Friday, 19 February 2016

The Engineer’s Bench Podcast; Video top tips

Introduction

We wanted to have three sessions of the podcast with some top-tips for the new year; useful tricks that allow the jobbing engineer to solve problems that we’ve come across over the years; this list is by no means exhaustive but all these tips are things that have saved our bacon!

This time it’s video, but coming up we’ve got data/power - that last one seems like a strange but generally if you’re doing a lot of work in data centres or comms rooms the thing you have to deal with alongside fibre and copper data is the mains.

1. 4k and UHD cabling and signal standards

I've had to dig into signal transport for 4k/UHD recently. Essentially I have a test-signal generator (SRI Visualizer TG100) running at a maximum raster of 4096x2160 at a maximum of 25 progressive frames/sec (and only 4:2:2 colour sampling; Y, Cr, Cb) with a 6G single-link output (so really 4 x 1.5G links) and HDMI 1.4 (so the same raster as the SDi). The monitors are the 24" and 30" Canon IPS 4k native monitors.

The Canon monitors will take quad-link HD/SDi and (in the case of the 24") HDMI. So, feeding the SRI single-link into a Blackmagic 4k multiplex (to produce quad-link) and then into the Canon produces four quads in the wrong colour-space!

The original 4k-over-four-BNCs standard

The more recent standard; each link looks like an HD version

Clearly the converter is producing 2SI but the Canon expects SD quad-link. In fact the guys at Canon tell me they have a firmware update early in 2016 to address this. The other error is that the Canon has mistaken the 4:2:2 video as RGB - but it has at least got the raster correct.

So, what to do? Well, by throwing in another converter and taking the HDMI out of the SRI means the BM mux will get an older SD quad-link input;

This produced what we need; clearly HDMI has not concept of mutliplexed pixels and so we're now fully in SD quad-link;

tugging BNC no.4 shows the monitor is now in quad mode

2. Perceptual video quality & compression - PSNR measurements

Compression is a fact of life, there have only been two production VTRs that stored uncompressed video - D1 and D5; they are no longer used because they were both SD (and D1 was only eight-bit video). So, the vast majority of the material we handle is compressed and so there should be a way of quantitatively judging it. There three methods of numerically analysing how good pictures are, but for the most part and engineer or editor proclaiming "...those pictures look a bit soft" is what still passes for picture quality analysis!

This comes from a discussion with a couple of industry colleagues earlier in the week about the quality of satellite contribution circuits. They'd got into a bit of a to-and-fro with the carrier over bit-rates and chroma-sampling structures ("but 4:1:1 isn't as good as 4:2:0 for the same data rate" etc.) which for my money entirely misses the point. You have to assess the picture quality of a compressed link (satellite, IP, etc.) not on encoder settings but on perceived picture quality. A few thoughts;

- Modern codecs perform better than older codecs when considering data rate vs quality
- Progressive pictures compress nicer than interlaced pictures
- Statistical multiplexing always produces better results over multiplexed connections
- Long GOP codecs outperform iFrame codecs by a factor of 5:1 typically.

The three methods used for determining video quality are;

1. Peak signal-to-noise ratio - PSNR

2. The "Just Noticeable Difference" or JND; a very BBC-type measurement and commonly derived by asking a crowd of observers to assess picture quality. I like the idea of this and when I was at the BBC you'd often hear people specifying "...half a JND" as being a required spec; that was also know as a "gnats" as in "gnat's whisker"! It's problematic because it describes perception rather than the effect of exposure - an editor friend told me that he preferred to work on DigiBeta pictures over DVCam footage not because he could spot the difference between shots from the more expensive format over the cheaper format (all other things - camera, lens, lighting etc being equal), but the more compressed pictures just made him feel more tired by the end of the day. The JND takes no account of the cumulative effect of looking at compressed footage that may at that moment look just as good but more subtly takes it toll on the viewer.

3. Mean Opinion Score - MOS; very similar to the JND but with a numeric score. I won't talk about this.

So, PSNR - to steal from Wikipedia;

> *...the ratio between the maximum possible power of a signal and the power of corrupting noise that affects the fidelity of its representation. Because many signals have a very wide dynamic range, PSNR is usually expressed in terms of the logarithmic decibel scale.*

PSNR is calculated as a rolling set of differences between source material and the compressed version and is most easily defined via the mean squared error (MSE). Given a noise-free m×n monochrome image I and its noisy approximation K, MSE is defined as:

$$MSE = \frac{1}{m\,n}\sum_{i=0}^{m-1}\sum_{j=0}^{n-1}[I(i,j) - K(i,j)]^2$$

The PSNR (in dB) is defined as:

$$\begin{aligned} PSNR &= 10 \cdot \log_{10}\left(\frac{MAX_I^2}{MSE}\right) \\ &= 20 \cdot \log_{10}\left(\frac{MAX_I}{\sqrt{MSE}}\right) \\ &= 20 \cdot \log_{10}(MAX_I) - 10 \cdot \log_{10}(MSE) \end{aligned}$$

The point is that it can be calculated from the pixels. No observer bias is involved.

Engineers love quick rules of thumb, and PSNR for video images are no different;

1. For idential images the MSE is zero and hence the PSNR is infinite (more dBs = better pictures!)
2. 40dBs is considered to be indistinguishable from uncompressed production quality (that's where the BBC JND lives!)
3. 32dBs is considered desirable for quality broadcast link circuits
4. High twenties is what you can expect for over-the-air transmission - the 10Mbit DVB-T2 pictures you watch on Freeview or Sky.

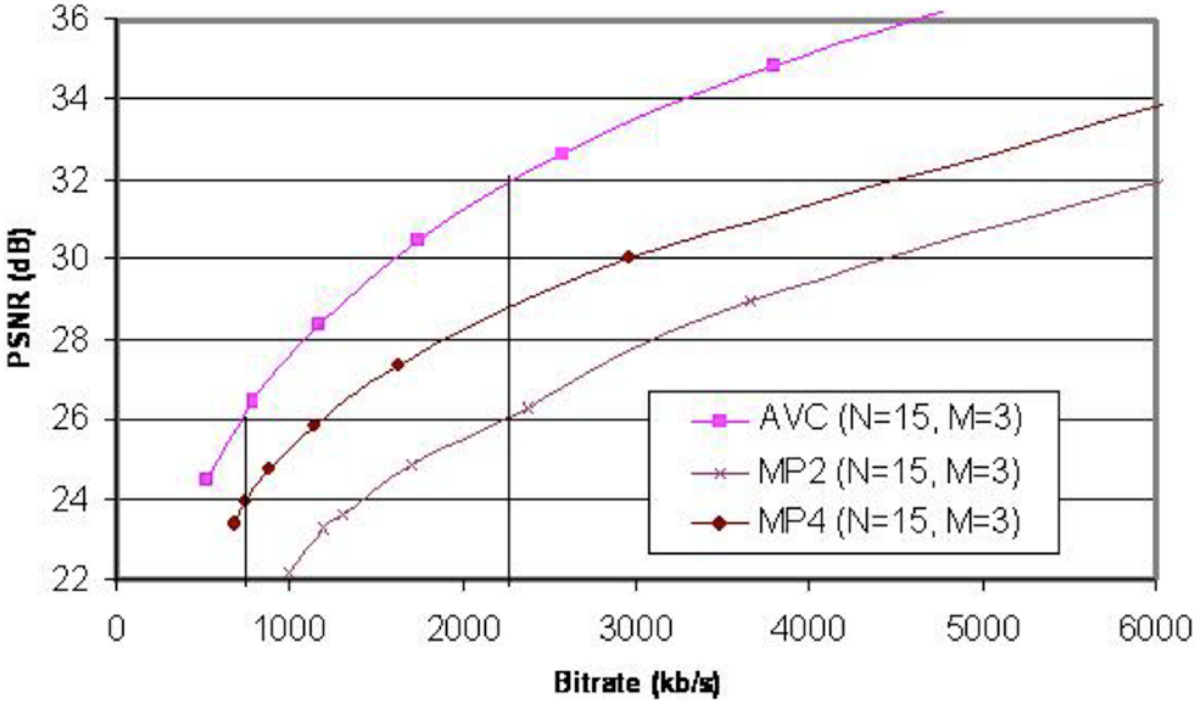

So, to return to my point 1 (above) - here is a graph showing data rates against codec types for the same SD pictures. Venerable old MPEG2 (from the mid-90s!) up against vanilla MPEG4 (late nineties) and AVC (AKA H.264 / MPEG4.pt10 - early noughties). A full 6dBs of quality (twice as good in layman's terms?) lie between those two codecs at 2.2Mbit/sec (all other things being equal - use of a Stat Mux etc). You could even dive in further to MPEG2 and see the difference between the implementations from twenty years ago and what folks like Main Concepts are doing in 2014). The decoders particularly are now much better at hiding macro-block edges and recovering from corrupt frames.

Point 2 (above) seems obvious, but only when comparing 1080i with 1080p pictures AT THE SAME FRAMERATE; so perhaps best to say 1080i vs 1080PsF; Interlaced pictures will always be a challenge as pixels (and hence macro-blocks) move within a frame, unlike progressive pictures. BUT, you still get better motion rendition withing interlaced frames for the same framerate. Eventually we'll have moved to 1080 50/60P and so it'll be a moot point.

This graph shows data rate for HD pictures, we expect over-the-air HD to be at ten megabits in the UK.

So, how to make these assessments if you're worried about a contribution circuit or transmission path that you're responsible for? If you work in coding and mux then you

Friday, 19 February 2016

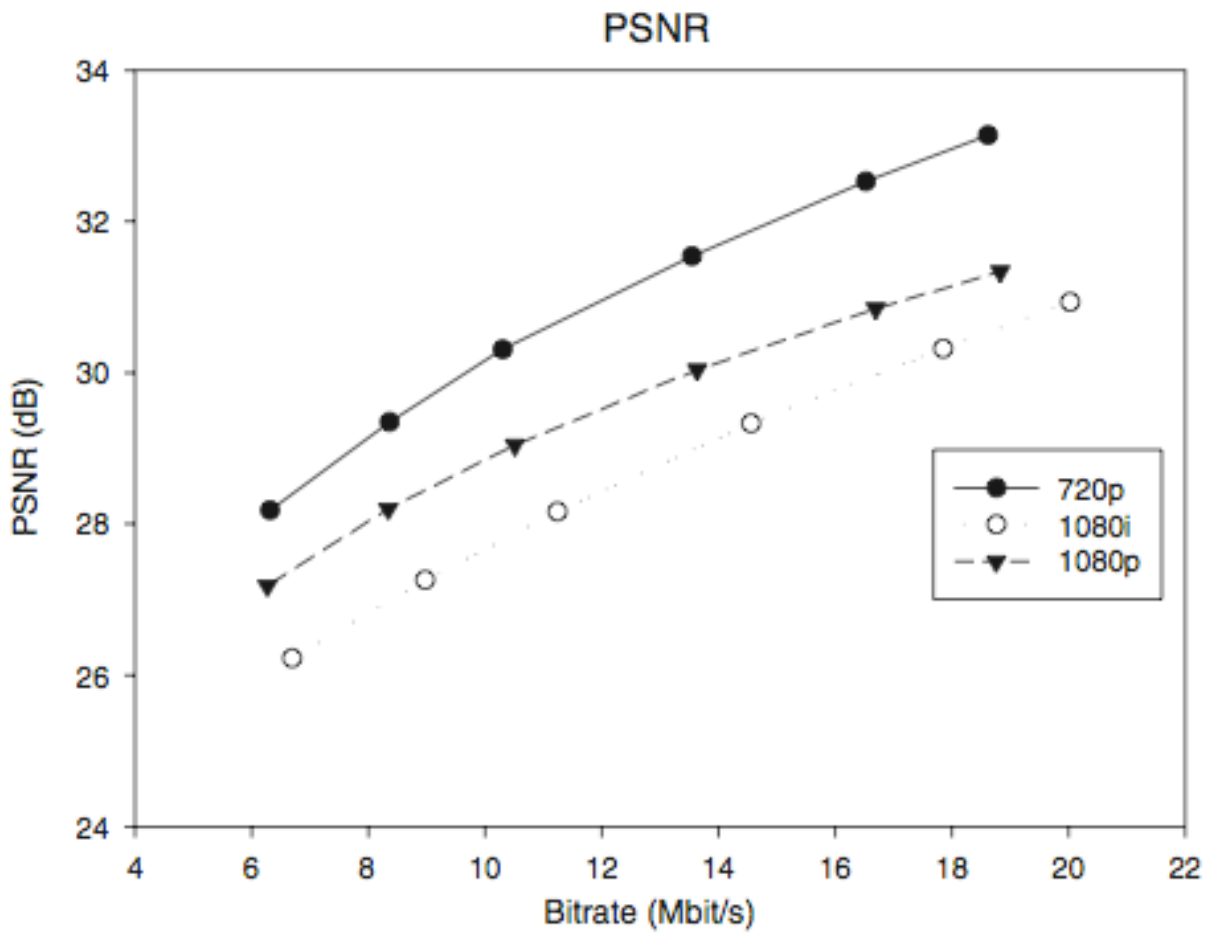

probably already have tools to assess. PSNR is such an important part of delivery specs/SLAs in broadcast (you need to keep the accountants at bay after all!) that you'll have a Tektronix PQA600 or a Rohde & Schwarz DVMS-series test set.

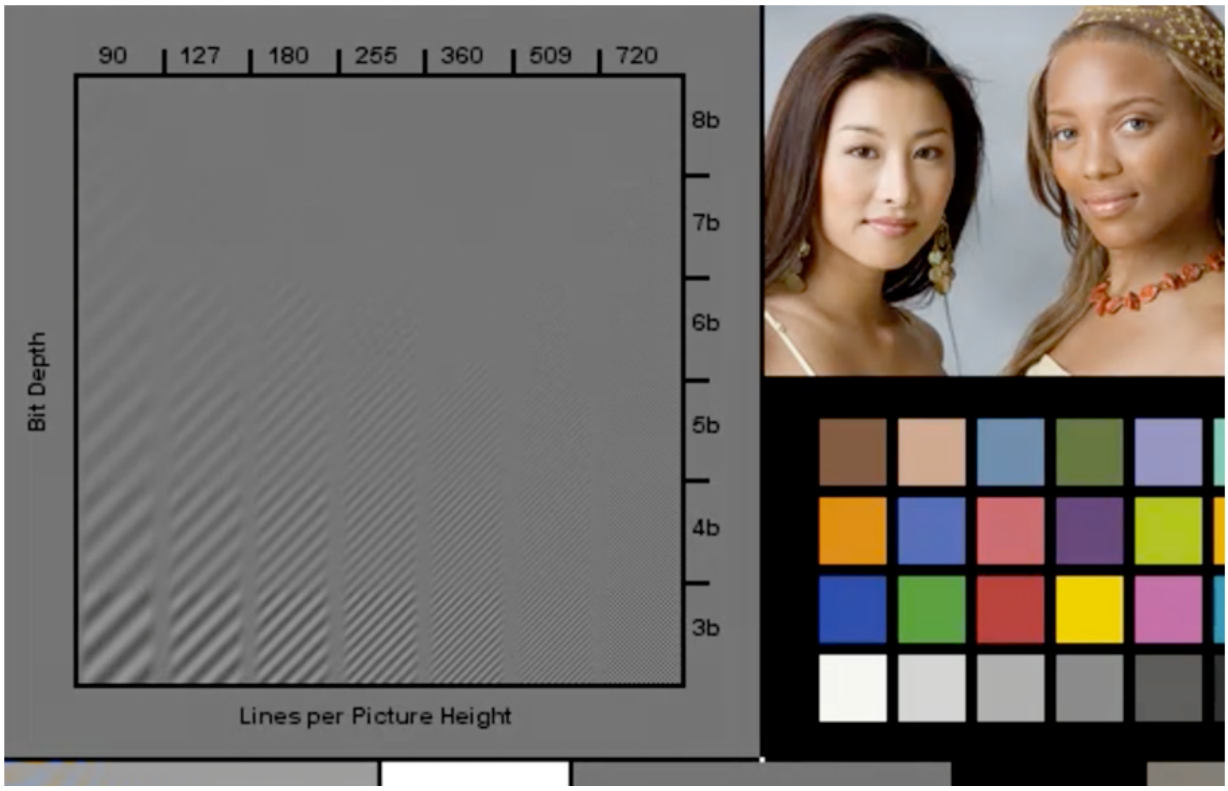

However, "traditional" video quality measurement needs access to the compressed signal and the original which may not be possible; particularly in the case of my pals who are at loggerheads with their satellite provider. What you need is a test signal that you can feed over the connection and then make an assessment from the picture content as to how badly the pictures are being degraded. I've banged on about the SRI Visualiser before but it has a compression multiburst that shows you lines-of-TV-resolution against perceived bit-depth. You can then relate the two worst-case lines/bit-depth figures to the table of PSNR values.

---

3. HD/SDi physical layer measurements - gotta be Tektronix

Here's a brilliant document from Tek - tell you all you need to know about eye pattern, jitter etc.

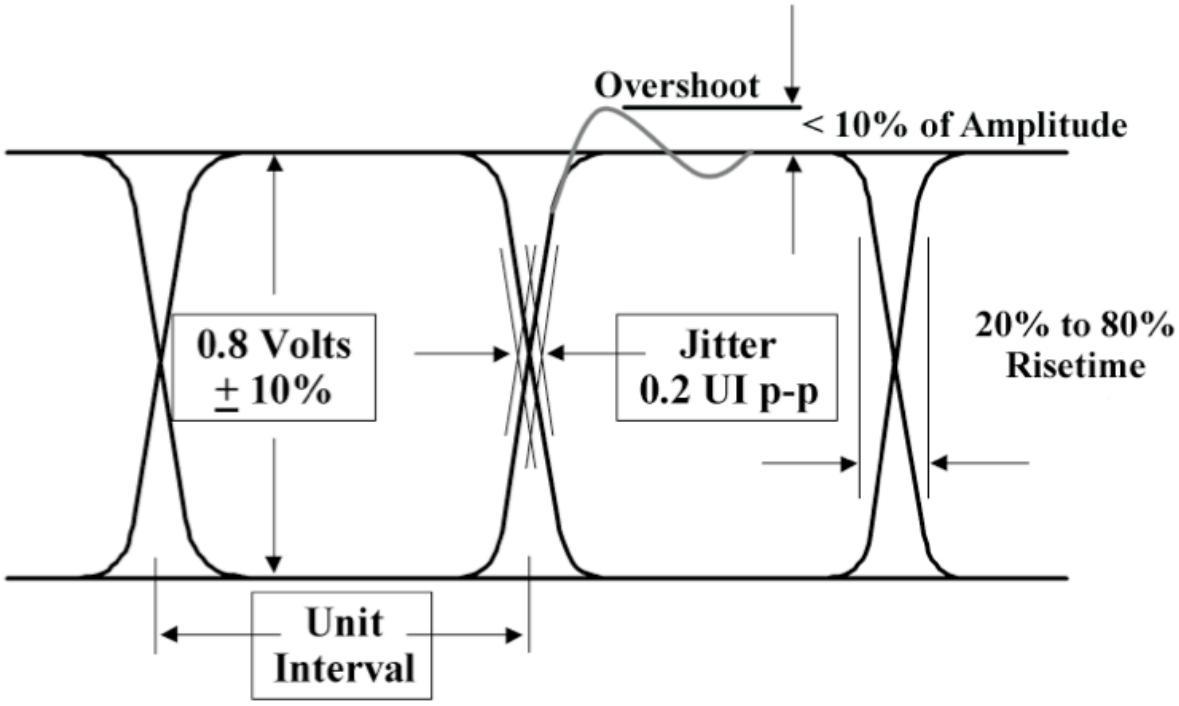

http://www.tek.com/document/how-guide/sdi-eye-and-jitter-measurements…

---

4. Cheap HD/SDi parts, Return Loss and the danger of short cables

Return loss is the loss of signal power resulting from the reflection caused at a mismatch with the terminating load or with a device inserted in the line. It is usually expressed as a ratio in decibels (dB);

$$RL(\mathrm{dB}) = 10\log_{10}\frac{P_\mathrm{i}}{P_\mathrm{r}}$$

where RL(dB) is the return loss in dB, Pi is the incident power and Pr is the reflected power.

Return loss is related to both standing wave ratio (SWR) and reflection coefficient (Γ). Return loss is a measure of how well devices or lines are matched. A match is good if the return loss is high. A high return loss is desirable and results in a lower insertion loss. In the case of newer budget HD/SDi equipment return losses can be as bad as 12dBs (I've measured Blackmagic boxes thus) whereas the spec for 3G is 16dBs and in the case of proper broadcast manufacturers 18dBs or better is often measured (Sony, Tektronix).

I recently tested some real budget SDi parts from a reseller who is thinking of importing them from the Far East. Here are a couple of eye patterns of a 1.5G signal at the input and via the high-impedance looping output of a distribution amplifier.

If you take the first measurement as the base (it's not brilliant but that probably due to the poor signal generator) and then see what you get when the feed is connected to the DA's input and measured at the Hi-Z loop-through you see the effect of return loss - the ability of the sending equipment to drive the line impeded by reflections at the receiving piece of of equipment because of sub-optimal termination.

One thing to look out for is using short cables with budget SDi parts because the reflected portion of the signal doesn't have the chance to dissipate and interferes with the incoming signal. Replacing the short coax cable with a longer one solves the problem. This seems counter-intuitive as all engineers are painfully aware that signals get more compromised by long cables. We all carry a set of figures around in our heads; Gigabit Ethernet over cat6; 100m, 3G HD/SDi over Belden 1694; 60m, 10gig Fibre Channel over OM3; 600m etc etc.

I was left scratching out heads with the following setup;

HD/SDi MTX, optical o/p -> 100m OS1 line -> optical->SDi converter -> SDi-HDMI converter -> TV

We were getting nothing at the TV. Checking for SDi just pre the HDMI adapter using a Tektronix WFM5200 (no eye patterns, unfortunately!) was fine, but the giveaway with sticking a signal generator just pre the HDMI converter and using the same short-length cable showed no signal at the TV. Replacing the short cable with a longer one (15cm with 1m) made it all come good.

Clearly the Tektronix had a much better return loss on it's input and so wasn't phased by being fed with the short cable. The cheap HDMI converter was a lot more sensitive. Interestingly switching between 270MBit/s SD and 1.5GBit/s HD made no difference.

the short cable was the bogeyman!

5. Crimp tool for all current video ends?

All the wiremen want me to buy them one of these! This tool can do SD01 standard def coax (Image 360), 1694 HD Coax and the newer Din1.0/2.3 mini-HD connectors.

6. Synchronous ain't always best!

Engineers have been trained since time immemorial that with multiple sources of video the best thing is that they are always locked together and timed to a common reference - ideally station black & burst (or TriSyncs nowadays). The reasons for this are numerous, but a couple;

Studios cameras into a vision mixer have to be locked to achieve clean cuts between the cameras - it would look rubbish if you had a frame roll every time you cut between sources. The same is true of sources into a continuity suite etc.

In the case of Avid for the longest time you had to make sure the VTR was locked to the same reference as the Media Composer (all the way from v.7 ABVB systems in the mid-90s to the last revision of Adrenaline in 2007!) otherwise the audio and video capture portions of the machine would free-run WRT each other and within minutes you'd be loosing frames of sync.

So - I know a big facility that had won an archive digitising project and they are using the BlackMagic DeckLink 4K card to allow them to ingest 4 x DigiBeta tapes at once. The capture software is ToolsOnAir and they found that after the first clip was captured the second clip would be a frame out of sync, and progressively worse after that - unless you re-booted between captures! It turns out that if the four VTRs are allowed to free-run then you don't get the problem. Perhaps processing the vertical syncs places a burden on the card/software and if it happens simultaneously on all four inputs trouble ensues?!

It reminded me of a situation with a big broadcaster who was distributing their regional variations over Astra on a single multiplex. The stat-mux was very unhappy with material that was (for the most part) identical; only the ad-breaks differed. Most video cuts occurred at the same time across all six SDi feeds. The solution was to apply a two frame delay between all the sources (so o/p 6 was now 12 frames late WRT o/p 1).

7. Rise times in HD TriSyncs

These are two traces from two separate TriSync generators - The blue trace represents a correct waveform and as every superhero will realise you're looking at the line timing pulse. Here is the diagram from rec-709 (the spec for HD video);

Line synchronizing signal waveform

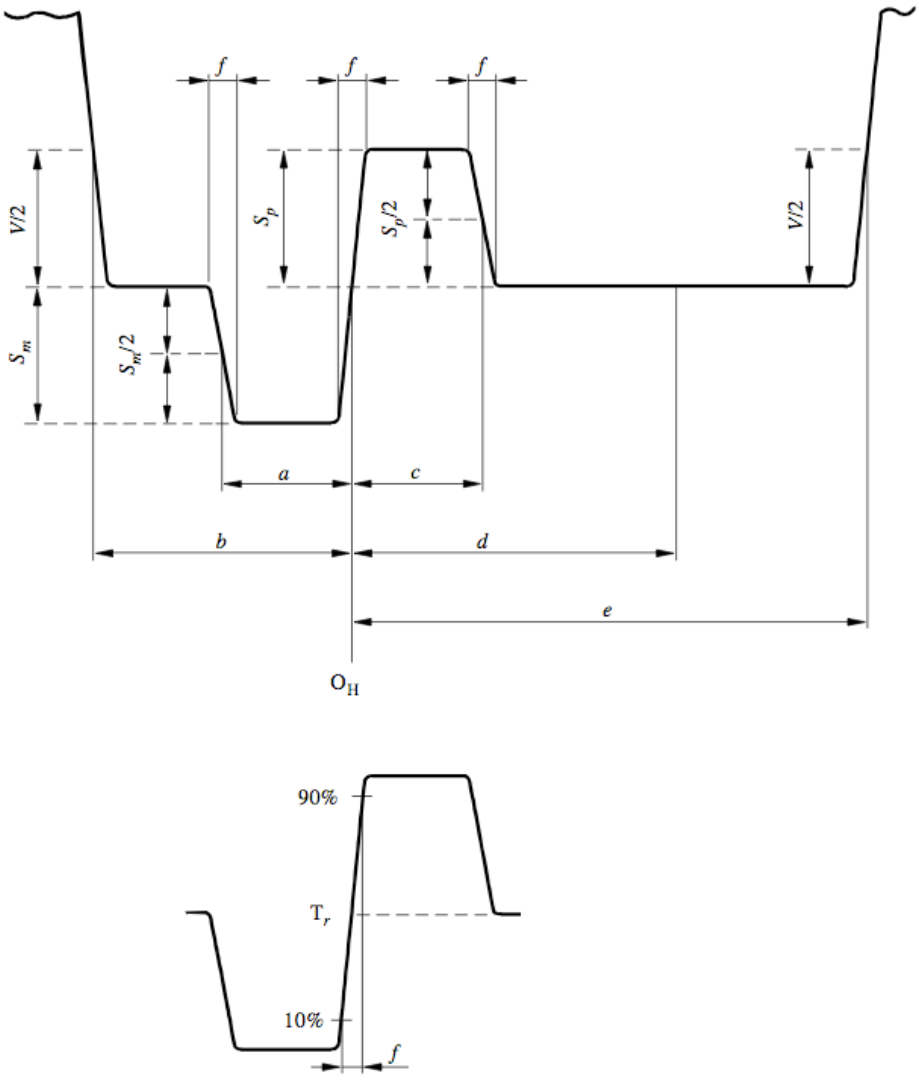

(The waveform exhibits symmetry with respect to point T_r)

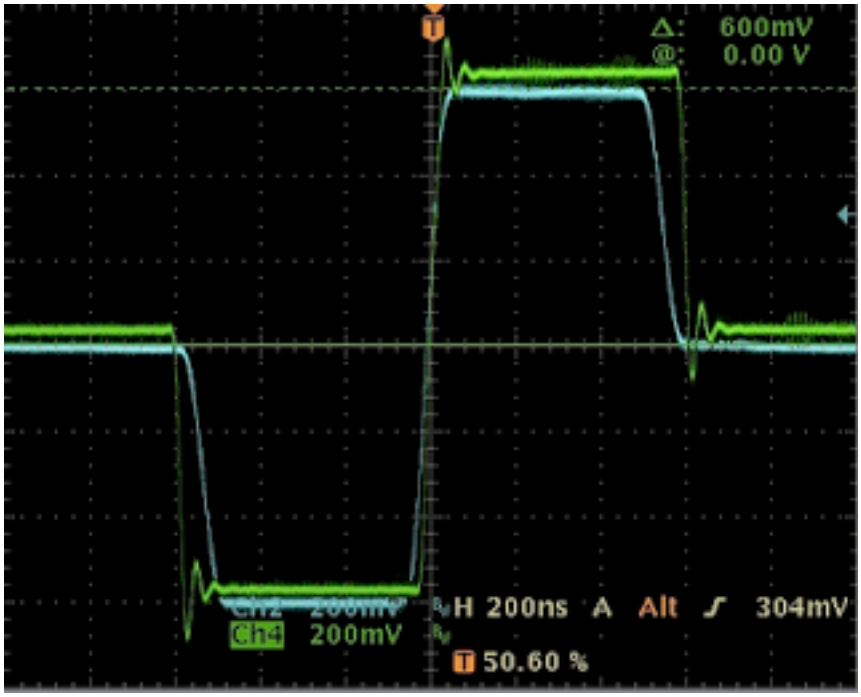

It clearly shows that rise time is to be equal between the start, middle and end of the line-sync pulse. The rep from the manufacturer of the green pulse insisted that his waveform was a lot sharper - but given the ringing on it I think they just aren't filtering it properly to comply with the rise time spec - seen here (again from rec-709) as being 4 clock cycles +/- 1.5

8. High Def 25Psf video

Two decades ago when Mr Sony was developing what would become HDCam (with some small contribution from the previous 1" open-reel HDVRS analogue & digital formats) they realised that progressive video was the future but existing HD equipment (typically the BVM-D series monitors) couldn't lock to such a slow framerate (24/25/29.97 as opposed to 48/50/59.98 fields). The answer for progressively-sourced pictures was the Psf standard which makes progressive frames look like interlaced video. So as to make film people think that this was better than video they have a new name for a field - the segment. In fact Psf is interlaced video (but there is no movement between the fields) - it just shows that good old interlaced video is able to faithfully reproduce progressive pictures (but the reverse is not true as progressive video with the same frame-rate has only half the motion rendition as interlaced video).

So - let's dismiss a couple of misconceptions;

There is no difference between a Psf signal and an interlaced signal from a technical standpoint.

Sending 1080 pictures via Psf doesn't degrade them in the slightest - in fact if you're laying off 1080 to HDCamSR then anything below a 5800 (in 50/60P mode @880mBits) is recording Psf!

Now then - below are screen-grabs from my trusty WFM7120. The first shows the output from a Symphony NitrisDX BOB. The footage had come from a Sony EX3 cameras recorded at 35mBits 1080/25P onto Memory Stick and imported straight into a

progressive timeline. The Avid plays back Psf which the Tek shows as 1080i (for the reasons discussed above). Laying this off to HDCamSR (a 5500 deck) gives a 25Psf recording on tape. The second screen shot is the Quicktime sample movie imported into a new 25P timeline - it just serves to confirm that the BOB output is always Psf.

This last picture is the down-convert output of a Leitch X75 which is a great little get you out of trouble box (basically does everything->everything with a few extra tricks thrown in - profanity delay etc.) but it's not a multi-frame broadcast standards converter (like a Snell & Wilcox Ukon).

---

9. Barrel and a terminator for easy line identification

Although you expect video cables to be numbered occasionally (or more often than you'd like!) you might have to identify long cables where you can't just tug them and see!

So - 75ohm termination with a BNC barrel (or a MUSA-BNC adaptor if you're at the front of the video patch) and your multimeter set to the lowest resistance range and you can squeak the line; but, beware! Video equipment terminates (and sends) in 75ohms so make sure you pull and re-connect the terminator to identify the line (or maybe make up a terminator with some unusual value of resistance; maybe 200 ohms to not be confused with another signal termination standard).