

HD Formats

Contents:

Introduction

HD Formats: The Infamous "Table 3"

Video in HD

Audio in HD

Recording Formats for HD

Storage Requirements for HD Capture

Backup Systems for HD Content

Conclusion

Introduction

High-definition television (HDTV) first arrived on the national stage in the late 1980s, but even today only a minority of consumers in the United States and a much smaller minority in other industrialized nations have HDTV systems. However, high-definition (HD) production for video and film is increasing rapidly, as is the installed base of high-definition-capable displays. Consumers are demanding higher-quality content that takes advantage of these better displays. In addition to the content delivered over the airwaves, a significant amount of content will be delivered to the displays through computers. This demand will help to further drive the increasing availability of HD content.

To help you better understand HD content requirements, this article describes basic HD formats and concepts, including audio and video in HD, common HD recording formats, and storage and backup systems for HD. These fundamentals are essential to understanding topics, such as interlaced profiles and interframe encoding, which will be discussed in future articles in this series.

HD Formats: The Infamous "Table 3"

The Advanced Television Systems Committee (ATSC) sets the standards for broadcast in the United States. The infamous ATSC table "le 3" outlines the HD formats in use today. The HD formats from this table are shown below. Because ATSC is a U.S. standard, I've also added the equivalent phase alternating line (PAL) frame rates to the table. PAL frame rates are well defined by the Digital Video Broadcasting (DVB) group and are used in other parts of the world.

Table 1. HD formats from the ATSC table 3 with PAL frame rates added

Resolution	Aspect ratio	Interlaced frames per second (fps)	Progressive fps
1280 x 720	16:9		23.976, 24, 29.97, 30, 59.94, 60
1920 x 1080	16:9	25 (50i), 29.97 (59.94i), 30 (60i)	23.976, 24, 29.97, 30


Video in HD

The following section describes key video elements, such as frame sizes and frame rates, and how they apply to HD.

Frame Sizes

HD source formats are almost always either 1920 x 1080 resolution or 1280 x 720 resolution. Note that a substantial difference exists between the 1080 and 720 standards. The 1920 x 1080 resolution contains 2.25 times more pixels than 1280 x 720 resolution at the same frame rate. This difference substantially increases requirements for processing 1080 content in terms of encoding time, decoding speed, and storage.

Figure 1. Screen resolutions



Progressive vs. Interlaced Frames

The 720 formats are all progressive. The 1080 format has a mixture of progressive and interlaced frame types. Computers and their displays are inherently progressive, whereas television broadcasting has been based on interlaced techniques and standards. For Windows Media, progressive offers faster decoding and better compression than interlaced, and should be used whenever possible. I will provide more information for interlaced profiles in a future article.

For HD terminology, we indicate progressive with the letter "p" and interlaced with the letter "i" by putting the letter between the frame size and the frame rate. There are various ways to express this value. For example, a frame size of 1080 pixels and an interlaced frame rate of 60 fps would be expressed as 1080/60i. For progressive content, the frame rate indicates frames per second. For interlaced content, it indicates fields per second. So, 60i is 30 frames per second, each with two fields.

Frame Rates

HD frame rates are 24, 25, 30, or 60 fps, with optional National Television Systems Committee (NTSC) variants that run at a 0.1 percent lower frame rate, to match the 30 fps versus 29.97 fps difference. Note that those derivatives are best calculated as 30000/1001 (commonly written as 29.97), 24000/1001 (commonly written as 23.976), and 60000/1001 (commonly written as 59.94).

Frame rates also have a substantial effect on performance. A value of 60p requires 2.5 times more bandwidth and processing power throughout than a value of 24p.

8-Bit vs. 10-Bit Formats

Video formats are either 8-bit per channel or 10-bit per color channel. The 8-bit formats have 256 steps from black to white, and 10-bit has 1024 steps. This extra detail can improve video quality when the source is heavily processed during editing or post-production. Today, the final Windows Media file is 8-bit, so there isn't any need to capture HD content in 10-bit format if the video is going to be compressed as a Windows Media file.

There are many resources that describe color spaces and sampling techniques. In fact, it's not hard to write entire books about this topic. Color information can be expressed in RGB values, whether the color is at the point of capture (for example, image sensor) or at display (for example, CRT monitor). Another color space, known as YUV or Y'CbCr represents brightness, or luminance (Y), and the color difference signals (Cb,Cr). This Y'CbCr color space is better suited to describe the way color is perceived.

The eye is more sensitive to brightness in green tones, less sensitive in red tones, and least sensitive in blue tones. The Y'CbCr color space can take advantage of that fact by allocating more bandwidth for Y, and less bandwidth for Cb and Cr. When color is sampled at 8 bits or 10 bits per sample per component, depending on the technology, it can use varying amounts of bandwidth for each component.

For example, video may use one of the following sampling rates:

- **4:4:4.** Y, Cb, and Cr are sampled equally along each field or frame line with one Cb sample and one Cr sample for each Y sample. High-end video editing systems use 4:4:4, often together with uncompressed video.
- **4:2:2.** Y is sampled at every pixel, but Cb and Cr color information is only sampled at every other pixel. While this sampling rate significantly reduces the bandwidth requirements, only a slight loss of color fidelity occurs. Many editing systems process this level of color information.
- **4:1:1.** Y is sampled at every pixel, but the Cb and Cr color information is only sampled at every fourth pixel, saving even more bandwidth. This sampling rate is used in consumer DV cameras and is currently the default sampling rate for the interlaced mode in the Microsoft Windows Media® Video 9 codec.
- **4:2:0.** This sampling rate uses a different technique known as a spatial sampling that takes a 2 x 2 square of pixels. 4:2:0 is used by default in the progressive mode of the Windows Media Video 9 codec and by most HD tape formats, as described in the "Storage Requirements for HD Capture" topic later in this article.

One other complication is that the values of Y, Cb, and Cr are displayed somewhat differently with HD formats than with standard definition (SD) video formats. Specifically, HD formats use the 709 color format instead of the Consultative Committee for International Radio (CCIR) 601 color format. I will describe how to deal with this difference in color formats in a future article.

Audio in HD

HD can be authored with virtually any audio format. The old standards of 48 kHz 16-bit audio are used in many cases. Higher bit depths like 20-bit and 24-bit are common, and higher sampling rates like 96 kHz occur as well.

HD is typically mastered in multi-channel, like 5.1 (five speakers plus a subwoofer) or 7.1 (seven speakers plus a subwoofer). Most HD tape formats support at least four channels of audio, and many support eight channels. The Windows Media Audio 9 Professional codec can compress streams that are up to 24 bit and 96 kHz at data rates well below the current audio standards. For more information about audio in HD, see the upcoming articles about 5.1-channel audio production and the How-To Create section on the Windows Media 9 Series Web site.

Recording Formats for HD

Today, the HD industry uses a variety of digital tape recording formats for professional HD production, including formats developed by Sony and Panasonic. These formats use the existing physical tape formats of earlier standard definition formats, but with new compressed bit streams. Fortunately, HD is a relatively new technology, and no significant archive of analog HD tape content exists. With HD, the dream of a wholly digital workflow has been realized.

The following digital tape formats are commonly used for HD recording:

- **HDCAM.** The Sony HDCAM format supports 1080 resolutions at frame rates of 24p, 25p, 50i, and 60i. HDCAM stores the video at 1440 x 1080, which is a 33 percent reduction horizontally from 1920. It also uses a unique color sampling of 17:6:6, which means that HDCAM has only half the color detail of other HD formats. HDCAM is 4.4:1 compressed and is 8-bit, but supports 10-bit input and output.
- **D5-HD.** The Panasonic D5-HD format uses the D5 tape shell. Unlike HDCAM, D5-HD can do 720/60p, 1080/24p, 1080/60i, and even 1080/30p. D5-HD compresses at 4:1 in 8-bit mode and 5:1 in 10-bit mode, and supports 8 channels of audio.

- **DVCPRO-HD.** This Panasonic HD format, sometimes called D7-HD, is based on the same tape shell used for DVCAM and DVCPRO. D7-HD does 720/60p and 1080/60i, with 1080/24p in development. It uses 6.7:1 compression, and supports 10-bit input and output per channel. DVCPRO-HD supports 8 channels of audio.
- **HDV.** This format is one of a number of emerging formats that are being used in lower-cost cameras. HDV was introduced with JVC's groundbreaking professional consumer (prosumer) HD camera, the JY-HD10, which records highly compressed MPEG-2 on a mini DV tape.

HDV is a MPEG-2 transport stream that includes a lot of error correction. Its video uses interframe-compressed MPEG-2, at 19 megabits per second (Mbps) for 720p and 25 Mbps for 1080i. Audio is encoded with 384 Kbps MPEG-1 Layer 2 stereo. The interframe encoding enables HDV to achieve good quality video at lower bit rates, which means much more content per tape, but it increases the difficulty of editing the content. The next article in this series will provide additional details about interframe encoding.

In addition to digital tape recording formats, there are new RAM-based storage media such as DVD-RAM and the JVC JY-HD10. For more information about recording formats, see the "Table of HD VCR formats" page at the Video Expert Web site .

Storage Requirements for HD Capture

Playing back HD content is relatively easy, but getting that content onto the computer can require large quantities of fast storage. The following section provides an overview of the storage issues for HD-sourced content and how to approach developing reliable, cost-effective storage solutions.

The Storage Numbers

How large? How fast? Let's look at some storage numbers, which are broken down into compressed and uncompressed 10-bit formats in the typical color spaces. In the following table, data rate is measured in Mbps, which is 1,000,000 bits per second. Note that many drives will measure the data rate in Megabytes (MBps) per second, which is 8 times higher than Mbps. A gigabyte (GB) is 1,000,000,000 bytes. Many applications, though, assume it is 2³⁰, which is approximately 7 percent larger and is called a Gibibyte, or GiB.

Table 2. Storage requirements for compressed and uncompressed data

Format	Data rate (Mbps)	GB per hour of video
Compressed formats		
720 x 480 DV 4:1:1	25	11
720 x 480 DV50 4:2:	50	22
HD-DVCPRO	100	44
Uncompressed 10-bit formats		
720 x 486 4:2:0	210	92
1280 x 720/24p 4:2:0	332	146
1280 x 720/60p 4:2:0	818	364
1920 x 1080/24p 4:2:0	746	328
1920 x 1080/60i 4:2:0	932	410

There are significant differences between digital formats. Uncompressed data rates are many, many times higher than the rate for a typical file compressed for HD delivery by using Windows Media 9 Series codecs. Windows Media files can also have data rates far lower than MPEG-2 based formats.

Speed

The numbers in the previous table represent the real-world throughput needed to write a single stream of video to the hard disk in real-time. If the hard disk can't keep up with the video stream, frames of video are dropped. Many editing systems have higher requirements, because they need bandwidth to read more than one video stream at the same time for real-time effects. This kind of performance is not required for capture only. To get the kind of bandwidth required for capture, use multiple hard disks as part of a redundant array of independent disks (RAID) solution. For more information about RAID, see the "RAID" topic later in this article.

Storage Capacity

Make sure enough storage is available on your hard disks for the intended use. Acquiring a hard disk is a lot easier than editing. Because HD requires so much storage, editing an offline version at standard definition and then only capturing the required sections in HD for finishing, can be a great option. Although a single 250-GB hard disk can now be purchased, a single disk can quickly become too small. In this case, a multidisk RAID solution is recommended.

SCSI vs. IDE

The debate regarding high performance storage between small computer system interface (SCSI) and AT Attachment (ATA) for IDE drives has been going on for more than a decade. Historically, SCSI drives have been faster than ATA drives, and ATA drives required more CPU power to run. These days, Serial ATA has largely caught up with SCSI. I've seen 1920 x 1080/60i 10-bit format working off a very fast IDE RAID. Because the cost of fast IDE drives is so much less than SCSI, and cabling is easier with Serial ATA, I expect the industry to continue the transition to ATA even in the HD world.

Either way, HD performance requires a stand-alone RAID controller card, and potentially several cards for throughput. The onboard RAID controllers available on many motherboards aren't up to the task of 1080 uncompressed video yet.

RAID

Regardless of whether you're talking about SCSI or IDE/ATA, the key to making your storage solution work is RAID. The basic idea is to tie multiple disks together into a single volume. This makes multiple disks look like one, and increases the storage capacity. It also enables writing to multiple physical disks at the same time, which increases throughput.

Redundancy means that data is duplicated on more than one disk, letting a single volume survive the failure of another volume. After a disk fails, you must rebuild the array, potentially taking it out of commission for a period of time, but in most cases, no data is lost.

There are many different flavors of RAID, which combine drives in different ways. The following RAID configurations are most relevant for HD video work:

- **RAID 0.** Two or more drives are combined, with the data split between them. RAID 0 isn't really redundant. If one drive fails, all the data in the set is lost. However, all drives are used for storage and performance is very fast. If source files are backed up, the RAID 0 performance and price can be worth its fragility.
- **RAID 5.** Three or more drives are combined, with one entire drive's parity data used for redundancy. The volume can survive the failure of a single disk. RAID 5 doesn't perform quite as well as RAID 0, but offers better fault tolerance.
- **RAID 5+0.** Two or more RAID 5 sets are combined into one RAID 0. Substantially more storage is available and redundancy is maintained (one drive per RAID 5 set can fail without data loss). This RAID configuration is the optimal mode for most HD uses.

For more information about RAID, see the Redundant Array of Inexpensive Disks (RAID) page at the PC Guide Web site .

Shared Storage

What happens when you don't want to tie up your expensive HD authoring station for hours of compression? Or you are editing HD content on a Macintosh and you need to transfer the source to a Windows-based computer for compression? To help you resolve these dilemmas, consider the following methods for shared storage:

- **Ethernet.** Today, the simplest method for moving files is through the Ethernet. Transferring files between computers running the Microsoft Windows® operating system is easy by using the server message block (SMB) file sharing protocol, as it is from computers running the Mac OS X 10.3 operating system to Windows-based computers. Earlier versions of Mac OS X will need to use an FTP client like Fetch to transfer files larger than 2 GB to Windows-based computers.

Ethernet speeds vary widely. For HD work, you'll want at least a one-gigabit Ethernet, which can realistically achieve several hundred Mbps of bandwidth, and close to 1000 Mbps with good cables and a fast switch.

- **SAN.** When a workgroup wants to share the same content for editing, not just transfer content, a Storage Area Network (SAN) is the solution of choice. This network enables multiple computers to access the same storage as if it were a local drive on each computer. This means HD content can be captured straight to stored media, without having to transfer the content after it is captured. SANs are expensive, but work well.
- **Sneakernet.** The final solution for transport is the venerable "sneakernet." This method involves transporting physical media, be it an IEEE 1394 drive, a tape, or a whole RAID array from one computer to another. For large quantities of data, transporting the physical media can often be less time consuming than sending the data by using a fast connection.

Backup Systems for HD Content

One huge problem with HD content is how to back it up. The sheer volume of content can be overwhelming on computer systems. The following section describes common backup methods.

Tape

It's certainly possible to use traditional tape backup systems to back up the raw HD data. And unlike compressed video, uncompressed video will actually get some use out of the hardware compression in modern tape decks.

Whereas personal backup systems clearly aren't up to the task, industrial-grade formats can have a lot of storage. The Sony SAIT-1 format stores a massive 500 GB of uncompressed data per tape (although you could buy a pretty nice car for the price of a complete backup system). At a more reasonable price point, Ultrium 200-GB tapes can be purchased for around \$100, with tape decks costing around \$4,000.

For the average facility, the simplest way to back up HD content is to employ the same method used to deliver the HD content—tape decks. The best news is that most HD facilities already use this method. Lightly compressed formats like D5 HD don't entail a drop in quality when going back to tape, although lossier formats like HDCAM do lose some detail when coming back from an HD nonlinear editing (NLE) tool.

Hard Disk Drives

Although hard disk drives aren't likely to be as stable for long-term storage as tape, they're inexpensive and useful. One option is to use a RAID configuration with removable disks and archive the entire set of disks when a project is finished. For an even more cost effective solution, data can be copied to IEEE 1394 or USB 2.0 drives. Drives are much faster to restore than tape backup systems, and they don't require the expense of buying a tape deck.

Conclusion

Hopefully this article has provided you with useful background information about HD formats, color sampling, and storage and transfer possibilities.