to encode. So, at the same bit rate, deinterlaced video shows fewer artifacts than interlaced video.



Differences between interlaced and deinterlaced video

The same video frame, shown prior to compression and after compression, as both progressive and interlaced. Note that there are more compression artifacts in the interlaced version, even though the file size is the same.

Note: MPEG-2 encoded for DVD or SVCD don't need to be, and shouldn't be, deinterlaced. MPEG-2 can encode frames in a specific interlaced mode.

Working in progressive mode when authoring for progressive output can give you substantially better results when exporting from a tool like After Effects and Adobe Premiere, where you can add a fair amount of graphics or other synthetic effects on top of the video. The tools automatically deinterlace interlaced sources while they do all of the additional processing in progressive mode, providing higher quality than if you deinterlace everything during export.

Understanding telecine and inverse telecine

A special case of progressive scan is film source and film-like 24p production. Traditional motion picture film runs at exactly 24 frames per second (fps) progressive scan, which doesn't convert to the 59.94 fps of NTSC or the 50 fps of PAL.

Converting from film to PAL is easy. The video is sped up 4% to 25 fps and converted into progressive PAL.

NTSC is a little more complex to perform. The film is first slowed down to 23.976 fps, and then it is converted to the 59.94 fps by using 3:2 pulldown. The first frame of film becomes three fields of video, and the next frame becomes two fields of video, and the next becomes three fields, and so on. This process results in two out of five frames having fields that come from different source frames.

Fortunately, 24p video is evolving into a production standard of its own, so editing tools can increasingly just author in 24p and then convert from that universal master to PAL, NTSC, film, or web formats as appropriate. When possible, it is best to convert to 24p on import and then do all the work on the project in 24p mode.

Cropping

Cropping is excluding areas of the source frame from inclusion in the output frame. Content composed for video normally confines critical elements to what's called *safe areas*. The safe areas exclude the edges of the screen, which may be cut off or distorted on some televisions (especially older ones).

In some tools, you can crop an image by using a specific filter. In others, like After Effects, the easiest way to crop an image is to scale the video up so the unwanted edges are no longer part of the frame. If an asymmetric scale is desired, you can move the anchor point of the video as needed.

One special case of scaling is going from 720 pixels wide source to 352 pixels wide output for VideoCD. The specified method for this scaling is to crop eight pixels from the left and right of the source, leaving a 704 pixel wide frame that you can reduce by exactly 50% to 352 pixels wide.

Another special case is going between 486 line and 480 line sources in NTSC. Motion JPEG commonly uses 720 x 486 and MPEG-2 commonly uses 720 x 480. The different line count isn't due to different scalings, but it is due to six more lines of the video signal captured in the 486 line mode. When converting from 486 line to 480 line (for example, in DVD authoring), you'll need to crop six lines. Generally four lines off of the top and two lines off of the bottom works well. Conversely, if you're going from 480 line to 486 line, you'll need to add six lines (again four lines on the top and two lines on the bottom are good choices). Doing odd numbers can switch field order, resulting in substantial problems.



Safe areas

The inner boundary is the title-safe area, and the outer boundary is the action-safe area.

Scaling

Scaling is resizing the cropped source rectangle to the proper output frame, for example, going from a 720 x 480 source frame to a 320 x 240 output frame for web video.

Scaling up on either axis forces the tool to make up data that's not in the source. It's always preferable to scale down on both axes if possible.

Aspect ratio correction

One important element of scaling is aspect ratio correction. Most production formats, like DV, are nonsquare pixel. For example, DV in NTSC is 720 x 480, which, by looking at pixels alone, would imply an aspect ratio of 3:2. However, DV isn't ever 3:2; instead, it's either 4:3 or 16:9, depending on how it was produced. In 4:3 mode, the pixels are more narrow than square, and in 16:9 mode, they're more wide than square. Nonsquare pixels are also used for MPEG-2 for DVD and MPEG-1 for SVCD and with different shapes for MPEG-2 for SVCD.

Conversely, most web video formats use square pixels, which makes it easier to do the math. For example, 320 x 240 is 4:3 in square pixels.

The key to getting aspect ratio correction right when going to square pixel sources from non-square pixel sources is to make sure that the output frame size matches the source aspect ratio. So, if you use a 4:3 source, a 4:3 frame size such as 160 x 120, 320 x 240, 400 x 300, 512 x 384, or 640 x 480 are all good choices, even if the source frame size is 720 x 480, 720 x 486, 640 x 480, or 352 x 480.



The importance of aspect ratios
This graphic shows a 4:3 and a 16:9 720 x 480 image; each transformed into a 4:3 and a 16:9 output. Getting the output aspect ratio correct is critically important.

Noise reduction

Noise reduction encompasses a variety of techniques used to remove noise, especially grain, from the image, making it easier to compress. Very simple noise reduction algorithms are just blurs—they can hide grain but can make the image quite soft as well. The commonly available Median filter is a good blur for noise reduction. More advanced algorithms try to blur only the parts of the image with grain and may take advantage of the differences among frames.

On the high end, there are tools that can perform automatic or manually assisted scratch removal and other improvements for degraded source. Normally these are stand-alone tools or plug-ins to a video tool and need to be used earlier in the process.

You don't always need to apply noise reduction. Very clean source isn't helped by noise reduction, and noise reduction may hurt the source audio. Typically, you should use only the minimum noise reduction required to suppress poor quality in the source.

Image processing

Image adjustment is the process of using Photoshop-like controls to change the overall image. Because of different color spaces used by different codecs, image adjustment used to be required for most projects. Today, for the most part, modern formats and tools automatically take care of image adjustment so it is mainly done to improve the suitability of the content for compression and isn't required for most projects.

Luma range expansion

In the past, a major issue in video compression was making sure that black was pure black and white was pure white. In essence, some tools, formats, and codecs considered black to be 0 and white to be 255, but others used the ITU standard of black as 16 and white as 235. For the most part, modern tools automatically keep these consistent. But if blacks and whites appear either crushed or faded, there might be a luma range mismatch during the authoring process. Normally all RGB codecs use the 0-255 range, Y'CbCr codecs use the 16-235 range, and video tools automatically convert from Y'CbCr or RGB to the other. Beware of codecs that use studio RGB, which also uses the 16-235 range (you may need to add additional contrast when converting to a delivery codec from these codecs).

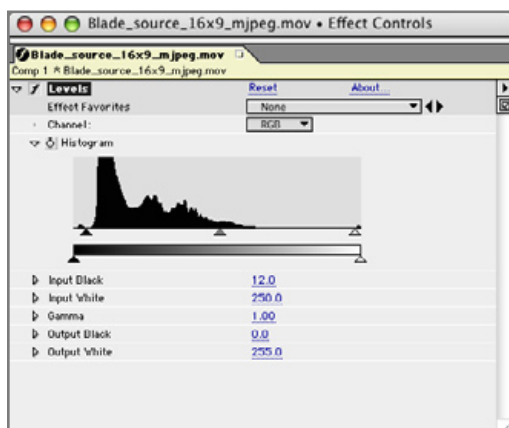
Black levels

Another goal of controlling luma is making black and white in the source mathematically black and white in the output, respectively. For elements like digitally rendered title cards, achieving this goal should happen automatically. But video captured from analog tape always has noise, which randomizes the luma values of pixels and leaves some black a higher value than black.

While you can get values down to black by either reducing brightness or increasing contrast alone, generally using both techniques moderately provides a subtler experience than doing only one or the other. Applying noise reduction before adjusting luma range can help you return some of the brighter black pixels closer to black, so you can use a less aggressive luma expansion.

The classic Adobe Levels filter with histogram is invaluable for this type of processing. The key is to make sure that the region, which should be black, is being processed as black. If brightness and

contrast aren't enough, you can also use a black or white restore filter. These filters set a threshold below which all pixels become black, and those pixels equal to or above it become white.



Levels histogram

The Levels filter in After Effects 6 with typical settings to set the luma black and white levels beyond that of the noise floor and ceiling.

Gamma correction

Gamma is a measurement of how the 0-255 luma range turns into brightness when it is displayed. With different gamma values, 0 and 255 appear the same way, but values in the middle of the range appear differently. Different platforms (Mac, SGI, and Windows) have different default gammas. The net effect is that some files appear to have brighter midtones on a Mac than on Windows. The good news is that most modern video codecs are able to automatically convert to the proper gammas upon playback.

Unfortunately, this is a transitional period in the industry with a lot of older codecs that don't perform gamma correction still in use. For content that you need to play back correctly on a platform other than the one it is authored for, it's important to spot check samples by using that combination to make sure they appear correctly.

Audio processing

Audio may only get a fraction of the bits of video, but it makes up half the experience. While audio processing is generally easier than video processing, it is important to correctly process both audio and video.

Resampling audio

For low data rates, especially those targeting modems, generally you need to reduce the audio to a lower sample rate, typically to 22.050 KHz or lower. Modern tools automatically use high-quality algorithms for rate resampling, which provide much better quality than the defaults of a few years ago.

Setting volume

Like digital audio for video in general, the peak volume should be slightly below the maximum volume allowed. By using a compressor/limiter to reduce dynamic range, you can sometimes improve audio quality at very low data rates. In general, audio mixed for television broadcast should encode just fine.

Adjusting channels and bit rates

Mono content encodes better than stereo at lower data rates. Consequently, it's preferable to convert the audio from stereo to mono when targeting lower bit rates (typically anything below 32 Kbps).

Given the choice of setting audio to stereo at a lower sample rate or setting it to mono at a higher rate, you should choose mono because sample rate matters more than stereo for the content sounds.

A few high-end systems also support 5.1 and 7.1 discrete surround-sound audio. The data rates required for this surround sound (128 Kbps and above just for audio) generally preclude their use in real-time streaming, but they can be quite usable for disc-based playback.

Dolby Pro Logic encoded-surround sound (which stores surround sound in a stereo pair) doesn't survive the encoding process into many audio codecs. If Pro Logic encoding is required for a given project, test carefully to make sure that the information is surviving the encoding process. With some codecs, raising the data rate or turning off options like Joint Stereo might help. Some codecs can't encode Pro Logic.

Reducing noise

There are also noise reduction filters for audio, although these filters are normally part of professional audio tools like Adobe AuditionTM, not a part of the compression tools. Providing clean, undistorted, hiss- and pop-free audio source results in better compression.

Encoding

If preprocessing is the art of compression, encoding is its science. Encoding means defining the correct parameters, like data rate, frame size, and frame rate, to get optimum results for the project in question. Balancing these parameters can be difficult because there are many trade-offs. As always, it's most important to understand the goals that the compressed video is meant to achieve. This section starts with parameters that apply to all codecs and continues with parameters that are specific to video and audio.

Choosing a data rate mode

Different codecs and formats can provide different data rate modes. Some only support one data rate mode, other codecs provide a half-dozen different options. What's most important is to understand how the compressed media is going to be used—in most cases, there is a single optimal mode for a particular use.

You can use any of these data rate modes:

- **Average limited** The traditional and still the most common kind of data rate limitation is average limited. In this mode, you specify the average data rate, and the codec tries to make sure that the file ends up at the requested average. The codec varies the rate up or down to match the changing difficulty of the video.

Most CD-ROM and QuickTime codecs still work this way. Average limited is used for most progressive download files where file size and data rate are the limiting factors.

- **Peak limited** A peak limited encode is generally meant for hardware devices or high-data-rate content where the speed of the decoder (CPU, memory, or disc) is the limiting factor. Specifying a particular peak limited data rate doesn't mean the peak of the file will actually be at the requested bit rate. The peak limited data rate is the maximum rate supported, but the encoder is free to use a lower rate if appropriate for the content.

Peak limited also works well for high-resolution and high-data-rate web and CD-ROM files that need to play back on lower-end computers. Some codecs support encoding with both average and peak limiting.

-
-
- **Buffer limited** For real-time streaming, buffer limited encoding is the method of choice. In buffer limited, the average data rate isn't applied to the overall file, but it is applied to any given section of video. So, when video is encoded with a three-second buffer, any arbitrary three-second chunk of the file has a data rate at or below the target data rate. Most streaming encoded files use average limited and buffer limited in combination. Buffer limited is often called *constant bit rate* or CBR.
 - **Quality limited** Quality limited is different from the other methods because it doesn't specify a data rate. Instead, it specifies a quality target, and each frame uses as many or as few bits as needed to reach that quality target. This mode isn't compatible with real-time streaming. However, it can work very well for archiving content or for downloading missing critical video where assuring quality is more important than the file size or download time. For example, a stock video company may use quality limited encoded files for short web samples to make sure that the content looks as good as possible and does not waste any bits on unnecessarily high data rates that would make the content easy to encode.

Most quality-limited codecs specify the quality on a 0-100 range. However, actual visual quality produced by the same value can widely vary among codecs. Some codecs support average and quality limited together; the quality limit sets a minimum quality for each frame. If a frame falls below the quality limit based on available bit rate, the bit rate for that frame is raised so it would meet the minimum. Typically, one or more frames after the low-quality frame are dropped to preserve the average bit rate. This mode sacrifices smooth motion in favor of image quality.

Choosing between encoding speed and quality

Many codecs offer a control that specifies a trade-off between encoding time and quality. The faster modes typically perform much less exhaustive motion estimation. For the most part, big speed changes cause much smaller quality gains—an eight times slower encode might only have 20% better compression efficiency. The correct trade-off between speed and quality varies depending on the project. For DVD projects, titles under an hour won't gain much by slower, higher-quality encoding because the data rate is so high. Web-distributed content almost always benefits from the slower, higher-quality encoded files because there are almost never enough bits for web video.

Setting frame size

Frame size is the number of pixels high and wide that form the frame of the final video. It's also sometimes called *resolution*. Roughly speaking, the data rate should change in proportion to the area of the frame (height plus width). Thus, converting from 320 x 240 to 640 x 480 would require about four times the data rate, but converting to 192 x 144, you could reduce the data rate down to one-third of the amount it was before. The effect isn't quite as strong as the math would suggest; however, because there are fewer pixels, pixel quality becomes more important.

Most codecs and formats require the height and width of a video to be divisible by two. Most codecs give best compression efficiency when height and width are divisible by 16.

Specifying aspect ratio

An important facet of frame size is aspect ratio. Most web and CD-ROM formats are square pixel by default, and many formats only support square pixel. Conversely, the disc-based VCD, SVCD, and DVD formats are always nonsquare pixel. It's important to figure out the target

aspect ratio first, and then make sure that the frame size you choose gives the correct aspect ratio given the picture shape in which the content is delivered.

| 4:3 | 16:9 |
|-----------|-----------|
| 160 x 112 | 208 x 120 |
| 192 x 144 | 256 x 144 |
| 256 x 192 | 336 x 192 |
| 320 x 240 | 432 x 240 |
| 384 x 288 | 512 x 288 |
| 512 x 384 | 688 x 384 |
| 640 x 480 | 848 x 480 |

Specifying the correct aspect ratio

Suggested frame sizes for 4:3 and 16:9 aspect ratios when working with square pixels.

Adjusting the frame rate

Frame rate is the number of frames per second at which the video runs. Higher frame rates appear smoother, and lower frame rates appear choppy. Anything below 20 fps doesn't appear smooth, and below 10 fps, the video starts appearing more like a filmstrip than moving video.

Reducing or increasing the frame rate has much less effect than frame size does on the data rate because reducing the frame rate requires encoding twice as much motion for each frame, which requires more bits for each frame. Also, each frame appears for twice as long, so there's more time to notice each artifact; thus, the image quality per frame must be higher to provide the same perceived quality.

Consequently, you should try to use the source frame rate if possible, and no less than half the source frame rate, unless you target real-time streaming to modem users. Raising frame rates generally pays off more than raising frame sizes in improving user experience.

Note that NTSC frame rates are 0.1% lower than an integer number. So, "30 fps" is really 30/1.001, which is conventionally rounded to 29.97. A 29.97 fps output is more accurate than 30 fps when dealing with NTSC source. Also note that interlaced video has two fields per frame, and you can turn these fields into two frames on output. So you can turn 29.97 fps interlaced source into 59.94 fps progressive.

There is another *very* important rule when changing frame rates—the output frame rate must be an integer division of the source frame rate. Therefore, the output, can be one-half or one-third the source but not three-fifths. Good frame rates for each type of source appear in the table below.

| NTSC | 24/23.976p | PAL |
|-------|------------|------|
| 59.94 | 23.976/24 | 50 |
| 29.97 | 12 | 25 |
| 15 | 8 | 12.5 |
| 10 | 6 | 8.67 |
| 7.5 | N/A | 6.25 |
| 6 | N/A | 5 |
| 5 | N/A | N/A |

Suggested frame rates

It's important to use frame rates that are evenly divisible by the source frame rate.

Negotiating frame dropping

Some codecs, optionally or automatically, can increase the number of bits spent on difficult frames to maintain image quality by dropping the frame rate to keep the target average data rate. This process generally takes the form of a control that sets the minimum image quality per frame. If the bits that the codec would normally allocate to a given frame aren't sufficient to hit the target quality, it raises the data rate for the frame to the minimum required for quality; then drops one or more frames that follow to keep the average data rate on target.

Whether this process is appropriate depends on the content. For entertainment content, like movies and especially music videos, having the video smoothly in sync with the audio is important, so choosing not to use frame dropping is appropriate. Conversely, presentation video, especially with a whiteboard in the background, might be better with many frames dropped and on-screen text remaining legible.

How keyframe rate comes into play

Keyframe rate controls the number of keyframes (I-frames) that are inserted in the video. Depending on the tool or the format, keyframe rate is measured in a keyframe every x frames or a keyframe every x seconds. MPEG-1 and MPEG-2 need a lot of keyframes; typically two a second. Web formats might have only one keyframe every 10 seconds.

Most formats also insert additional keyframes, called *natural keyframes*, when the video changes dramatically, like at scene cuts, or when there is extremely fast motion, like in explosions. These natural keyframes reset the “keyframe every” value, so if a video has a lot of rapid cuts, it doesn't matter if the automatic keyframes are set to every 100 or 1000 frames—all of the keyframes are natural. For codecs that also do natural keyframes, the “keyframe every” value really just sets a maximum distance among keyframes.

How keyframes use 1-pass and 2-pass algorithms

Historically, most codecs worked in a single pass—they read in the source file and wrote out the compressed data frame-by-frame. This process certainly worked, but it required the data rate mechanism to guess the makeup of the future content. Many good heuristic algorithms have been developed for this process, but every algorithm guesses wrong in some cases. The worst situations are when a 1-pass algorithm interprets the video getting more difficult and raises the data rate to keep quality constant, or when another, more difficult sequence is following a few seconds behind the current frame; then just when the video gets the most complicated, the codec has to reduce the data rate below the average to compensate for the previously raised rate.

In a 2-pass algorithm, the codec first analyzes the video and builds a table of the relative encoding difficulty in each frame. Then it calculates the actual bits per frame needed to achieve the highest possible average quality and applies these bits in the second pass. The net effect is the same as a quality-limited encode where the correct quality value reaches the target data rate.

One benefit about 2-pass encoded files is that the quality is very consistent, so that the easier and the harder parts more closely match. This consistency is less jarring for the user than a video where some scenes look great and others look terrible.

You can use 2-pass encoding with bit-rate-, peak-, and buffer-constrained files if they are supported by the format and codec. With buffer limited, larger buffers give the algorithm more flexibility to distribute bits.

A few codecs have 3-pass or N-pass systems. With these systems, the codec analyzes the results of the second pass, builds a new data rate allocation table, and encodes the file with hopefully

better results. In N-pass systems, the codec can repeat this process multiple times, although obviously diminishing returns set in quickly (generally, a good 2-pass implementation is already very close to the optimal results).

Modern formats like Windows Media and RealMedia have 2-pass integrated into the user interface. AVI and QuickTime codecs that support 2-pass have more difficulty making this work. With AVI, the codec typically needs to create two different encoded files—one dummy file and one full file.

How sample rate affects the quality of audio

Sample rate is the number of audio measurements made per second: 44.1 KHz for CD, 48 KHz for DV, and 8 KHz for analog telephone.

Anything less than 32 KHz sounds noticeably degraded compared to the source. For most broadband and CD-ROM use, you should use at least 32 KHz.

Unless the specs of the target computers are available, use 44.1 KHz for computer playback. It is a safer choice than 48 KHz. Many sound cards, especially older and cheaper models, only play 44.1 KHz internally, so by using 48 KHz, you force a software sample-rate conversion. This conversion isn't as good as a high-quality filter that you can apply at compression, and so final output quality is worse. 48 KHz is a fine choice with audio source at 48 KHz or higher and an output computer or device with native support. A few codecs, most notably WMA9 Professional, can also encode at higher sampling rates, like 96 KHz.

Note: As mentioned previously, almost all audio codecs are 16 bit. WMA Professional is a recent exception that can usefully take advantage of sources that are 20 or 24 bit. There are no problems in encoding from 16-bit source to 24-bit source (nor is there much advantage).

Channels

For most codecs, channels are a choice between mono and stereo. Some codecs are mono only, stereo only, or optionally offer multichannel modes like 5.1 or 7.1.

Older codecs, like IMA, require the data rate to double when going from mono to stereo. Modern codecs are able to encode the content the same among channels only once, so the additional bit rate required to provide stereo at the same perceived quality is more like 20-50% of the mono bit rate, depending on the codec and how different the channels are.

As discussed previously, channels are less important to perceived quality than sample rate or the presence of audible artifacts. A 32-KHz mono sounds better to most users than 22.050-KHz stereo. Generally, you should use stereo only if you can use at least at 32 KHz and ideally 44.1 KHz.

Multichannel (5.1 and 7.1) audio codecs also exist. The most common of these codecs are Dolby Digital (AC-3) for DVD and WMA9 Professional. Multichannel is supported in RealAudio 10.

Delivery methods for video

The only reason to compress content is to deliver it; the way you compress content depends enormously on how you will deliver it. It's useful to understand the different environments in which people use video and the trade-offs each environment entails.

Disc-based delivery

The first half of compressed video history was in the CD-ROM era, which went from tiny slide shows to full-screen and full-motion web video. Today, video on mobile devices is becoming available. The good news is that disc playback keeps getting better even as the focus on the industry progresses.

Disc-based delivery works well, and the data rates available are radically higher than those for web video. Even an ancient 8x IDE CD-ROM can provide a sustained data rate of 8000 Kbps! The limit on CD-ROM and DVD-ROM isn't a question of data rate but of disc space available and processor speed to decode it. The largest CD-ROM discs today are 800 MB. With DVD-5 giving 4.7 GB and DVD-9 up to 8.4 GB, DVD gives vastly more room.

When targeting faster personal computers, disc-based playback can do incredible things. You can play video at 1280 x 720 on most computers shipped in 2002, and you can play 1920 x 1080 on the high-end computers of 2003. You can easily play video at 640 x 480 on computers several years old, especially if you use older, performance-tuned codecs (with a corresponding reduction in compression efficiency).

Downloadable files

The most basic type of video on the web is a simple, downloadable file. The file is downloaded to the local hard drive like any other file.

The advantage of downloadable files is that users don't have any expectation that they'll get real-time playback while streaming. Instead, their browser tells them how large the downloadable file is. Thus, you can use very high data rates. Also, because the file is local, users might pass the file among themselves instead of every person downloading the file, reducing the cost of bandwidth per viewer. The drawback is that users can't view any of the file until the download is complete, which creates the risk for users to spend a lot of time downloading a file they would have realized in the first 10 seconds that they don't want.

Downloadable files are often used for marketing or noncommercial use. They are very common in video game samples where very high data rates are required to encode game play in sufficient quality.

Progressive download files

During progressive download, a file is transmitted through standard web protocols like FTP or HTTP, but users can view the part that has already been transmitted while the remainder is still transmitting.

With most formats, the video starts playing automatically when the amount of time remaining in the download is less than the duration of the clip. This format allows the clip to finish downloading while it is playing. However, in some cases, the video playback catches up with the downloaded data, causing the playback to pause for a moment or longer.

Because progressive download uses protocols like HTTP and FTP to ensure that all packets arrive, it isn't affected by the video and audio glitches associated with real-time streaming. However, progressive download can't provide reliable real-time performance. A 400-Kbps file playing over a 200-Kbps connection takes nearly a minute of buffering per minute of video. Of course, that same 400-Kbps file couldn't stream over a 200-Kbps connection.

Progressive download is the easiest, cheapest, and most forgiving type of web video to author and deploy. It's a great way for users new to web video to start.

Real-time streaming

The next level of web video is real-time streaming. This format offers the traditional experience of seeing the video after only a few seconds of buffering, but it risks image glitches, poor quality, and limiting the data rate to what's available.

So, why put up with that? Because real-time streaming provides real-time performance and random access. Progressive download works well with a movie trailer, but not for movies, especially if the user wants to skip to the last 20 minutes. With real time streaming the file starts playing in under a minute, no matter the length, and users can skip to exactly where they want to. Also, most real-time streaming formats provide scalability, which allows streams to dynamically adapt to the user's connection speed.

The user experience of real-time streaming has improved compared to only a few years ago. Codecs are much better and so are servers. Although server quality isn't often a direct focus, it has huge payoffs for the user. Modern servers are able to aggressively manage the buffer, sending frames minutes before they're needed, which provides plenty of time to retransmit dropped packets and adjust bandwidth based on changing network conditions.

People in the compression community often debate when to use real-time streaming and when to use progressive download. There are cases where the choice is obvious, and other cases where the choice is not so obvious. A good rule to consider is duration—clips under three minutes are often better as progressive, and clips over three minutes are often better as streaming.

One important feature of modern real-time streaming is scalability. Data rates on the Internet aren't reliable. Connectivity through a cable modem can substantially vary throughout the day. Many users share the T1 connection and faster pipes to the workplace, leading to rapid changes in available bandwidth per user.

It's not enough to measure the data rate at the start of a stream. Instead, the player and server need to be able to communicate in real time to take advantage of all the available bandwidth without sending more data than can be received at any given time.

All of today's scalable formats work by putting multiple versions of the video and audio data as different streams inside the same file. These scalable solutions are often generically called *multiple bit rate*, or MBR. The details vary and can matter as much as the codecs themselves in providing a good streaming experience.

Choosing a compression format

Depending on the project, picking the correct format can take two seconds or two weeks.

Sometimes the customer mandates the format; other times a compressionist has to make a recommendation. Sometimes the project requires features that are only available in a particular format. Perhaps the requirements haven't been determined. After you choose a format, you need to determine the versions of players that support the format. This section guides you through some questions to ask when trying to decide on a format.

- **Features** Often the biggest determining factor of a platform is the features that users need. If a format can't meet the project requirements, it's obviously not an option. In every project, different features vary in the degree of importance. Compression efficiency is almost always critical. Real-time streaming scalability is another major feature.
- **User base** Depending on the project and audience, user base may be very important or not a concern. With most advertising content, users are unlikely to download a whole new player or codec just to watch someone's marketing clip. But for really compelling content, users jump through enormous hoops. For example, Apple® logged several million downloads of QuickTime in the 72 hours after the *Phantom Menace* trailer was released.

-
-
- **Platforms** Some formats are only available on a single platform or have different features depending on the type, or the version, of the operating system. Most of the modern media players are available for at least Mac OS® and Windows; some are available on Linux® and other operating systems. If very broad compatibility is required, an older standards-based format, like MPEG-1, might be appropriate.

In many cases, especially when you create content for the Internet, it may make the most sense to deploy content in multiple formats so that the user has the choice.

Still, most consumers of digital media have already installed the players for the major formats. While it can be tempting to offer variants in five formats at four data rates each, it is generally better to focus on as few formats as possible and make three great files instead of twenty low-quality files.

Exporting to Windows Media format

Windows Media is the Microsoft digital media technology. Originally introduced as NetShow, it has matured very nicely, and the Windows Media 9 platform, released early 2003, is a very strong proprietary digital media platform.

Using Windows Media for streaming delivery

Windows Media originally focused on real-time streaming, and even though it has become a more diverse technology over time, real-time streaming is still a core focus. Windows Media 9 has an excellent streaming architecture, excelling with MBR technology, servers, and codecs.

The MBR solution in Windows Media is called *Intelligent Streaming*. Prior to Windows Media 9, it was relatively limited, handling up to five video data rate bands per file, but all of these data rate bands needed to share the same resolution and audio track.

Intelligent Streaming was enormously improved with Windows Media 9. Now, it can use over a dozen data rate bands for both video and audio. The video bands can vary resolution as well, making it possible to provide a single link that can scale from audio-only for modem up to high definition playback for a local area network (LAN). The only significant restriction to Intelligent Streaming now is that all bands must share the same codec and encoding mode (for example, Windows Media 9 with a 2-pass CBR).

To use the new Windows Media 9 Intelligent Streaming options, you need to run Microsoft Windows Server™ 2003. If you are making a file that needs to work on the older Windows Media Server 4.1 for Microsoft Windows NT® and Windows 2000, use the same resolution for all video bands and only use a single audio band. Older players (earlier than Windows Media 9) also need to use a subset of the available features.

Using Windows Media for progressive download

Earlier versions of Windows Media weren't suitable for progressive download, but this changed with Windows Media 7.1 and later. Now the players reliably play partially downloaded material, and the encoders support download optimized modes, like 2-pass bit rate limited VBR.

Using Windows Media for disc-based delivery

Windows Media works with CD-ROM and DVD-ROM, as long as Windows Media Player or another Microsoft DirectShow-based player is used. However, because data rates are much less limited on discs, peak data rates that CPUs can decode may limit playback performance. The peak constrained VBR mode can be used in those cases to guarantee performance on the minimum playback platform.

Extensions

There are many different file types and extensions available for Windows Media. Here are some of the most common extensions and their meanings:

.asf

The file format used for Windows Media files is called the Advanced Streaming Format, and the original extension was .asf. However, this extension didn't discriminate between audio-only and video files, and so the extensions .wmv and .wma were added with Windows Media 7. Today, .asf files are typically legacy content using the Microsoft MPEG-4v3 video codec. However, the file format itself hasn't changed.

.wmv

A Windows Media Video file (.wmv) is an .asf file that contains video.

.wma

A Windows Media Audio file (.wma) is the same as an .asf file, but it only contains audio.

.asx

An .asx file is a metafile for Windows Media Streaming. The .asx file goes on the web server and points toward the media on the streaming server. The .asx file format itself hasn't changed, but with the advent of WM7, the .wvx and .wax file extensions have been used more commonly.

.wvx

A .wvx file is a metafile that contains video and, optionally, audio.

.wax

A .wax file is a metafile that contains only audio.

The Macromedia Director CD-ROM authoring environment can't use DirectShow for video playback out of the box. To use Windows Media files in Director, you need a third-party plug-in.

There is a growing number of titles using high definition Windows Media 9 content on DVD-ROM discs. While discs require a fast computer for playback, Windows Media 9 is proving to be the first viable high-definition disc format.

Windows Media players

There is a wide variety of different Windows Media players available for different platforms, making compatibility a moving target. Recent players aren't currently available for some common operating systems, which forces a lowest common denominator that provides multiple versions of the content or excludes some platforms. The following list describes the most commonly deployed players.

Windows Media Player 6.4 for Windows and Windows NT

Windows Media Player 6.4 shipped with Windows 98 and Windows 2000, and it is the most recent version available for Windows 95 and Windows NT. Windows 95 doesn't have an appreciable market share for digital video consumers, but Windows NT is still common on corporate desktops and matters as a delivery platform. Windows Media Player 6.4 is very limited in general, both in performance and features. Connecting to the same content over the same connection, Windows Media Player 6.4 has lengthier buffering times and less reliable playback than the newer versions.

Windows Media Player 6.4 is shipped with Microsoft MPEG-4 version 3 as its most advanced codec. However, the player can automatically update through Windows Media 9. This update requires administrator privileges on the computer and a firewall that allows the download. There is also an enterprise deployment pack that Information Technology departments can use to provide the updated codecs in a network-wide system image. Windows NT systems typically run on older hardware, and so there may be performance issues playing modern codecs on those systems.

Windows Media Player 6.4 can't work with nonsquare pixels or interlaced video like Windows Media Player 9, but those features typically require a faster CPU than Windows NT computers are likely to be running on.

Windows Media Player 7.1 for Windows

Windows Media Player 7.1 for Windows is available for Windows 98, Microsoft Windows ME®, and Windows 2000, and is shipped with Microsoft Windows XP®. Windows Media Player 7.1 for Windows can be automatically or manually updated to the latest codecs, but it is missing a few features of the latest Windows Media Player 9. Specifically, it can't do more than two-channel output (5.1 and 7.1 sources are mixed down to stereo for playback) or interlaced and nonsquare pixel playback. Still, for web video applications, Windows Media Player 7.1 is nearly as good as Windows Media Player 9.

Windows Media Player 9 for Windows

Windows Media Player 9 for Windows runs on the same platforms as Windows Media Player 7.1 and may displace it over time. It only started shipping as a standard application on Windows XP in the summer of 2003.

Windows Media Player 9 supports the full suite of features for Windows Media Video 9 and is the only player with support for the new features like multichannel audio, interlaced video, and nonsquare pixels.

Windows Media Player 9 for Mac OS X (10.1 and later)

In late 2003, Microsoft released an updated Mac OS player with support for the WMV 9 and WMA 9 codecs. Otherwise, the update was quite similar to the older Mac OS players, which lacked support for nonsquare pixels and interlaced video. Performance for lower resolution content is quite good, but even the fastest G5 computers can't play high definition Windows Media Player 9 content yet. Note that none of the Mac OS X players have supported the older, but still common, ACELP.net speech audio codec.

Windows Media Player 7.0 for Mac OS “classic” (Mac OS 8.1-9.x)

The early versions of Windows Media Player for Mac OS were problematic. But with version 7, the Mac OS solution became quite good, to the point where Windows Media became a popular alternative platform for delivering real-time streaming content to Macs.

The current Mac OS 9 solution is similar to the Windows Media Player 6.4 player because it uses the old scripting and Digital Rights Management (DRM) model and supports codecs through Windows Media Video 8, Windows Media Audio, and ACELP.net.

Windows Media Player for Pocket PC 2003

Current Microsoft Windows Mobile Pocket PCs can also play Windows Media with WMV7, 8, or 9, and WMA. Note that many Pocket PCs have slow processors and video performance, so resolution, frame rate, and data rate need to be much lower than for computer playback if you target a wide variety of devices. Ideally, you should target only a particular model so users can tune settings for ideal playback.

For higher-end Pocket PCs, reasonable settings for WMV9 would be 320 x 240 at the native frame rate and 300 Kbps data rate. Any WMA9 audio is fine, up to 48 KHz stereo.

Pocket PC playback is normally local, but the Pocket PC player is capable of real-time streaming if it has a wireless Internet connection.

Windows Media Player for Pocket PC 2002

WMP for Pocket PC 2002 is only supported through WMV8. It also offers lower performance because it was available on lower-performance hardware.

WMP for Pocket PC 2000

WMP for Pocket PC 2000 is similar to that for 2002, but it only supports WMV7, not WMV8. Also, Pocket PC 2000 devices tend to be slower than 2002 devices, which require even more conservative compression settings.

Windows Media encoding modes

Windows Media has a variety of encoding modes, each of which is appropriate for different types of uses. The terminology can be confusing, so it's important to understand the modes that are optimal for the given uses.

Unlike other formats, these data rate modes apply to both audio and video for many codecs. Some codecs are CBR only, however.

CBR

CBR is an abbreviation for *constant bit rate*. CBR is the mode optimized for real-time streaming by use of Windows Media Services CBR. CBR always needs to be used for streaming. CBR is also needed for reliable playback on the legacy Windows Media Player 6.4 for Windows NT and Windows 95.

When encoding CBR, the user defines a buffer size (by default three seconds). Larger buffer sizes increase compression efficiency but also increase latency. The latency increase is minimal with Windows Media Player 7 and later, but the minimum latency with the Windows Media Player 6.4 is always at least as high as the buffer size, even for local playback.

CBR encoded files also use a value of 0-100 that is often labeled as a quality measurement, but is better understood as specifying a trade-off between image quality and providing an accurate frame rate. With a value of 0, no frames are dropped, and the image quality is as good or poor as is possible given the content and data rate. Values greater than 0 set a minimum quality threshold for each frame. If the quality for a given frame drops below that value, its data rate is raised until it meets the quality target. If the data rate of the clip is exceeded, one or more frames are dropped to maintain the data rate. Dropping frames this way can result in a video with a jerky appearance.

1-pass

The default mode for Windows Media encoding is 1-pass CBR. The advantage of 1-pass encoding is that it is twice as fast as 2-pass CBR. However, it doesn't offer the compression efficiency of 2-pass CBR.

2-pass

2-pass CBR mode offers better compression efficiency than 1-pass CBR, especially with highly variable content and larger buffer sizes. Whenever you create CBR files, use 2-pass CBR, if time allows.

Quality-limited VBR

Quality-limited VBR is a 1-pass VBR mode for which image quality is specified instead of the data rate. Instead of requesting a data rate, a quality value from 0-100 is assigned. These values are the same as from the control in CBR for controlling image quality and frame rate. The final data rate of the file depends on the complexity of the content.

Quality-limited VBR is the only 1-pass VBR mode, and it is best used where data rate isn't an important constraint, like making an archive for later use. As a VBR mode, it isn't compatible with streaming or playback in Windows Media Player 6.4. Now that products like After Effects and Adobe Premiere Pro can use WMV files as source files, archiving to WMV with a high file rate can be a good way to compactly store broadcast quality video.

Bit-rate-limited VBR

Bit-rate-limited VBR is a 2-pass encoding mode for file-based use. Unlike CBR, it doesn't have a buffer size or quality limit.

The effect of a bit-rate-limited VBR is essentially the same as a quality-limited VBR. The quality value that provides the target data rate is used. In effect, the first pass simply determines the quality value that would produce the requested file size.

Bit-rate-limited VBR is the optimum mode for file-based projects for which total file size is the limiting factor, mainly affecting progressive download.

Note that VBR-encoded files can have playback issues on Windows Media Player 6.4. They won't crash the player generally, but playback performance can be poor. Use 2-pass CBR instead when targeting this player. The other major drawback of bit-rate-limited VBR is that it can have high-data-rate peaks, which can cause playback issues. Peak-limited VBR addresses that issue.

Peak-limited VBR

Peak-limited VBR is a special type of bit-rate-limited VBR. It works the same as bit-rate-limited VBR, but it has the additions of peak data rate and buffer size parameters. These additions are used to specify a maximum data rate that won't be exceeded within any section of the file up to the buffer size.

Consider peak-limited VBR when playback performance or peak data rate is a concern, like with very high-data-rate files, or playback on slower computers or slower optical drives. There is no drawback to peak-limited VBR versus bit-rate-limited VBR if the peak bit rate is high enough. As consumer electronic devices adopt Windows Media, documentation should be available defining appropriate peak values for different uses.

Windows Media video codecs

Windows Media has a broad set of codecs available, but there is generally only one codec that is the obvious choice for a particular project.

Windows Media Video 9 codec

WMV9 is a leading video codec with excellent results in a variety of uses that scale from modem to high definition. It provides great compression efficiency and good speed for both encode and decode processes.

So far, the biggest drawback for WMV9 is that it is only preinstalled with recent versions of Windows Media Player on Windows and Mac.

WMV9 is the first video codec in Windows Media that supports interlaced and nonsquare pixel video. While not useful for web applications, this codec makes Windows Media much better for consumer electronics and video authoring workflows.

Windows Media Video 8 codec

WMV8 doesn't offer the compression efficiency of WMV9 (maybe 20% lower), nor does it have the nonsquare pixel or interlaced modes. However, it does have the various 2-pass and VBR modes of WMV9. WMV8 is probably the best choice for most mainstream Windows Media projects if the player version isn't known in advance.

The advantage of WMV8 is that it is supported on almost all platforms of interest. These platforms include Windows 95 and Windows NT, Mac OS 9 and X, and Pocket PC 2002. The codec is preinstalled on the Mac players, Pocket PC 2002 and later, and Windows Media Player for Windows 7.1 and later.

Windows Media Video 7 codec

Most players that can support WMV7 also support WMV8 now. The one exception is Pocket PC 2000, which can't upgrade to WMV8.

Microsoft MPEG-4 version 3 codec

The older Microsoft MPEG-4v3 codec is less advanced than WMV8 and WMV9 in compression efficiency. However, it is the most recent codec available by default in the stock Windows Media Player 6.4. It is generally the oldest codec available in modern encoders.

ISO MPEG-4 version 1 codec

The International Organization for Standardization (ISO) MPEG-4 codec is an implementation of the standard simple visual profile of MPEG-4. However, the Windows Media format doesn't support the MPEG-4 file format, streaming, or audio codecs. WMV7 includes the MPEG-4 codec to provide users with access to MPEG-4.

Windows Media Video 9 screen codec

WMV9 screen codec is optimized for screen recordings, like those used in software tutorials. Unlike screen codecs from other vendors, it has an explicit data rate control, which makes it work well for streaming.

Screen codecs work best for source that is like screen animation, in other words, RGB source where most of each frame is identical to the frame before it. It can also work for some types of flat color motion graphics.

Windows Media Screen version 7 codec

Windows Media Screen version 7 is an older version of WMV9 Screen. It has somewhat lower compression efficiency, but otherwise it works similarly to WMV9 Screen.

Windows Media audio codecs

Windows Media also has a variety of audio codecs tuned for different uses.

Windows Media Audio 9 codec

Windows Media Audio 9 codec is the latest version of the WMA encoder, and its output is backward-compatible with the earlier WMA2 decoder. Any version of Windows Media Player still in use should be compatible with WMA9. Now that there are other codecs with the WMA9 prefix, the original WMA is now generally called *WMA Standard*, although it doesn't appear that way in encoding tools. In this document *WMA9* refers to the standard codec.

WMA9 now supports all of the WMA9 bit rate modes. However, you should use CBR mode if compatibility with Windows Media Player 6.4 is required. Note that the bit rate VBR mode has a more limited data rate range of 48-192 Kbps, but it offers CBR, which can range from 5 to 320 Kbps. Also, the quality-limited VBR mode doesn't provide a full 0-100 range, but it offers a few discreet options in the codec data rate pop-up menu.

There is a 0 Kbps option for WMA9 CBR as well. This option is for making audio-free streaming files compatible with the legacy Windows Media Services 4.1 used prior to Windows Server 2003. Intelligent Streaming for that server requires each file to have one and only one audio track, so the 0 Kbps option provides a way to have an audio-free file.

ACELP.net codec

ACELP.net is a low-bit-rate speech codec. It sounds better than WMA9 for speech content below 24 Kbps. However, the decoder isn't available on all platforms—it's missing from both Mac OS X and Pocket PC.

Windows Media Audio 9 Voice codec

WMA9 Voice is a low-bit-rate speech codec and replaces ACELP.net.

WMA9 Voice performs decently with music content at very low bit rates when played in the music or mixed modes. WMA9 Voice is also known as *WMA 9 Narrowband*. It is generally a better choice than WMA9 for music at 20 Kbps and below.

Note that WMA9 Voice is mono and has a CBR only with a maximum sample rate of 22 KHz and a maximum data rate of 20 Kbps. It's only appropriate for low-bit-rate applications.

Windows Media Audio 9 Professional codec

WMA9 Professional is a new codec for multichannel and high-resolution audio. It requires data rates of at least 128 Kbps, so it isn't appropriate for most web content. It was the first codec in one of the major architectures to support more than two audio channels, sample rates higher

than 48 KHz, and bit depths greater than 16 bits. You should use it only with source that requires one of those features given its high-data-rate requirements. High-bit-rate WMA9 does just fine with stereo 16-bit sources at 48 KHz or less.

Multichannel WMA9 Professional plays back on systems with only stereo speakers, combining channels intelligently. It can also fold down well in other cases, like 7.1 to 5.1, or to left, right, and center gaming systems. Still, multichannel is best to use when you anticipate a significant number of listeners who have multichannel audio playback.

WMA9 Professional supports all of the data rate modes of Windows Media 9, including 2-pass CBR and all three types of VBR. As always, you should use 2-pass CBR for streaming and VBR for file-based playback.

Windows Media Audio 9 Lossless codec

The WMA9 Lossless codec is a mathematically lossless codec for archiving audio. Because a lossless compression can't guarantee any particular data rate, it is only 1-pass VBR. WMA9 Lossless works on the same players as WMA9 Professional. Lossless is overkill for delivery and you should only use it for archiving and mastering. It is also popular for digitally archiving audio CD content.

Exporting to RealMedia format

RealMedia is the oldest real-time streaming web format; the pioneering RealAudio shipped in 1995. RealNetworks has continued to enhance it; the updated Real 10 platform was announced in January 2004.

Unlike Windows Media, most RealMedia tools specify the total data rate from which audio data rate is subtracted to leave video data rate. However, some tools don't use this method—it's important to know the mode a particular tool uses.

Using RealMedia for streaming delivery

RealMedia was the first streaming platform. RealMedia does great in a streaming environment. Its MBR solution is called SureStream. Unlike Windows Media, it can mix and match codecs in different streams. All streams must share the same frame size, which can limit scalability.

Using RealMedia for progressive download

While RealMedia is mainly considered a streaming format, it has supported the excellent .rmvb 2-pass VBR progressive mode for several years. These files actually use a different extension than the standard .rm files, indicating that the files won't work when placed on a streaming server.

Using RealMedia for disc-based delivery

The .rmvb files also work fine from CD-ROM or DVD-ROM. While most software tools, like Microsoft PowerPoint, don't allow RealMedia files to be directly embedded in them, the files play well from the disc in the RealOne player. RealMedia is also a good high-definition playback solution. Its codec has lighter decode requirements than WMV9, which allows high-resolution playback on lower-end computers. The new Real 10 also includes multichannel audio support.

RealMedia video codecs

RealMedia supports a smaller set of video codecs than the other formats. The most recent codec is always the best choice.

RealVideo 10 codec

The new RealVideo 10 is a backward-compatible enhancement to RealVideo 9. Compared to the older codec, it offers significantly improved compression efficiency and a control to adjust compression quality and encoding time.

RealVideo 9 codec

The RealVideo 9 codec is an excellent general-purpose video codec. It scales well from modem data rate and mobile device playback on the low end up to high-definition playback at the high end.

RealVideo 9 comes standard with RealOne Player (Windows, and Mac OS X) and is an automatic codec upgrade for RealPlayer 8 (the last version on Mac OS 9).

RealVideo 8 codec

RealVideo 8 was the previous version of RealVideo. RealVideo 9 is almost as common as RealVideo 8 and offers better compression efficiency, so it is generally better to use RealVideo 9 instead.

RealVideo G2 codec

RealVideo G2 was the first modern RealVideo codec. It is only required by old RealPlayer installations.

RealMedia audio codecs

Unlike other formats, a RealAudio codec is defined by a combination of a codec, data rate, sample rate, and channels. Instead of picking a codec and other parameters, the user picks an inclusive combination. Another twist with the RealMedia audio support is that many tools require audio codecs to be specified for video and music content, for both video and audio files and audio-only files. The user picks the one based on the mode the file was encoded with. It's important for the user to pick the correct setting for the mode that is actually being used.

RealAudio 10 codec

RealAudio 10 uses the old RealAudio 8 Stereo Music codec for bit rates below 128 Kbps and AAC-LC for higher bit rates (Sony's ATRAC3 was used for higher bit rates in previous versions). The codec is backward-compatible to RealPlayer 8 at lower bit rates but the player requires an update for higher bit rates.

RealAudio 10 Lossless codec

This codec requires RealPlayer 10. Data rates are estimated to 700 Kbps for 44.1 stereo, so this codec isn't a web or streaming solution.

RealAudio 10 Multichannel codec

This codec is a 5.0 or 5.1 channel codec that requires RealPlayer 10. It requires high data rates, so it is not a streaming solution.

RealAudio 8 Stereo Music codec

RealAudio 8 Stereo Music is the leading general purpose audio codec in the RealMedia format. Its compression efficiency is better than the mono RealAudio Music codec, so it works well even for mono audio and at data rates as low as 12 Kbps.

At some data rates, a "high-frequency" variant of the codec is available. This high-frequency variant provides a broader frequency range. However, with complex audio, it can also result in more artifacts. You need to experiment to find the settings that work well for different types of content.

RealAudio 8 Surround codec

The RealAudio 8 Surround codec is a variant of the RealAudio 8 Stereo Music codec. It isn't actually a true 5.1 or 7.1 system. Instead it stores a stereo pair, but it can maintain Dolby Pro Logic surround sound encoding, which the base RealAudio 8 Stereo Music codec can't. While

this codec doesn't provide the full multichannel experience, it is compatible with a wide range of equipment.

This codec is appropriate if the source is a Dolby Pro Logic surround sound mix. For normal stereo music, the RealAudio 8 Stereo Music codec provides better quality and more flexible data rate options.

RealAudio Speech codec

The RealAudio Speech codec is a relatively standard speech codec; it is intelligible at the lowest data rates and can sound excellent at higher data rates. Unlike most speech codecs, it provides 32- and 64- Kbps options with excellent fidelity for content such as audio books.

RealAudio Stereo Music codec

The RealAudio Stereo Music codec is the original RealAudio stereo codec from the G2 release. It offers lower compression efficiency than the RealAudio 8 Stereo Music codec. This codec provides extensive backward-compatibility.

RealAudio Music codec

The RealAudio Music codec is the original RealAudio mono codec from the G2 release. Even for mono content, the RealAudio 8 Stereo Music codec provides better quality. However, the RealAudio Music codec is available at lower data rates than the RealAudio 8 Stereo Music codec, so it can be useful in some cases for low bit rates.

Exporting to QuickTime format

QuickTime is the oldest of the mainstream digital media formats; it was introduced back in 1991. It predates the Internet and even the CD-ROM drive as standard computer equipment, and it has matured to support a wide variety of applications (although it is less advanced than RealMedia and Windows Media for real-time streaming). The greatest strength of QuickTime is that it supports rich interactivity and rich media.

By default, QuickTime data rate is set as total kilobytes per second, not the total kilobits per second of RealMedia or the kilobits per stream per second of Windows Media. However, many tools that support multiple formats provide the rate in Kbits. For consistency, data rates are listed in Kbits below.

Note this section is about the QuickTime format (.mov), not about the QuickTime architecture, which can author and play back many other formats, including MPEG-4 and AVI.

Using QuickTime for real-time streaming

QuickTime doesn't have a real-time scalability solution, like SureStream or Intelligent Streaming. QuickTime can be a poor fit for real-time streaming on the Internet because data rates are so variable. However, QuickTime can work well in an intranet environment. One advantage is that the streaming server is the free, open source Darwin Streaming server, available for a vast variety of platforms.

QuickTime added streaming support relatively late in its development. Unlike the other formats where files are always streamable, in QuickTime, you have to add hint tracks to files you want to stream. This extra data gives instructions to the server about how to stream the data. Note that the data itself isn't sent, so the file size of a hinted movie is larger than the amount of data actually being sent to the client.

Movie alternates

QuickTime takes a very different approach to scalability from Windows Media and RealMedia. Instead of making a single file that contains multiple versions of the data at different data rates, QuickTime uses multiple files at different data rates, which are switched when the movie begins playing back.

There are several drawbacks to this approach. First, the switching only takes place when the file begins playing, so there is no switching supported during playback. Second, authoring is more complex—only a few advanced compression products, like Adobe Media Exporter, support authoring the alternate movies while encoding. Third, QuickTime relies solely on the user to correctly configure QuickTime preferences to determine the data rate. So if a user sets up QuickTime on a laptop with a broadband connection at work, and then connects by means of a modem from a hotel room, QuickTime tries to transmit the broadband version of the file and either fails dramatically if it is real-time streaming or causes a very lengthy download for progressive files.

For these reasons, it is best to offer users direct links to each media file, even if you also offer a general movie alternates link.

There are some unique features of movie alternates. For one, they work with progressive files, not just streaming files. For another, they can filter in more parameters than just data rate. You can provide different versions for Mac OS and Windows users, for computers using different languages, and for computers with different CPU speeds.

Using QuickTime for progressive download

QuickTime pioneered progressive download, and this format remains the most popular way to distribute web video with QuickTime.

One complexity with progressive QuickTime movies is that they need to be fast-starting movies. Every QuickTime movie has a movie header, which contains information about the file's structure and is required for playback. Whenever you modify a file, QuickTime moves the movie header to the end of the file, which means that it won't start quickly and only the QuickTime logo will play until the entire movie is downloaded. Even editing metadata, like the movie title, can force this change. The simplest way to move the header to the beginning of the file is to perform a "Save As" from within QuickTime Player Pro.

Don't use hint tracks with progressive files because they aren't used, and they increase the file size. Instead, use the Compressed Movie Header mode with progressive files because it makes the header smaller.

Using QuickTime for disc-based delivery

QuickTime is a great solution for CD-ROM and other disc-based architectures. The QuickTime API makes it very easy to add QuickTime playback to applications, and QuickTime is the best supported format inside Director. However, PowerPoint on Windows doesn't natively support QuickTime.

QuickTime video codecs

Given the long heritage of QuickTime for flexibility and support for third-party codec vendors, the range of codecs supported in QuickTime is broader than in the other formats. However, the best codecs are often not provided for free.

Sorenson Video 3 codec

The highest quality codec available for QuickTime today is Sorenson Video 3 from Sorenson Media. It is a very flexible codec and can work in both progressive download and streaming environments. The basic version doesn't include a variety of advanced features that improve compression efficiency and streaming performance.

QuickTime comes with the free, basic version of Sorenson Video 3. However, for professional use, the Sorenson Video 3.1 Professional upgrade is available. Sorenson Video 3.3 Pro offers improved compression speeds, especially on dual-processor computers, and improved quality, especially with real-time streaming protocol (RTSP) streaming. The Pro version is necessary for any commercial QuickTime encoding. The Pro encoder makes files fully compatible with the bundled decoder—no user upgrade is required to view the files.

Sorenson Video 3 Pro options

Sorenson Video 3 has been included with QuickTime since version 5.0.2, which shipped in the summer of 2001.

- **Quick Compress** This option slightly speeds up the encoder at the expense of slightly reducing quality. Because neither effect is significant, you can almost always leave this option deselected.
- **Bidirectional Prediction** To use B-frame encoding in Sorenson Video 3 and add compression efficiency and scalability, choose this option. With scalability, the B-frames can drop if the processor isn't fast enough to play back all of the frames or if the data rate isn't sufficient to stream all the frames by means of RTSP. This option automatically drops frames and reduces

CPU requirements by at least 50% and reduces data rate by about 25%, at the cost of cutting the frame rate in half. This option is especially useful in streaming because without B-frames, the server drops all frames except the keyframes when under bandwidth stress.

So, it would be great to use bidirectional prediction in all cases. However, there is one problem with B-frames—they cause video to be delayed two frames relative to audio. At 29.97 fps, this delay is only mildly distracting, but at lower frame rates, the video can appear dubbed. Also, B-frames can sometimes cause instability in versions previous to QuickTime 5 (this problem appear to have been fixed in QuickTime 6).

There is an AppleScript for Mac to fix the sync in an already encoded file.

- **Rate Control Mode** The unlabeled rate control box for this option has three additional options. The option Neither Feature performs a standard CBR encode.

The Sorenson One-Pass VBR option loosens up the data rate. This is the optimal mode for encoding for RTSP and for most non-RTSP encoded files where peak data rate isn't a concern. Note that if the Streaming option is selected, the VBR window is constrained to the appropriate value for streaming.

The Minimum Quality option sets a minimum quality threshold for each frame on the 0-100 scale. If the Drop Frames To Maintain Data Rate option is selected, the frame rate is reduced when the quality threshold is exceeded, and the data rate is maintained. If this option is not selected, the data rate rises when the minimum quality isn't met, meaning that the final data rate is overshot.

- **Image Smoothing** This option activates a deblocking filter, which improves the quality when it would otherwise be blocky. It doesn't have much effect at higher data rates, and it is automatically deactivated on playback if the computer isn't fast enough to run the operation. Generally, you should activate it.
- **Media Key** The Media Key option assigns a password to the video that the player application needs to provide for the video to play back.
- **Streaming** You should select this option for all files that will be RTSP streamed. The default settings are appropriate for most projects.

Don't modify the Slice Picture Into Packets Of Bytes option. It needs to be modified in sync with the hinting settings, and except in some private networks, it isn't helped by other values.

The Force Block Refresh Every option ensures that every 16 x 16 block of the video is refreshed at the specified interval. This refresh helps with error recovery (any glitch from a dropped packet is erased at least that often). By using this option, you can lower the number of keyframes in the file. This option slightly reduces compression efficiency with static images (which isn't as important), but it doesn't have much of an impact with fast motion (where blocks are naturally refreshed by motion).

- **Masking** Use this option to include a real-time alpha channel in the video. The codec can either use a provided alpha channel, or it can generate one itself based on an internal chroma-keyer. The internal keyer is limited, so using a professional tool like Adobe After Effects to generate the alpha channel can offer you much better results. To use an alpha channel in the source files, you need to set the color to Millions+.

-
-
- **Watermarking SV3.1 Pro** If you want to add a real-time composited watermark to the video, use this option. This watermark can provide much higher quality than adding a watermark to the video source because the fine details of a source watermark are often lost in compression.

The watermark can have opacity, or it can be colored. Note the file format for the watermark is very finicky—carefully read the section in the codec manual describing how to make the file. If you use Photoshop to convert an alpha channel to a Sorenson watermark, a good starting point is to apply the existing alpha channel to a 50% gray background and then flatten the image. Note that the color the codec uses to key the transparency is in the upper right corner.

Sorenson Video codec

QuickTime also includes the older Sorenson Video (no number is listed, but it's version 2) codec. This codec is a completely different codec than the current Sorenson Video 3 because it has lower compression efficiency and quality. You can use it if you require compatibility with versions of QuickTime prior to version 5.

There was an upgrade of the original Sorenson Video codec called Sorenson Video 2.2 Developer Edition, but it is no longer commercially available.

MPEG-4 codec

QuickTime includes a MPEG-4 codec. This codec is available by default to all Mac OS encoders. However, it is only available on a few Windows encoders that have paid a license fee to enable it.

QuickTime supports the Simple Visual specification for playback, and you can use third-party MPEG-4 encoders to make a file that plays back within QuickTime. Third-party encoders generally offer better compression efficiency, with techniques like 2-pass VBR encoding.

Sorenson Video 3.1 Pro offers better compression efficiency than MPEG-4 and better streaming performance. The biggest advantage of MPEG-4 in a QuickTime file is the decoder speed, which is faster than Sorenson Video 3's.

MPEG-4 has been included in QuickTime since QuickTime 6, which shipped summer of 2002.

H.263 codec

H.263 is a standard video conferencing codec that has been included with QuickTime since QuickTime 4. It streams well and provides good streaming performance. It can (rarely) beat the free version of Sorenson Video 3 for streaming, but it isn't as good as Sorenson Video 3.1 Pro.

One unique feature of H.263 is that it only internally uses the three listed resolutions. If a different resolution is picked, it uses the next smallest resolution internally, and scales the video to the correct resolution in real-time. If desired, you can assign the specific resolution here, but this assignment isn't needed in most cases.

The Cycle Intra Macroblocks feature performs a similar function to Sorenson Video 3's Block Refresh and improves streaming performance, especially in lossy networks. It should be active only for RTSP files.

Cinepak codec

Cinepak was the dominant CD-ROM codec for most of the 1990s. It has been included in QuickTime since version 1.5 over a decade ago. Cinepak's hallmark was a very fast decoder—video at 320 x 240 played back well on 60 MHz computers. Still, the simplicity of the encoder results in compression efficiency well below that of modern codecs. Cinepak is one of the few delivery codecs that supports a native grayscale mode, which gives faster decode performance with black-and-white sources.

The main use of Cinepak these days is for high-resolution playback on old, slow computers, like those often seen in K-12 schools. A video at 640 x 480 and 29.97 fps can play without a hitch on a five year old computer albeit at data rates of at least 5000 Kbps.

VP3 codec

On2/The Duck Corporation has been making commercial codecs for a decade. Their more recent VP5 and VP6 codecs aren't available in commercial products, but the VP3 codec is open sourced and is freely available for QuickTime. While it isn't part of the standard install, it's available through component download, which means QuickTime automatically offers to download the codec when it detects content that uses it. To force an install, run the QuickTime Updater utility—the VP3 codec is one of the third-party options.

The VP3 codec doesn't have a native packetizer, so it isn't an appropriate tool for RTSP streaming. Also, the VP3 codec is several years old and has compression efficiency below that of Sorenson Video 3.3 Pro, and also has unpredictable data rates, so it isn't a good fit for web distribution. However, the decoder is faster than Sorenson Video 3, making the VP3 codec a viable option for high-resolution CD-ROM and kiosk work.

Because the VP3 codec is open sourced, different implementations vary. For example, the DirectShow version of the codec for making AVI files on Windows has a few options that aren't currently available in the Mac OS encoder, although the Mac OS player can play files encoded with those features.

The Allow Dropped Frames option in the VP3 codec instructs the encoder to drop frames to maintain data rate. The VP3 codec has a fairly high minimum quality, so when this option is not selected, the target data rate might be substantially exceeded. Still, dropped frames are distracting, so the VP3 codec isn't the best one to use for high compression efficiency.

The Quick Compress mode substantially speeds up the encode with only a minimum hit in quality. Unlike other codecs, it's safe to run the VP3 codec in Quick Compress mode by default. Still, if the time is available, encoding with this mode inactive guarantees the best quality and most accurate data rates possible.

VP3 has a more complex keyframe insertion mode than most codecs. Not only can you set a threshold, but you can specify the minimum and maximum number of keyframes. The default settings for Threshold and Minimum work well for most projects. If your video includes very fast editing, you might want to set a lower minimum, and if you're working with very high frame rates, you can set a higher maximum. Make sure that the keyframe rate set in the encoding application outside of this dialog box is greater than the maximum number of frames.

Animation codec

Apple's Animation codec is an interframe Run-Length Encoded (RLE) RGB codec. Hence, for good compression efficiency, it relies on long horizontal lines of exactly the same color for each frame. The Animation codec doesn't have any data rate control; it only has quality sliders. Hence, it's not appropriate for encoding video originated as natural images where two adjoining pixels are rarely identical. It can do a good job with recorded screen captures, although there are some workflow tricks required to do those well, which are outside the scope of this document.

The Animation codec is lossless when the quality slider is at 100%. At lower levels, it simply makes horizontal lines that are almost identical actually identical. This process can produce poor quality with many sources, which results in visible banding. Note that despite the name,

the Animation codec isn't very good at encoding three-dimensional optimization, where lighting effects prevent regions of identical pixels.

You can also use the Animation codec to store intermediate files among files. The Animation codec stores intermediate files well when using an RGB workflow with RGB native tools, like After Effects, especially with synthetic images. However, the data rates the Animation codec produces are barely better than uncompressed data rates for natural images. Also, some compression tools don't work well with interframe encoded files, so you should use a keyframe in every frame in those cases. The Portable Network Graphics (PNG) and Microcosm codecs provide better compression with identical quality in those cases.

Graphics codec

The Graphics codec is a lossless RGB codec, like the Animation codec. It's different because it offers better compression efficiency, but it only works in 8-bit color. It delivers screen recordings well, but only if you convert the file to 8-bit first, which is a tricky process.

PNG codec

PNG is an RGB-only codec that was originally developed as a patent-free alternative to the GIF format. However, it evolved to have better compression efficiency and support for more than 8-bit index color. It is very efficient at encoding synthetic content with areas of flat color, like screen shots. It isn't useful as a delivery codec for moving video because it is interframe only, and it doesn't have a data rate control mode. PNG supports the Millions+ mode, so it can transport an alpha channel.

PNG's interlaced mode isn't for interlaced fields; instead, you use it for progressive decoding of the image. When you choose this option, the image initially shows a lower resolution proxy while it downloads. This codec does not work well for archiving because it slightly reduces compression efficiency.

PNG is always lossless, but the Filter mode controls the trade-off between encoding time and final file size. For maximum compression efficiency, use the Best option. For maximum speed while still getting some compression, use the Up option.

Microcosm codec

The Microcosm codec is another lossless RGB interframe-only codec. Its big advantage over PNG is its support for 16 bits per channel, so this codec can losslessly archive content when you use it with tools that support 16 bits per channel, like After Effects. This 16-bit support really excels when delivering to film or broadcast, especially when working with Cineon or a source with more than 8 bits per source. The Microcosm codec can also transport an alpha channel.

A free decoder for the Microcosm codec is available through component download, but you need to purchase the encoder.

None16 codec

None16 is a free codec from the creators of Microcosm. It offers 16-bit-per-channel RGB and an alpha channel. However, it is fully uncompressed, so the file sizes are huge.

Photo-JPEG codec

Photo-JPEG is a codec implementation of the standard JPEG still photography codec. It's a lossy codec that supports 8-bit-per-channel 4:2:0 Y'CbCr, or grayscale; however, it is visually lossless at quality settings above 85 for most content (Photo-JPEG lacks an explicit data rate control).

Photo-JPEG can be a good way to deliver content for slide shows, but Motion JPEG is more suitable for most video archiving and transport uses. Photo-JPEG doesn't recognize fields, so it works best with progressive content.

Motion JPEG codec

Motion JPEG is a 4:2:2 variety of JPEG with support for fields. It's the native format of many NLE systems; many of which can losslessly transcode from their own native format to the generic QuickTime Motion JPEG codecs.

There are two types of Motion JPEG codecs: A and B. These originally corresponded to different chipsets supported by early NLE cards. They're functionally equivalent; so it's probably easiest to default to using A. You need to manually set the correct field mode and field dominance (which is field order mislabeled) in the Options dialog box.

Component codec

The Component codec is an uncompressed Y'CbCr 8-bit-per-channel 4:2:2 codec. Because it's uncompressed, it is extremely fast, but data rate requirements can be substantial. It is mainly used as an intermediate format for users with a lot of computer drive space. The Component codec does not support an alpha channel.

None codec

The None codec is an uncompressed RGB. It offers all of the RGB QuickTime color depths, including alpha channels, except for 16-bit-per-channel modes. Unless you need pure speed, you rarely need to use the None codec, given the variety of lossless RGB codecs in QuickTime.

QuickTime audio codecs

QuickTime audio codecs don't have a specific data rate control in their main user interface. Instead, you set the data rate, if controllable, in the Options dialog box.

MPEG-4 codec

The best general purpose audio codec in QuickTime today is MPEG-4 AAC-LC. Encoding is only included by default on Mac OS. Windows users don't have access to this codec in most applications. The MPEG-4 codec lists options for AAC-LC (a general purpose codec) and Code Excited Linear Prediction (CELP), used for low-bit-rate speech. CELP isn't actually available after QuickTime 6.5.

Like MPEG-4 video, MPEG-4 audio is based on the open MPEG-4 standard. You can use a third-party AAC-LC encoder to make a file that you can import into QuickTime and use as a video track. MPEG-4 has a native packetizer and works well for both real-time streaming and file-based delivery.

AAC-LC is a very efficient modern codec. Data rates above 128 Kbps, even for 48-KHz stereo, are overkill for most content; music is good enough to dance to at 48 Kbps and 44.1 KHz mono, and sometimes below those measurements.

The option listed in bold is the recommended option for optimal quality—generally choose this option. The options available depend on whether you selected mono or stereo in the Main Audio Export dialog box—mono gives you higher sample rate options at lower data rates.

Note that the codec automatically converts from the sample rate specified in the Audio dialog box to the rate specified in the Codec dialog box. If encoding to AAC-LC at either 44.1 or 48, set the sample rate in the Audio dialog box to match. If using a lower sample rate, set the sample

rate in the Source dialog box to match that of the source (generally 48 KHz). Avoid doing two sample rate conversions.

The Encoder Quality option sets the quality and speed trade-off for the codec. For a 16-bit source, Better is the appropriate default. Best only helps with sources more than 16-bit and is slower than Better.

QDesign Music 2 codec

QDesign Music 2 (QDM2) was the leading streaming audio codec before QuickTime 6 and MPEG-4. It remains the best codec available by default in QuickTime for Windows. The QDM2 codec has a native packetizer and it is a fine option for streaming.

QDM2 does very well with instrumental music, but it can have objectionable artifacts when encoding speech, even at higher data rates. If it creates artifacts, try reducing channels from stereo to mono, and then try reducing the sample rate. The Pro version and data rates greater than 48 Kbps are required to compress 44.1 stereo with mixed content. The free version of QDM2 is limited to 48 Kbps maximum data rate.

QDesign Music 2 Pro codec

The Pro version of QDM2 offers better compression efficiency, much faster encoding speeds, and data rates greater than 48 Kbps. The basic version of QDM2 supports 8-48 Kbps. The Pro version can compress up to 128 Kbps.

This codec includes suggested sample rate and channel combinations. These suggestions work well for instrumental music but are overly aggressive for content containing voices. You can use the codec in either Quality or Speed mode. The Pro version is much faster than the Basic version, and so you should use the Quality mode.

For most content, you can use the default Attack and Spread settings of 50/50. If the content is compressing with a lot of artifacts, tweaking these settings on a trial-and-error basis can sometimes yield better results.

Qualcomm PureVoice codec

Qualcomm PureVoice is the codec that Code Division Multiple Access (CDMA) cell phones use, and it sounds a lot like a cellular phone. This codec is quite intelligible with speech, but it has difficulty with other types of content. The PureVoice codec supports mono audio only.

The codec works best at 8 KHz, especially with real-time streaming. If you have sufficient data rate to run it at 22 KHz or greater, a different codec would likely provide a better overall experience.

PureVoice has a native packetizer, so it is a fine streaming codec. The packetizer works best at 8 KHz, so you should use that sample rate for streaming projects.

The Qualcomm PureVoice codec has two modes. Full Rate mode offers a 9:1 compression over pulse code modulation (PCM), so 8 KHz is around 15 Kbps. Quality is similar to a cell phone with a good connection. Half Rate mode offers a 19:1 compression over PCM, so about 7.5 Kbps at 8 KHz. While speech should be intelligible, quality is noticeably degraded from the Full Rate mode. The quality is similar to a cell phone with a weak connection.

The Optimize Compression For Streaming option instructs the hinter not to transmit blank data in each audio block when streaming. This option can somewhat reduce real-world data rates, but it performs this reduction unpredictably. You should activate it only for RTSP.

IMA codec

IMA is an old CD-ROM codec created by the Interactive Multimedia Association. It offers a 4:1 compression over the None mode. Quality is good at 44.1 KHz, but it lessens at lower levels. IMA is useful because of its very fast decoder and wide compatibility (ranging back to QuickTime 2.1 and it works in AVI as well). It is a decent option for CD-ROM and kiosk uses where playback performance is more important than compression efficiency.

QuickTime also supports the IMA Adaptive Differential Pulse Code Modulation (ADPCM) variant created by Microsoft. It's effectively the same codec, but it is more compatible with older AVI players.

MP3 codec

While QuickTime doesn't include its own MP3 encoder, it has a very capable MP3 decoder, and it can play MP3 files encoded by a wide variety of tools. Some encoding tools have native support for encoding MP3 as an audio track.

A good MP3 encoder offers better performance for mixed content than any other QuickTime codec, other than MPEG-4. Quality is quite good at 128 Kbps for 44.1 stereo, and compression is transparent for most listeners at 160 Kbps.

QuickTime 4.1 and later can handle VBR MP3 just fine. However, it crashes with QuickTime 4.0, so you should use CBR only if you require compatibility with QuickTime 4.0.

Although there is an official RTSP mode for MP3, QuickTime doesn't support it. You shouldn't use MP3 for RTSP streaming.

None codec

The None codec is an uncompressed PCM audio track. It is rarely appropriate for a delivery codec, but it is the best option for intermediate and transport use. The None codec supports 8 bit and 16 bit, but you should only use 16 bit.

24- and 32-bit integer codecs

QuickTime also supports PCM in the higher sample depths of 24 and 32 bit. These codecs are for transporting between two tools that support these higher sample rates.

Floating Point codecs

The floating point codecs are occasionally used to export audio to some high-end audio or signal processing tools. You wouldn't use these in a normal workflow.

MACE codec

Macintosh Audio Compression Engine (MACE) is a 1980s pre-QuickTime audio codec with very poor 8-bit quality.

Exporting to MPEG-1 format

MPEG-1 was the first open multimedia standard. The original vision for MPEG-1 for computers was that users would buy add-on MPEG-1 decoder boards for their computers for playback. However, this vision didn't become popular, so MPEG-1 playback didn't happen widely until computers became fast enough to perform all of the decoding in software.

VideoCD, a precursor to DVD, also used MPEG-1. While it didn't get used much elsewhere in the world, VideoCD became enormously popular in Asia.

MP3 was also a technology spin-off from MPEG-1.

Using MPEG-1 for progressive download

Although its compression efficiency may only be half that of today's cutting edge web codecs, MPEG-1 still plays a role in progressive downloads where vast compatibility is more important than compression efficiency.

MPEG-1 data rates are higher than other formats. A good option is 256 x 192 video at 336 Kbps coupled with 64-Kbps, 32-KHz mono audio. That data rate produces decent quality with lots of content when coupled with a good 2-pass VBR MPEG-1 encoder. MPEG-1 for the web should always be square pixel.

Using MPEG-1 for real-time streaming

MPEG-1 reasonably supports real-time streaming and includes built-in support in Apple's free Darwin Streaming Server and in a variety of products from Cisco. Note the compression efficiency of MPEG-1 is not as high as other modern codecs.

Using MPEG-1 for CD-ROM delivery

MPEG-1 is often used for CD-ROM work, where decode performance and compatibility is more important than compression efficiency, especially for cross-platform CD-ROM projects where it isn't possible to install QuickTime. MPEG-1 is the only good cross-platform file format for PowerPoint, which can't easily use QuickTime files on Windows or Windows Media on Mac.

MPEG-1 encoded for CD-ROM should be encoded in square pixel because older versions of Microsoft DirectShow handle nonsquare pixels incorrectly.

Using MPEG-1 for VideoCD delivery

The VideoCD format is mainly popular in Asia, but it's also playable on most personal computers without additional software and in DVD players.

Note that VideoCDs, like DVDs, come in PAL or NTSC players. A computer can play both formats, but a set-top DVD or VCD player may only play discs in its own format.

VideoCD has a very tightly constrained set of parameters. Resolution is 352 x 240 in NTSC and 352 x 288 in PAL, and data rate is also locked down. PAL is always 25 fps and NTSC can either be 29.97 fps for video or 23.976 fps for film sources. The data rate is locked down to a CBR rate exactly that of an audio CD, so for every minute of audio a CD can hold, it can hold one minute of VideoCD video.

Note that you can make VideoCD on CD-R, but CD-R VideoCD doesn't work on some early model DVD players due to an incompatibility with the laser. Manufactured CD-ROM media works fine.

There are many VideoCDs in use, either older discs or discs encoded with a cheap real-time or software encoder. Modern professional grade encoders can be quite good. So while the resolution might be low, the image and sound quality can come out nicely with a variety of content.

MPEG-1 video and audio codec

MPEG-1 only supports one video codec, also called MPEG-1. This codec also includes three audio modes: Layer I, Layer II, and Layer III.

MPEG-1 codec (for video)

MPEG-1 offered incredible compression efficiency for its day, and modern encoders can perform well at data rates far below the data rates the format's creators thought possible. MPEG-1's biggest barrier for use is that it lacks postprocessing, so at low data rates, it gets blocky instead of

soft, which results in a relatively high minimum bit rate for decent results. 2-pass VBR substantially helps MPEG-1 and is available in most implementations.

MPEG-1 codec (for audio)

In practice, only the Layer II audio mode is used. Layer I had lower compression efficiency. Layer III never became standard as part of the decoders. However, it continues to live under the name of MP3.

Layer II offers compression efficiency about 50% less than MP3. VideoCD uses 224 Kbps, and generally sounds quite good. Using lower Kbps values can also work well for web content.

Layer II supports joint stereo, which improves compression efficiency, especially at lower data rates. Most Layer II encoders don't default to having joint stereo activated, so you may need to set this manually. Use joint stereo unless you need to preserve Dolby Pro Logic surround sound encoding.

Exporting to DVD and MPEG-2 formats

Measured in eyeball-hours, MPEG-2 is easily the dominant digital-video format worldwide. It's used in virtually all digital broadcast systems (digital cable, satellite, and broadcast), as well as DVD.

One great strength of MPEG-2 is the way it handles progressive and interlaced content. Entire streams can be interlaced or progressive. Interlaced streams can have progressive frames in them, and interlaced frames can have progressive blocks in them. Note that, depending on the player, interlaced MPEG-2 files may or may not be deinterlaced when played back on computers.

The DVD format's MPEG-2 subset doesn't actually support progressive streams. Telecined film on NTSC DVD is stored with the usual 3:2 pulldown pattern. However, by marking the frames that are progressive and those that are interlaced, and by using "field repeat" tags that indicate the 12 fields every second that are duplicates, you can use progressive DVD players to perform a real-time inverse telecine and restore the video to 23.976 fps. When you attach a progressive DVD player to a progressive display, the video can play back at the native frame rate, which improves image quality and smoothness.

Good MPEG-2 encoders automatically take 23.976p or 24p source and make a correct "progressive" DVD out of it.

Using MPEG-2 for DVD delivery

DVDs use the MPEG-2 video codec, but they don't use the MPEG-2 file format. Instead, they use a VOB (Video Object) file, which uses other audio codecs and includes other types of information, such as subtitles.

Many compression tools can't directly export a VOB file. To get around this limitation, you can produce an MPEG-2 video file (called an *elementary*) stream (generally M2V) and a stand-alone audio file (WAV or AC3), and then use an authoring tool to combine them into a VOB file.

Using MPEG-2 for CD-ROM/DVD-ROM delivery

MPEG-2 hasn't become a viable disc-based playback standard because decoders aren't bundled with all personal computers, and you can't legally distribute them without paying a significant licensing fee.

For kiosk applications, MPEG-2 can work well. It has lower processor requirements per pixel

than many other formats, so if disc space isn't a concern, you can use MPEG-2 quite nicely to play back high-resolution, high-quality video.

Working with GOPs on DVD

For DVD, the standard group of pictures (GOP) size for PAL and telecined film content is 12, and NTSC uses 15. The standard allows 15 for the maximum PAL and 18 for the maximum NTSC.

GOPs can be either open or closed. In an open GOP, the last B-frames of the GOP can use the next I-frame for reference. The ability to use the next I-frame gives a small, but welcome, improvement in compression efficiency, but it makes random access a little more complex because each GOP is not wholly self-contained.

For DVD authoring, Open GOP is fine except for multiangle titles. Multiangle titles need Closed GOPs and need to have all the I-frames align among the angles to provide smooth switching among angles.

MPEG-2 audio codecs

While there are default MPEG-2 audio codecs, most implementations use other codecs.

Layer II audio codec

The theoretically standard audio codec used in stand-alone MPEG-2 files is Layer II, the same as MPEG-1. The AAC audio codec that MPEG-4 uses also originated as an MPEG-2 codec, but it is virtually unused in MPEG-2 files. In an MPG or MP2 file used for computer-based playback, Layer II is the only audio codec used.

However, DVDs don't use Layer II; they use other audio codecs that are not part of the official MPEG-2 specification.

PCM codec

Consumer DVD authoring packages use uncompressed PCM audio. The quality is flawless, but the data rates are high. For a one-hour home video, quality is not an issue, but when you try to add a lot of content on a single disc, the Dolby Digital compression efficiency allows you to add more high-quality video to a single disc.

You can use PCM for high sample rates and channels, but it is almost only ever used at 48-KHz 16-bit stereo.

Dolby Digital codec

Dolby Digital, also called AC-3, is the standard for professional DVD audio. Including AC-3 encoding is the hallmark of a professional DVD authoring tool.

The historical standard AC-3 data rate was 384 Kbps, but the DVD standard allows for 448 Kbps. The DVD standard is best for most multichannel projects unless data rate is at an absolute premium. Stereo encoding can use less—192 Kbps is generally sufficient.

DTS codec

Some high-end titles also include the Digital Theater Sound (DTS) audio format. DTS allows for much higher data rates (although these higher rates are partially due to lower compression efficiency—DTS doesn't use psychoacoustics). DTS is generally only used on A-list Hollywood titles, for several reasons. First, the mastering tools are quite expensive and are not bundled with compression tools. Second, DTS decoding isn't mandated as part of the DVD specification, so

DTS titles also include Dolby Digital tracks as well. Having both DTS and Dolby Digital tracks cuts into available disc space, and the audio quality gains are slight at best compared to 448 Kbps AC-3.

Exporting to MPEG-4 format

MPEG-4 is the follow-up ISO standard to MPEG-1 and MPEG-2. Unlike those formats, which were created before the web era and were strictly focused on audio and video playback, MPEG-4 is a much broader and extensible format. The format is only now becoming well known.

Profiles and levels

There is no simple thing called “MPEG-4 compatibility.” The MPEG-4 standard has a massive scope. To manage the support needs of particular products and industries, it uses the concept of profiles and levels. A profile defines the features you can use, like B-frames. A level defines the maximum parameters that you can use with those features, like maximum frame rate, frame size, or data rate.

The following list contains the profiles that handle video compression that are most commonly seen, or at least discussed, for future implementation. Not listed are the rich media and other, more advanced, profiles that are more a matter of authoring instead of compression.

- **Short Header** This profile is the same as a baseline H.263 and is a subset of the Simple Visual profile. So far, Short Header mode is only required by some cell phones. Even a typical PDA can use more aggressive profiles. Still, Short Header can provide reasonable quality at lower data rates on the small screens these devices have.
- **Simple Visual** This profile is the current baseline for computer playback and is the highest video profile supported by QuickTime. It provides compression efficiency better than H.263, but it is less than competitive against the current generation of codecs.
- **Advanced Simple** This profile is an enhanced version of the Simple Visual profile, adding support for interlaced video (at higher levels), B-frames, and other factors that can improve compression efficiency. Most computer players support the Advanced Simple profile, with the exception of QuickTime.
- **Simple Scalable** This profile is the most basic form of scalability for real-time streaming in MPEG-4. Simple scalable has only one enhancement layer, and hence, two possible data rates. It has been adopted in PacketVideo’s suite of tools but not beyond that.
- **Fine Grain Scalability** This profile provides robust scalability for MPEG-4. It combines multiple enhancement layers plus the ability to use only part of a given layer to provide very flexible scaling that is able to optimize itself very tightly with the available data rate.

Unfortunately, no products use it yet. Until Fine Grain Scalability or similar Master Boot Record (MBR) technology is implemented, MPEG-4 won’t be competitive for real-time streaming over the Internet compared to Windows Media and RealMedia.

- **AVC (Advanced Video Coding)** This profile is the big new thing in the MPEG-4 world. Also known as *H.264*, *JVT*, *H.26L*, and *MPEG-4 Part 10*, it is a new codec built from the ground up that is not based on H.263. The hope is AVC will be a codec that lasts a decade, like MPEG-2, so its design emphasizes compression efficiency over performance.

At the NAB (National Association of Broadcasters) 2003 exhibition, developers first demonstrated the AVC encoders. Currently, AVC encoders are in the early stages and started appearing in products in 2004.

-
-
- **Studio** This profile is meant for studio use, like I-frame-only MPEG-2, Motion JPEG, and proprietary capture codecs from many vendors. Using MPEG-4 video technology, the Studio profile offers mastering quality at lower data rates than today's formats. This profile includes 10-bit-per-channel mode, 4:2:2 and 4:4:4, and other features desirable in high-end video authoring.

There aren't any Studio profile products yet, nor are any expected in the near future. It remains to be seen if this profile will become a significant use of MPEG-4.

- **N-Bit** This profile is for high-end uses, such as archiving medical images. MPEG-4 offers the N-Bit profiles that allow encoding in 16 bits per channel and other modes. N-Bit profiles are useful in their own domain, but they are unlikely to be used for compressed video delivery or even video authoring.
- **3GPP profile** This profile, which stands for Third Generation Partnership Project, is designed for mobile phones. There are two types, each with different audio codecs. However, most carriers are mandating a carrier-specific type of 3GPP. So far, it isn't possible to make a single file that will work on all the MPEG-4 playing mobile phones. As the mobile phones become more powerful, this limitation may change.
- **ISMA profiles** The most important group doing work for MPEG-4 for web-based computer playback is the Internet Streaming Media Alliance (ISMA). They published the ISMA 1.0 specification in late 2001 defining two basic profiles (confusingly, not the same as an MPEG-4 profile, but a combination of MPEG-4 profiles and level).

ISMA Profile 0 This basic profile is designed for very low-power mobile devices, like cell phones. So far, these kinds of devices are either using 3GPP or a carrier-specific format.

ISMA Profile 1 This basic profile was originally designed for computer playback, but modern XScale-based PDAs are also becoming fast enough to support it. Profile 1 is the one most commonly quoted for support. Unfortunately, many players don't support the whole specification. CELP support is only available in a few players.

MPEG-4 licensing issues

So far, the biggest barrier to widespread MPEG-4 adoption has been licensing issues. The license terms aren't prohibitive, but they weren't finalized for a long time, so concern occurred around the ambiguity.

Most formats involve some license fees to build encoders and decoders, and MPEG-4 is no different. Product vendors pay these costs; they aren't an issue for the compressionist or distributor. But MPEG-4 includes fees for distributing content in the MPEG-4 format. For disc-based content, replicators have been collecting MPEG-2 fees for making DVDs for years. Fees for distributing content by means of the Internet is new for MPEG-4.

These fees aren't as intimidating as they sound—most distributors won't have to pay a cent. First, the fees only apply to remunerated content, not marketing, advertising, and noncommercial content. For remunerated sites (mainly those doing subscription or add-funded content), the first 50,000 subscribers per year per legal entity are free. After 50,000, the distributor can pay their choice of \$0.25 per subscriber per year or \$0.02/hour of content streamed. These numbers are a small portion of the budget of any streaming project. Still, the simple issue of tracking use has limited adoption in some cases.

DivX

The AVI codec with the most buzz these days is DivX. It doesn't even use the .avi extension; instead, it uses .divx. DivX was originally based on an illegal hack of the Microsoft MPEG-4 version 3 codec, but since version 4, it has been based on standards compliant MPEG-4 for video (Simple in V4, Advanced Simple in V5).

For audio, DivX normally uses MP3s. To avoid confusion with all of the other .avi files out there, DivX has adopted the .divx extension for files, even though the file format itself is pretty much a plain AVI.

The reason it's not considered a plain AVI is that DivX files often contain VBR MP3 audio, where AVI historically didn't support VBR audio. This support has led to playback glitches on some players, especially QuickTime.

MPEG-4 audio codecs

Audio profiles in MPEG-4 combine sets of possible codecs. Because most implementations support particular codecs instead of whole profiles, it's generally more useful to describe the audio codecs instead of profiles.

AAC-LC codec

The audio codec used in most MPEG-4 implementations so far is Advanced Audio Coding-Low Complexity (AAC-LC). The "Low Complexity" aspect is a very high-quality codec. While there are other types of AAC, only one has shown significantly better compression efficiency than AAC-LC in listening tests.

The ISMA standards include AAC-LC, and every computer-based MPEG-4 player supports it.

HE-AAC codec

The new variant of AAC is HE-AAC (the HE stands for High Efficiency). HE-AAC uses a technology called *spectral band replication*, which is essentially a very efficient enhancement layer.

Currently, HE-AAC is the best audio codec by a fair margin. Listening tests demonstrate near transparency compared to CD at 48 Kbps and entertainment quality at 32 Kbps. HE-AAC has yet to become mainstream, although it is expected.

CELP codec

CELP is a low-bit-rate speech codec. Due to licensing issues, it hasn't been deployed in enough players to reach critical mass, so it isn't widely used. Like other similar codecs, it provides intelligible speech at low data rates. There are several other speech codecs in MPEG-4, but CELP is the most frequently used codec. A speech codec is not needed as much in MPEG-4 because AAC-LC performs well with low-bit-rate speech.

AMR codec

The only other speech codec that matters in MPEG-4 is Adaptive MultiRate (AMR). Like CELP, it's a low-bit-rate speech codec. Unlike CELP, AMR supports real-time scalability. AMR is mainly used in the 3GPP profiles for cell phones.

Exporting to AVI format

The Microsoft Audio Video Interleave format was a very early format for CD-ROM, and it is still widely used in a variety of applications.

Using AVI for web delivery

AVI is not a good web format because it doesn't support real-time streaming. Also, the set of codecs that are available by default on all computers is relatively limited. AVI files on the web are almost always distributed as complete files.

Using AVI for disc-based delivery

AVI is still strong for disc-based playback. It's well supported in applications, such as Director, and because codecs can be delivered with a disc, installations are much less of a worry.

AVI video codecs

There are many codecs available for AVI. Most of them require manual installation. The following list contains a few of the most common and useful codecs.

Cinepak codec

The classic AVI codec was Cinepak. As previously described in the QuickTime section, Cinepak

offers limited quality but broad playback compatibility and very fast decoder performance. Cinepak is the only video codec that you can assume is available to all AVI players.

Indeo codec

Intel® created the Indeo codecs, but Ligos® now owns them. Indeo 3, 4, and 5 were very popular in the CD-ROM era. However, they haven't been meaningfully updated in years. They used to be good default codecs, but Indeo playback isn't preinstalled in Windows XP anymore nor is it available for Mac OS X. The Indeo 3, 4, and 5 quality was poorer than modern codecs, especially because they used the YUV-9 color space, which only had one chroma sample per 4x4 block.

VP3 codec

As previously described in the QuickTime section, VP3 is an open source codec originally created by On2. Free encoders and decoders are available for it in a variety of platforms, although it is not preinstalled on Windows or Mac. (AVI playback of VP3 files works by means of QuickTime, which includes VP3 in its automatic component download program.)

The current version of the VP3 codec has a few features not available in the QuickTime version. The Lowest Allowed Quality control sets a minimum threshold for quality. To maintain data rate, you should use the minimum value of 63. The Sharpness control sets a target for whether a frame should have high detail but contain potential artifacts, be soft but contain no artifacts, or be in the middle. The best choice is specific to the source and how it's used.

VP5/VP6 codecs

These new codecs are also from On2 and are built on VP3 with much improved compression efficiency, especially in VP6. They are available as a free download for personal use from www.on2.com.

DivX codec

As described in the sidebar, DivX codecs had an ignominious beginning but have evolved into an important part of the AVI infrastructure. Today's DivX 5 is a good implementation of the MPEG-4 video codec, supporting either a Simple or Advanced Simple profile.

The DivX codec also contains a number of preprocessing features inside the codec itself, which can be useful if you are exporting with a tool that lacks those features.

One unique feature for DivX in AVI codecs is that it supports N-pass encoding, where it successively refines the compression in multiple passes. DivX performs N-pass encoding in a round-about fashion by building a dummy AVI file with the complexity data stored in it for the first pass, and then it uses that dummy file for the second pass. Some tools automate this process, others require you to perform it manually. It's rare when more than three passes continue to improve the compression; even the difference between two and three passes is quite slight.

DivX has a specific psychovisual mode in which an internal model selects the parts of the image that are of greater interest to the human visual system, and then the codec attempts to distribute bits to those parts. Using the Light setting is quite safe; however, the Extreme setting can produce odd results. Still, at very low bit rates, this mode can provide a meaningful improvement with some content.

DivX lets you manually set the threshold for triggering a new natural keyframe. The default of 50 is a great starting point for most content. DivX also allows five levels of trade-off between quality and speed of encoding. Time permitting, always use the Slowest mode.

AVI audio codecs

Audio codecs in AVI don't receive the same attention as the video codecs, but audio codecs are just as important, and there is a vast variety of them. The following list contains a few of the most important.

None codec

The most basic audio codec is None, an uncompressed mode like the QuickTime None codec. It is mainly useful for intermediates, not delivery.

IMA ADPCM codec

The IMA codec is the same as the one in QuickTime, offering a 4:1 compression over the None codec with decent quality. It's the most universal audio codec for AVI and has a very light decoder.

MS ADPCM codec

The MS ADPCM is a slight variant of IMA. It doesn't offer any advantages over IMA, and it's slightly less compatible, so you don't need to use it for authoring.

MP3 codec

The MP3 codec is popular with AVI. However, VBR MP3 audio isn't nearly as widely supported as CBR MP3. Windows ships with a very limited MP3 encoder for AVI that maximizes at 56 Kbps. However, a number of good third-party encoders are available.

Adobe Systems Incorporated • 345 Park Avenue, San Jose, CA 95110-2704 USA • www.adobe.com

Adobe, the Adobe logo, Acrobat, Adobe Audition, Adobe Premiere, After Effects, Illustrator, Imageready, Photoshop are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Mac and Macintosh are trademarks of Apple Computer, Inc., registered in the United States and other countries. Intel and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are the property of their respective owners.

© 2004 Adobe Systems Incorporated. All rights reserved. Printed in the USA.



Tools for the New Work™