DVC Pro FAQ from www.adamwilt.com/DV.html

DV, DVCAM, DVCPRO

What is DV?

DV is an international standard created by a consortium of 10 companies for a consumer digital video format. The companies involved were Matsushita Electric Industrial Corp (Panasonic), Sony Corp, Victor Corporation of Japan (JVC), Philips Electronics, N.V., Sanyo Electric Co. Ltd, Hitachi, Ltd., Sharp Corporation, Thompson Multimedia, Mitsubishi Electric Corporation, and Toshiba Corporation. Since then others have joined up; there are now over 60 companies in the DV consortium.

DV, originally known as DVC (Digital Video Cassette), uses a 1/4 inch (6.35mm) metal evaporate tape to record very high quality digital video. The video is sampled at the same rate as D-1, D-5, or Digital Betacam video -- 720 pixels per scanline -- although the color information is sampled at half the D-1 rate: 4:1:1 in 525-line (NTSC), and 4:2:0 in 625-line (PAL) formats. (See below for a discussion of color sampling.)

The sampled video is compressed using a Discrete Cosine Transform (DCT), the same sort of compression used in motion-JPEG. However, DV's DCT allows for more local optimization (of quantizing tables) within the frame than do JPEG compressors, allowing for higher quality at the nominal 5:1 compression factor than a JPEG frame would show. See Guy Bonneau's discussion of DV vs MJPEG compression for more details.

DV uses **intraframe** compression: Each compressed frame depends entirely on itself, and not on any data from preceding or following frames. However, it also uses adaptive **interfield** compression; if the compressor detects little difference between the two interlaced fields of a frame, it will compress them together, freeing up some of the "bit budget" to allow for higher overall quality. In theory, this means that static areas of images will be more accurately represented than areas with a lot of motion; in practice, this can sometimes be observed as a slight degree of "blockiness" in the immediate vicinity of moving objects, as discussed below.

DV video information is carried in a nominal 25 megabit per second (Mbps) data stream. Once you add in audio, subcode (including timecode), Insert and Track Information (ITI), and error correction, the total data stream come to about 36 Mbps. Roger Jennings' paper on the Adaptec website runs through the detailed numbers.

What's the difference between DV, DVCAM, and DVCPRO?

Not a lot! The basic video encoding algorithm is the same between all three formats. The VTR sections of the US$20,000 DVCAM DXC-D130 or US$17,000 DVCPRO AJ-D700 cameras will record no better an image than the lowly DV format DCR-VX1000 at US$4,000 (please note: I am *not* saying that the camera section and lens of the VX1000 are the equals of the high-end pro and broadcast cameras: there are significant quality differences! But the video data recorded in all three formats is essentially identical, though there may be minor differences in the actual codec implementations). A summary of differences (and similarities) is tabled in Technical Details.

The consumer-oriented DV uses 10 micron tracks in SP recording mode. Newer camcorders offer an LP mode to increase recording times, but the 6.7 micron tracks make tape interchange problematic on DV machines, and prevents LP tapes from being played in DVCAM or DVCPRO VTRs. Sony's DVCAM professional format increases the track pitch to 15 microns (at the loss of recording time) to improve tape interchange and increase the robustness and reliability of insert editing. Panasonic's DVCPRO increases track pitch and width to 18 microns, and uses a metal particle tape for better durability.

DVCPRO also adds a longitudinal analog audio cue track and a control track to improve editing performance and user-friendliness in linear editing operations.

How good are the DV formats compared to other formats?

DV formats are typically reckoned to be equal to or slightly better than Betacam SP and MII in terms of picture quality (however, DV holds up better over repeated play cycles, where BetaSP shows noticeable dropout). They are a notch below Digital-S and DVCPRO50, which are themselves a (largely imperceptible) notch below Digital Betacam, D-1, and D-5. They are quite a bit better than 3/4" U-matic, Hi8, and SVHS.

On a scale of 1 to 10, where 1 is just barely video and 10 is as good as it gets, I would arrogantly rate assorted formats as follows:

D-5 (10-bit uncompressed digital)	10
D-1 (8-bit uncompressed digital)	9.9
Digital Betacam, Ampex DCT	9.7
D-9 (Digital-S), DVCPRO50	9.6
DV, DVCAM, D-7 (DVCPRO), Digital8	9
MII, Betacam SP	8.9
1" Type C	8.7
3/4" SP	6.5
3/4", Hi8, SVHS	5
Video 8, Betamax	4
VHS	3
EIAJ Type 1, Fisher-Price Pixelvision	1

[I had previously placed D-2 and D-3 uncompressed composite digital formats just below BetaSP, lower than any of the component formats. My feeling was that while D-2 and D-3 are excellent first-generation formats for composite analog playback and NTSC broadcast, the compositing of color with luma (which includes a color bandwidth limitation even more severe than DV or BetaSP employ) makes clean multigeneration and multi-layer image compositing problematic at best (even such simple things as adding titles).

However, I was severely upbraided by several folks with extensive digital composite experience, who all rated D-2 and D-3 between DV and DigiBeta. If you've got a high-end all-digital postproduction chain, the quality in these formats holds up over multiple generations extremely well, much better than any analog format, be it component or composite. While this is certainly true, if you don't have that all-digital pathway, I'm doubtful about how they would fare... so assume that D-2 and D-3 fall somewhere in the range between 1" and DigiBeta, and go have a look for yourself!

What are the DV artifacts I keep hearing about?

DV artifacts [Pix: Artifacts] come in three flavors: mosquito noise, quilting, and motion blocking. Other picture defects [Pix: Defects] encountered are dropouts and banding (a sign of tape damage or head clogging).

The most noticeable spatial artifacts are *feathering* or *mosquito noise* around (typically) diagonal fine detail. These are compression-induced errors usually seen around sharp-edged fine text, dense clusters of leaves, and the like; they show up as pixel noise within 8 pixels of the fine detail or edge causing them. The best place to look for them is in fine text superimposed on a non-black background. White on blue seems to show it off best. The magnitude of these errors and their location tends to be such that if you monitor the tape using a composite video connection, the artifacts will be masked by dot-crawl and other composite artifacts.

A spatial *quilting* artifact can also be seen on certain diagonals -- typically long, straight edges about 20 degrees off of the horizontal. These are minor discontinuities in the rendering of the diagonal as it passes from one DCT block to the next; so minor that they're usually invisible. Watching such diagonals during slow pans is often the only way to see the artifact.

Motion blocking occurs when the two fields in a frame (or portions of the two fields) are too different for the DVC codec to compress them together. "Bit budget" must be expended on compressing them separately, and as a result some fine detail is lost, showing up as a slight blockiness or coarseness of the image when compared to the same scene with no motion. Motion blocking is best observed in a lockdown shot of a static scene through which objects are moving: in the immediate vicinity of the moving object (say, a car driving through the scene), some loss of detail is seen. This loss of detail travels with the object, always bounded by DCT block boundaries. However, motion blur in the scene usually masks most of this artifact, making this sort of blocking hard to see in most circumstances.

Finally, *banding* or *striping* of the image occurs when one head of the two on the scanner is clogged or otherwise unable to recover data. The image will show 10 horizontal bands (12 in PAL countries), with every other band showing a "live" picture and the alternate bands showing a freeze frame of a previous image or of no image at all (or, at least in the case of the JVC GR-DV1u, a black-and-white checkerboard, which the frame buffers appear to be initialized with). Most often this is due to a head clog, and cleaning the heads using a standard manufacturer's head cleaning tape is all that's required. It can also be caused by tape damage, or by a defective tape. If head cleaning and changing the tape used don't solve it, you may have a dead head or head preamp; service will be required.

This sort of banding dropout occurs fairly often; about once per DV tape in my experience. Usually it isn't even noticeable -- a single frame of banding due to a momentarily clogged head won't be visible unless there's motion in the scene to show off the frozen stripes. Have a look through your old tapes frame by frame (on a slow day, of course!) and you might be surprised how often you'll be able to find a single, subtly banded frame. For what it's worth, I've never found such a banded frame on any DVCAM footage I've shot, which indicates to me that DV is right on the edge of reliability. DVCAM, with its 15 micron track width, or D-7 (DVCPRO) with its 18 micron track, are sufficiently on the safe side of the bleeding edge so that this sort of droput is much less likely to occur.

Bear in mind that analog BetaSP typically has several dropouts per *minute;* the last time I measured visible dropout rates on Hi8 and S-VHS I got numbers in the range of a dropout every 3-5 *seconds* (Hi8) and every 7-20 seconds (S-VHS). One visible dropout per hour-long tape, on average, is not something to get flustered about. But if it does bother you, shoot DVCAM or D-7 instead.

4:2:2, 4:1:1, 4:2:0

What are 4:2:2, 4:1:1, and 4:2:0 anyway? [Pix: Sampling]

These are all shorthand notations for different sampling structures for digital video. They are also used for CIF and QSIF and suchlike MPEG frame sizes, but in the discussion that follows, I focus on the numbers for SDTV (standard-definition TV) digitized to the ITU-R BT.601 standards: 13.5 MHz sample frequency and 720 pixels per line.

The first number refers to the 13.5 MHz sampling rate of the luma: "4" because (a) it's nominally almost approximately sort of four times the NTSC and/or PAL color subcarrier frequencies, and (b) because if it's "4" the other numbers can be integers whereas if it were "1" the formats would be "1:0.5:0.5", "1:0.25:0.25", and "1:0.5:0" respectively, and which would you rather try to read off in a hurry? The 13.5 MHz sampling yields 720 pixels per scanline in both 525/59.94 and 625/50 systems (NTSC and PAL/SECAM). This number applies to D-1, D-5, Digital Betacam, BetaSX, Digital-S, and all the DV formats just the same.

The other two numbers refer to the sampling rates of the color difference signals R-Y and B-Y (or, more properly in the digital domain, Cr and Cb)

In **4:2:2** systems (D-1, D-5, DigiBeta, BetaSX, Digital-S, DVCPRO50) the color is sampled at half the rate of the luma, with both color-difference samples co-sited (located at the same place) as the alternate luma samples. Thus you have 360 color samples (in each of Cr and Cb) per scanline.

In **4:1:1** systems (NTSC DV & DVCAM, DVCPRO) the color data are sampled half as frequently as in 4:2:2, resulting in 180 color samples per scanline. The Cr and Cb samples are considered to be co-sited with every fourth luma sample. Yes, this sounds horrible -- but it's still enough for a color bandwidth extending to around 1.5 MHz, about the same color bandwidth as Betacam SP (which, were it a digital format, would be characterized as 3:1:1).

So where does **4:2:0** (PAL DV, DVD, main-profile MPEG-2) fit in? 4 x Y, 2 x Cr, and 0 x Cb? Fortunately not! 4:2:0 is the non-intuitive notation for half-luma-rate sampling of color in both the horizontal and vertical dimensions. Chroma is sampled 360 times per line, but only on every other line of each field. The theory here is that by evenly subsampling chroma in both H and V dimensions, you get a better image than the seemingly unbalanced 4:1:1, where the vertical color resolution appears to be four times the horizontal color resolution. Alas, it ain't so: while 4:2:0 works well with PAL and SECAM color encoding and broadcasting, interlace already diminishes vertical resolution, and the heavy filtering needed to properly process 4:2:0 images causes noticeable losses; as a result, multigeneration work in 4:2:0 is much more subject to visible degradation than multigeneration work in 4:1:1.

"Now how much would you pay? But wait, there's more!" In *US* implementations of 4:2:0, the color samples are supposed to be vertically interleaved with luma, whereas in *European* 4:2:0 they're supposed to be co-sited. Practically speaking, this is a headache for developers of codecs, encoders, and DVEs, but for DV purposes it's not especially exciting, since only European DV is 4:2:0.

1394/FireWire

What is 1394 and/or "FireWire"?

IEEE-1394 is a standard communications protocol for high-speed, short-distance data transfer. It has been developed from Apple Computer's original "FireWire" proposal (FireWire is a trademark of Apple Computer). Check out the 1394 Trade Association, white papers on Adaptec's website, and DVCentral's links for pointers to additional 1394 sites for detailed information.

Sony calls their implementation of 1394 "i.LINK".

Why are DV and 1394 always discussed together?

They appear to have been developed together. The data stored on DV tape appear to reflect the packet structure sent across a 1394 link to a frightening degree of exactness. Certainly the DV format and 1394 High Performance Data Bus co-evolved, such that the first consumer DV camcorder in the USA (the Sony DCR-VX1000 and its single-chip brother the VX700) was also the first 1394-equipped consumer product available.

What does a 1394 connection do for me?

Plenty of good things:

- You can make digital dubs between two camcorders or VTRs using 1394 I/O, and the copy will be identical to the original.
- You can do cuts-only linear editing over 1394, with no generation loss.
- You can stick a 1394 board into your computer (PC or Mac), and transfer DV to and from your hard disk. If your system can support 3.6 MBytes/sec sustained data rate -- simple enough with many A/V rated SCSI-2 drives and with most ATA/EIDE drives these days -- the world of computer-based nonlinear editing is open to you without paying the quality price of heavy JPEG compression and its associated artifacts, or the monetary price of buying heavy-duty NLE hardware and banks of RAID-striped hard drives.

Is 1394 that much better than Y/C or component analog?

Yes. A 1394 dub is a digital copy. It's *identical* to the original. That's *really* nice.

Yes, you can do *almost* the same thing with a SMPTE 259M SDI (serial digital interface) transfer. But VTRs with SDI cost big money. 1394 is built into many low-end cameras and VTRs, and the connecting cable -- even at Sony prices -- is only US$50; you can find it for US$20 if you shop around.

Also, transferring via 1394 is a digital copy, a data dump (as it is over the expensive SDTI interface on high-end DVCAM and D-7 VTRs). No decompression or recompression occurs. Transferring DV around as baseband video, even digitally over SDI, subjects it to the small but definite degradation of repeated decompression/recompression.

If a digitally-perfect copy is a 10, and a point-the-camera-at-the-screen-and-pray transfer is a 1, here's how DV picture quality holds up over different transfer methods:

IEEE-1394, SDTI	10
SDI	9.8
Analog Component (Y, R-Y, B-Y)	9
Y/C ("S-video")	8
Analog Composite	5
Point camera at screen and pray	1

What's the deal with DVCPRO gear and 1394?

DVCPRO, or D-7, is a DV-based format with a few subtle differences in its datastream. These changes were made by Panasonic's engineers to improve the robustness and reliability of the DVCPRO system when compared to DV, but they do mean that certain data header bits do not conform to Blue Book standards. Thus a direct data interchange between DVCPRO gear and DV/DVCAM gear is not possible in the same way that DV and DVCAM gear can interchange data; furthermore some nonlinear editor systems are not capable of accepting or generating a D-7-compatible signal.

As a result, DVCPRO gear with 1394 connections can only exchange data with other DVCPRO systems, not with DV or DVCAM gear. Since a 1394 transfer is a direct data dump, this is understandable; if a cross-format transfer were to be possible it would require that one deck or the other "translate" the signal to or from the DVCPRO data format to the Blue Book format.

As far as incompatibility with 1394 transfers to and from NLEs, this limitation is expected to diminish (and eventually vanish) as developers get a chance to work with DVCPRO over 1394, and to provide switches inside their programs to supply a Blue Book or DVCPRO datastream as required. Matrox and Canopus, for example, have D-7 compatible versions of their NLEs.

Remember, D-7 was designed first and foremost as an ENG format; robustness of the signal was paramount, and interconnection of gear in the ENG world is done via analog or via SDI (1394 is too limited an interface for the broadcast world, where the ability to switch and route video over thousand-meter runs is both necessary and taken for granted; 1394 has a length limit of 4.5 meters and requires a point-to-point session-level communication instead of a switchable open-ended transmission). 1394 was added to the DVCPRO lineup as an afterthought, at the prompting of customers, and as it becomes more prevalent (and if the marketplace demands it) you'll see more NLEs capable of dealing with D-7 data as readily as with Blue Book data, and possibly even realtime DV/DVCPRO format translators. It's early in D-7's evolution; there may yet be surprises up Panasonic's sleeves...

... for example, at NAB 2001 Panasonic introduced the AJ-D455 D-7 VTR, which can input and output Blue Book *or* D-7 data over 1394, regardless of the tape being played (DV, DVCAM or D-7) or recorded (D-7 only)!