From Analogue to Digital

Although we all know and understand the crude basics behind analogue and digital, John Watkins from Snell&Wilcox, takes us deeper into explaining the theory behind an increasingly more important digital video world

The world of television production and program delivery is going digital. There are good reasons for that which we will take a look at here. No-one goes digital for its own sake. Sure, it's glamourous, but there had better be more reasons than that; like economic reasons for instance. Quite simply, a carefully planned digital installation costs less to run today than a traditional analogue plant doing the same job. It can run with less down time and needs less scheduled and unscheduled maintenance.

The experimental phase is over and now that digital technology has matured to deliver reliable, cost effective hardware, it's analogue that suddenly looks uneconomic. It's not just the hardware which is changing, but the whole nature of television is under scrutiny.

The nature of television changed when the home VCR was developed. Viewers could time-shift broadcasts, or purchase or rent cassettes of their choice - an early form of Video-on-Demand.

Television changed again when competition from cable and satellite came along. Now it is set for a period of further change.

The aspect ratio will change from 4:3 to 16:9. Higher quality pictures will reach the home via advanced signal standards. Economics says that advanced television will have to use compression which also means it will be digital.

Mass storage technology from computers and video compression technology combine to allow true Video-on-Demand where the viewers choice is communicated to the server via a back-link. That back-link also allows interactive video and home shopping. Some personal computers can already receive television pictures and play compressed digital video disks.

So unless you want to start a sheep farm, it's not a question of whether or not to go digital, more a matter of when and to what extent.

Digital does the same job as analogue, but in quite a different way and with more operational freedom. Treat digital equipment like analogue and you'll get nothing but trouble and miss out on the freedom. Treat it right and you'll reap the benefits. One of the jobs of this article is to explain what right means.

Digital is complex and the mathematicians can have a great time spouting formulas and buzzwords which make a lot of sense to rocket scientists. This article is not for them. It's for normal people who want to make decisions without getting buried in details. So there's no math here, and buzzwords are only put in so we can explain what they mean rather than to show off what a bunch of smart-asses we are.

What is digital?

Digital is just another way of representing an existing television waveform.

Fig.1a) shows that a traditional analogue video system breaks time up into fields and frames, and then breaks up the fields into lines.

These are both sampling processes: representing something continuous by periodic discrete measurements. In digital we simply extend the sampling process to a third dimension so that the video lines are broken up into three dimensional point samples which are called *pixels* or *pels*. Both terms are what you get when you say "picture cells" in a hurry.

Fig.1b) shows how a television line is broken up into pixels. A typical 525/60 frame contains around a third of a million pixels. In computer graphics the pixel spacing is often the same horizontally as it is vertically, giving the so called "*square pixel*". In broadcast video systems

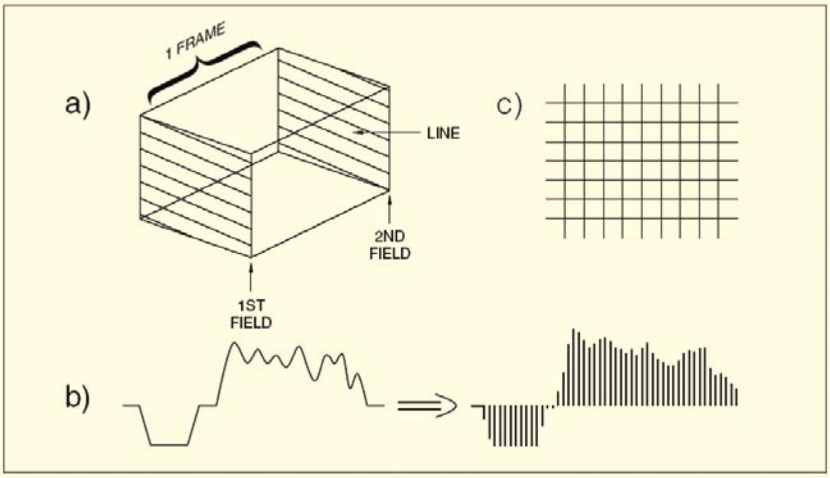

Fig 1a In analog television, time is sampled into fields, and each field is sampled vertically to make lines. If each line is also sampled, as in b), an array of pixels will be formed. Such an array can be made from a conventional analog television signal by sampling it a line at a time as shown in c).

pixels are not quite square for reasons which will become clearer later in this section. Once the frame is divided into pixels, the variable value of each pixel is then converted to a number.

Fig.1c) shows one line of analogue video being converted to digital. This is the equivalent of drawing it on squared paper. The horizontal axis represents the number of the pixel across the screen which is simply an incremental count. The vertical axis represents the voltage of the video waveform by specifying the number of the square it occupies in any one pixel.

The shape of the waveform can be sent elsewhere by describing which squares the waveform went through. As a result the video waveform is represented by a stream of whole numbers, or to put it another way, a data stream.

Any waveform can be digitized in this way. With analogue TV signals, there is only one waveform per channel. When the TV signal is digitized, the entire waveform, chroma included, is just expressed numerically. Digital TV still has the same characteristics as analogue, including the four field sequence and 0.1% frame rate offset because it is one and the same thing.

In the case of component analogue video there will be three simultaneous waveforms per channel. We need three convertors to produce three data streams in order to represent GBR or colour difference components.

By converting video (and audio) signals into data we have the freedom to use storage, processing and transmission techniques which have been developed for computers.

Digital transmission has the advantage over analogue in that it can reject noise and jitter. If we convert the whole numbers representing the video waveform voltage into binary, using the conversion table shown in Fig.2, our resultant binary digits, or bits, have only two states, 1 and 0. These two states can then be represented by any electrical, magnetic, optical or mechanical system which can exist in two states. Fig.3 shows binary represented electrically by two different voltages, magnetically by the reversal of the direction of magnetisation, optically by alternate opaque and transparent areas of a recording, and

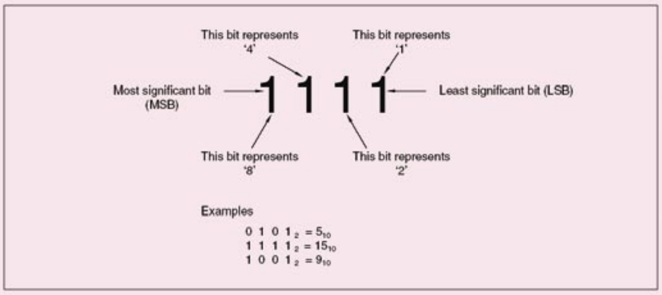

Fig 2. A binary conversion table allows any ordinary decimal number to be stored or transmitted by any system which can exist in two states.

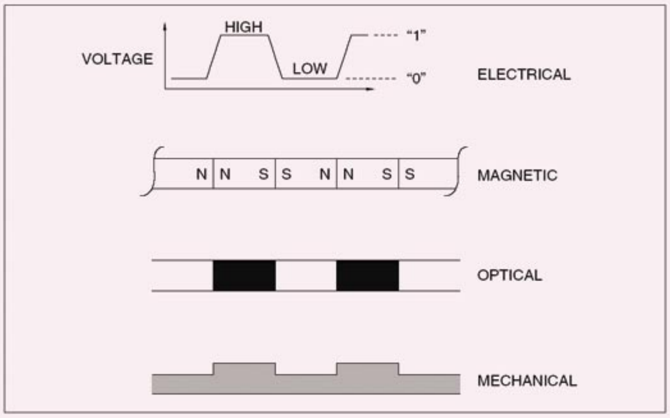

Fig 3 Several ways of handling binary data are shown here. In electronic circuits, two different voltages are needed, commonly achieved with a simple on/off switch. In magnetic recording, two flux directions are used. Optical recording may use alternating opaque or dark areas and mechanical storage may use raised or depressed areas on a surface.

mechanically by the presence of pits in the surface of a laser disk.

With only two states, more noise can be rejected than in any other system.

Fig.4a) shows how using the example of an electrical interface. Although the signal transmitted is a clean, two-level waveform, by the time it reaches the receiver it will have picked up noise and jitter. The receiver compares the voltage of the signal with a reference which is mid-way between the transmitted levels in a process called slicing. Any voltage above the slicing level is considered a 1 and any voltage below is considered a 0. This slicing process will *reject considerable amounts of noise and restore the signal to clean binary* once more.

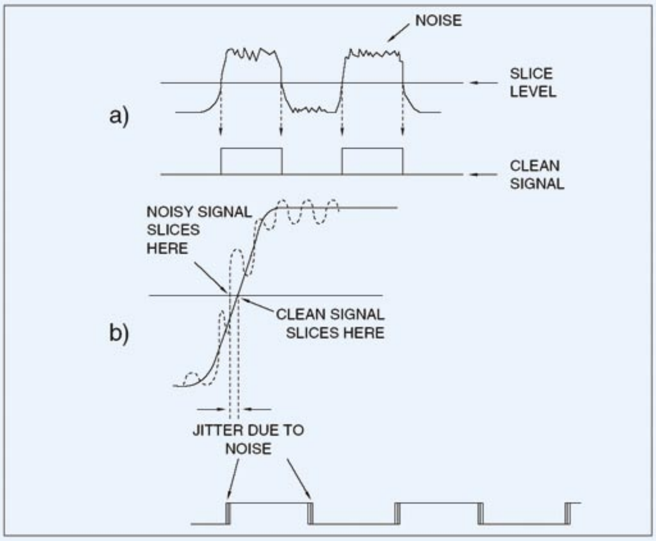

Fig 4 A binary signal has the advantage that a simple slicer shown at a) can be used to recreate a square signal. The result is that noise only causes jitter as shown in b). This can be removed with a phase locked loop.

Slicing will fix the noise, but it cannot fix jitter. We need another process to do that. Fig.4b) shows that jitter has the effect of displacing voltage changes, or transitions, along the time axis in a random manner.

However, the average timing of a large number of transitions is unaffected. A phase-locked loop (PLL) is an electronic circuit which can average the timing of many transitions to recreate a stable clock from a jittery one. It acts like the flywheel on a piston engine which averages the impulses from the pistons to produce smooth rotation.

The combination of a slicer and a phase-locked loop is called a *reclocker*.

The neat thing about a reclocker is that the waveform coming out of a reclocker is identical to the waveform at the transmitter. Noise and jitter have been rejected so there has been no loss of data due to the transmission. The really neat thing is that we can cascade or tandem as many reclockers as we like and the same data will come out of the end as is shown in Fig.5.

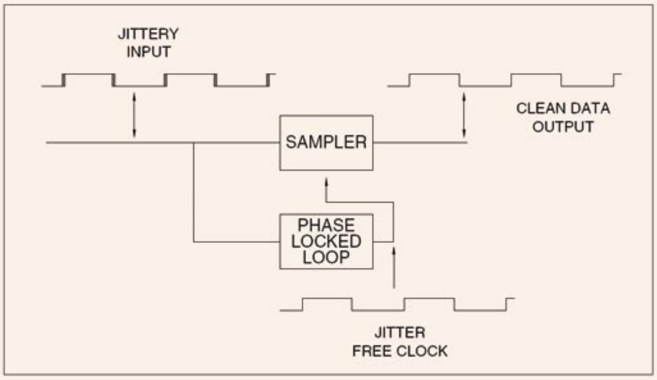

Fig 5 A reclocker uses a phase locked loop to recreate a clean clock which samples the input once per bit and produces a jitter free output. Any number of slicers and reclockers can be cascaded without quality loss.

The reclocker is found everywhere in digital systems. Digital VTRs use it to eliminate tape noise and jitter. Hard disk drives use it, CD players use it, routers use it. The universality of reclocking is due to the fact that it prevents generation loss. A device which contains a reclocker launches a clean signal which is just as robust as the signal launched from the previous device.

One of the reasons for the use of digital systems is that *no matter how many generations of recording or transmission are used, there will be no generation loss.* Note that this is only strictly true if we don't employ compression. Fig.6 shows a simple system with analogue inputs and outputs, analogue-to-digital (ADC) and digital-to-analogue (DAC) convertors, and a full bit-rate DVTR with digital inputs and outputs.

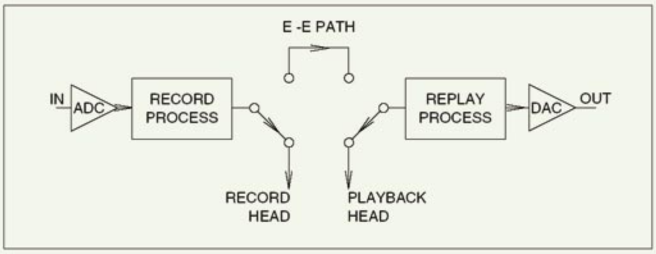

Fig 6 In a simple DVTR with effective error correction, the picture quality in playback is the same as in E-E.

The data leaving the ADC are reclocked on receipt at the DVTR and recorded identically on the tape. On playback, the off tape data are reclocked to eliminate tape noise and jitter. Any data which are incorrect because of tape dropouts will be put back to their original values by the integral error correction system prior to transmission to the DAC. The DAC reclocks the received data so that it is identical to that leaving the ADC.

The result is that the picture quality of the system output is determined only by the ADC and the DAC. The ADC could be connected straight to the DAC, or the DVTR could be put into E-E mode, bypassing the tape, and the picture would not appear any different. This is true provided that there is no compression system and that the data are not corrupted by noise in excess of what the reclockers can handle or by dropouts which are too big for the error correction system of the DVTR. Thus in digital systems there are two main areas of concern:

1. ADCs and DACs determine picture quality and the best available should be used. Cheap convertors are false economy.

2. Quality is determined by the convertors only if there are no data errors. Therefore the goal in digital systems is DATA INTEGRITY.

Data integrity cannot be measured with conventional analogue monitoring equipment.

How do I convert from analogue to digital?

This section deals with the ADC and the DAC which together largely determine the picture quality achieved in a digital system. Convertors use three fundamental processes: filtering, sampling and quantising.

Fig.7a) shows what happens when sampling is done right. The original waveform is preserved in the envelope of the samples.

Fig.7b) shows what happens if we screw up and put in a signal whose frequency is too great for the sampling rate in use. The envelope of the samples now carries an incorrect waveform. This is *aliasing*; the result of incorrect sampling.

Everyone has seen stagecoach wheels stopping and going backwards in cowboy movies. It's an example of aliasing. The frequency of wheel spokes passing the camera is too high for the frame rate in use. It is important to prevent aliasing in convertors and this is done by including a filter, called an anti-aliasing filter, prior to the sampling stage. Don't leave home without one. The anti-aliasing filter stops frequencies higher than one-half the sampling rate from entering the sampler and so aliasing cannot occur.

In addition to the sampling process the ADC needs a quantiser to convert the analogue sample to a binary number. Fig.7c) shows that a quantiser breaks the voltage range or gamut of the analogue signal into a number of equal-sized intervals, each represented by a different number. The quantiser outputs the number of the interval the analogue voltage falls in.

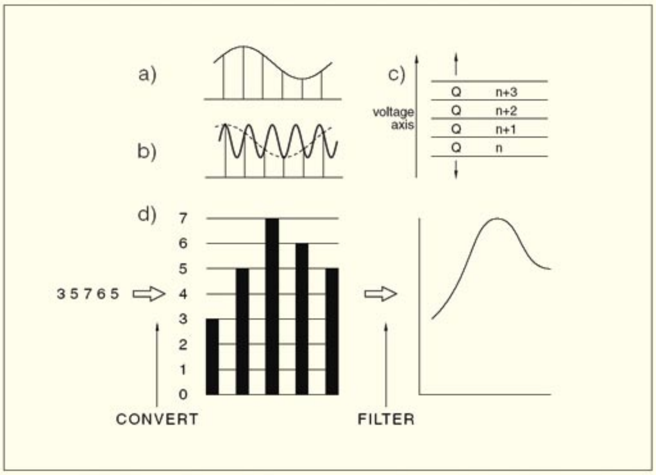

Fig 7 a) A high sampling rate is intuitively adequate, but if the sampling rate is too low, aliasing occurs as shown in b). c) Quantizing divides the voltage range up into equal intervals Q. The quantized value is the number of the interval in which the input voltage falls.
d) Numbers are converted back to voltage pulses, and these are filtered to produce a continuous signal.

The position of the analogue voltage within the interval is lost, and so an error, called a quantising error, can occur. As this cannot be larger than a quantising interval the size of the error can be minimized by using enough intervals. In an eight-bit convertor there are 256 quantising intervals because this is the number of different codes available from an eight bit number. This allows an unweighted SNR of about 50 dB. In a ten-bit convertor there are 1024 codes available and the SNR is about 12 dB better.

Equipment varies in the word-length it can handle. Older equipment and recording formats such as D-1 only allow eight-bit working. More recent equipment uses ten-bit samples.

Fig.7d) shows that to convert back to analogue, two processes are needed.

Firstly a voltage is produced proportional to the binary value of each sample, then these voltages are passed to a reconstruction filter which turns a sampled signal back into a continuous signal. It has that name because it reconstructs the original waveform. So in any digital system, the pictures we see have come through at least two analogue filters. In real life we may have to convert a signal in and out of the digital domain several times for practical reasons. Each generation, another two filters are put in series and any shortcomings in the filters tend to get magnified. Thus the best filters available are the ones to use. There is no point in using low-priced filters as their defects will then be passed on perfectly by the digital system.

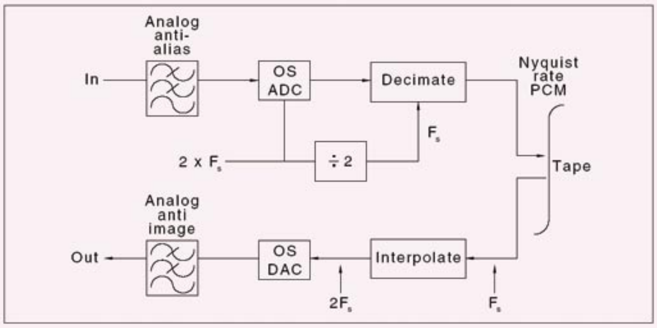

Fig 8 In an oversampling convertor the sampling rate is temporarily increased to ease the problems of filtering in the analog domain. The sampling rate is changed up or down in the digital domain where filtering is easier.

One way of overcoming the drawbacks of analogue filters is to use over-sampling. This can be applied to ADCs and DACs as Fig.8 shows. In an over-sampling ADC, the sampling rate is temporarily doubled, so that the analogue anti-aliasing filter need only stop signals of twice the video bandwidth. Following the ADC, the doubled sampling rate data are passed through a digital filter which performs a precise anti-aliasing function with better phase linearity than an analogue filter could manage. Once this is done the sampling rate can be reduced to the normal figure.

In an over-sampling DAC the sampling rate of the input data is doubled in an interpolator which is a digital filter that computes new values lying between the existing samples. The DAC itself now operates with twice the usual sampling rate and the analogue reconstruction filter finds it easier to separate the doubled sampling frequencies from the video frequencies.

How do digital interfaces work?

This section discusses the standards which exist for interchanging and routing data representing both component and composite video signals.

We also look at the problems of interfacing between composite and component signals.

Standardized composite digital video uses a sampling rate which is locked to four times subcarrier frequency (4FSc). The actual rate is about 14.32 MHz in NTSC. This makes life easier because the phase of samples is then stable with respect to the colour burst and Y/C separation in digital filters becomes simpler. Fig.9 shows how NTSC is sampled at 4FSc. Note that the sample numbering starts at the beginning of the active line and wraps around into the next line. This is because composite digital VTRs like D-2 and D-3 don't record the horizontal blanking interval. They start recording at sample 1 and on replay the horizontal interval is recreated artificially.

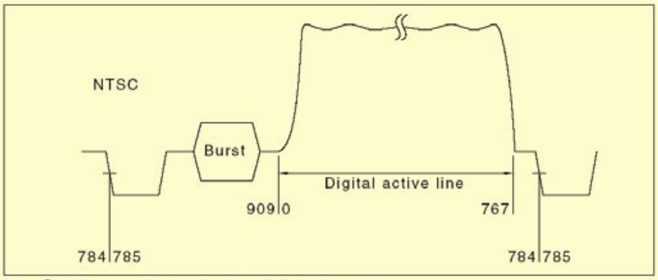

Fig 9 The sample numbering of digital NTSC.

However, the digital interfaces carry the entire NTSC signal. Fig.10 shows how the analogue NTSC signal fits into the eight- and ten-bit quantising ranges. The sampling is carried out on the I and Q axes and as a result samples do not coincide with burst peaks or crossings.

Component signals are sampled differently to composite. A common sampling rate was arrived at which allows 525/60 and 625/50 video to be sampled at a common rate locked to horizontal sync. The figure most often used for luminance is 13.5MHz. The price of this commonality is that the pixels are not square. Fig.11 shows how the U.S. standard TV line fits into 13.5MHz sam-

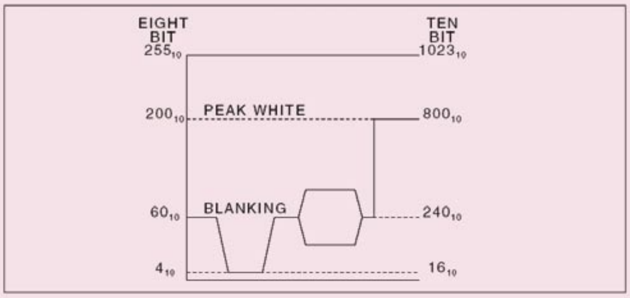

Fig 10 How the NTSC waveform fits into the eight- and ten-bit quantizing range.

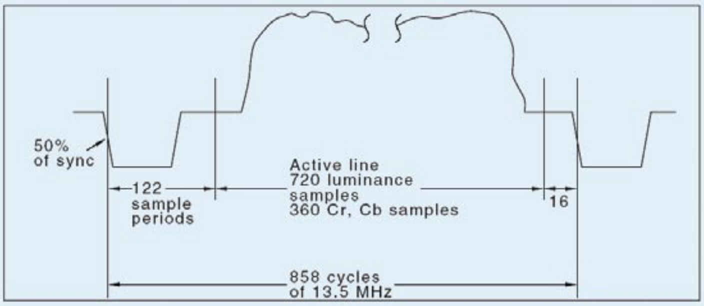

Fig 11 The sample numbering of digital component.

pling. Note that only the active line is transmitted or recorded in component digital systems. The digital active line has 720 pixels and is a tad longer than the analogue active line so the sloping analogue blanking is sure to be included.

In component systems, the colour difference signals have less bandwidth. In analogue components (from Betacam for example), the colour difference signals have one half the luminance bandwidth and so we can sample them with one half the sample rate, i.e. 6.75MHz. One quarter the luminance sampling rate is also used, and this frequency, 3.375MHz, is the lowest practicable video sampling rate, which the standard calls 1. So it figures that 6.75MHz is 2 and 13.5MHz is 4.

Most component production equipment uses 4:2:2 sampling. D-1, D-5 and Digital Betacam record it, and the serial digital interface (SDI) can handle it. Fig.12.a) shows what 4:2:2 sampling looks like in two dimensions. Only luminance is represented at every pixel. Horizontally the colour difference signal values are only specified every second pixel. Two other sampling structures will be found in use with compression systems. Fig.12.b) shows 4:1:1, where colour difference is only represented every fourth pixel horizontally. Fig.12.c) shows 4:2:0 sampling where the horizontal colour difference spacing is the same as the vertical spacing giving more nearly "square" chroma.

Fig.13 shows how component digital fits into eight- and ten-bit quantising. The analogue syncs can go off the bottom of the scale because we only use the active line. Colour difference signals are offset upwards so positive and negative values can be handled by the binary number range.

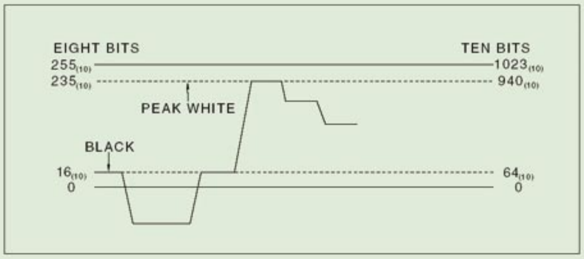

Fig 13 How component video fits into the quantizing range for eight-and ten-bit sampling.

Interfacing digital video can be done two ways: parallel and serial.

Parallel interfaces use a 25-pin D-connector and a heavy multi-core cable. This is an old technology and can't go more than about 50 m before the bits all arrive at different times and confuse the receiver. Maybe the worst news is that a parallel router has to switch every bit and the clock separately, so it is actually an eleven layer device - complex, expensive and impossible to make in large sizes. With the availability of serial chip sets parallel is now effectively obsolete, but only very recent equipment has serial inputs and outputs as standard. If you have some older digital equipment with only parallel connections, don't panic, just add a parallel-serial conversion unit. Adaptors of this kind are referred to as "glue" because they help stick things together.

The ***serial digital interface*** (SDI) allows component or composite digital to be passed down regular coax cable. Instead of sending all the bits of a sample at once on different wires, SDI sends them in turn down the same wire. This needs high frequencies - 270 million bits per second for 4:2:2. At that frequency just serializing the bits

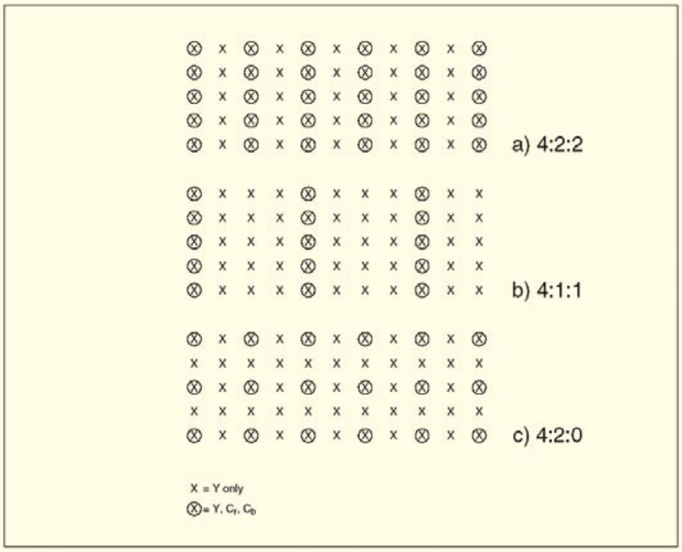

Fig 12 a) Post production standard 4:2:2 sampling in two dimensions.
b) 4:1:1 sampling cuts the horizontal colour bandwidth.
c) 4:2:0 sampling matches horizontal and vertical spacing.

won't work as some bit patterns produce signals that are difficult to handle. The solution is to use a technique called scrambling which breaks up these awkward bit patterns before sending. A descrambling process at the receiver completely restores the original bit patterns.

SDI signals have such a high frequency that they suffer from cable loss. The further they go, the smaller the signal gets and the more jitter it suffers. This is normal, and the system is designed around the reclocking technique mentioned in the previous section so that the losses and jitter are cut out. However, for reclocking to work reliably, the signal cannot be so small that the slicer can't decide whether it's high or low nor can it be so jittery that the phase locked loop can't figure out what to do. Fig.14.a) shows what cable lengths can be used in SDI. These are not negotiable so don't exceed them. Fig 14.b) shows what happens if you do. The performance of SDI drops dramatically after the maximum design distance is passed. The change is so dramatic engineers call it a ***crash knee***. For reliability you have to leave some margin in an SDI cable installation to keep away from the end of the world.

So how do you measure the margin? This is easy. When the system is first hooked up and working, simply extend the length of each cable in turn with another 15 m or so of cable. If the system still works fine, the margin is OK. If the system fails the margin is too small. The cable will have to be upgraded, or a repeater installed. Before going to that expense, you might want to check that the phase locked loop in the SDI receiver is adjusted properly.

In serial digital, component and composite signals are as incompatible as they are in analogue. However, the SDI interface signal is just a bit stream, and at that level the only difference is in the bit rate. So it's perfectly possible to build dual standard routers which can work with composite or component SDI. These devices are available today. A quality router will contain a reclocker, implemented with the aid of a phase locked loop. Often all that is needed to change from composite to component is to flick a switch which changes the centre frequency of the loop. This means that you can buy a router today to upgrade a composite system to digital, and still use it when you change over to components in the future.

SDI handles digital video which has the same timing as analogue video. Thus ***where analogue video has blanking, component digital has a gap in the data***.

These gaps can be used to transmit ancillary data along with the video. Ancillary data passes through routers transparently and video-only devices ignore it. Smart devices detect the special codes which herald ancillary data and strip it out of the bit-stream. Ancillary data can also be used with composite digital, although it is limited to the space available in sync tip because composite digital interfaces still carry syncs and burst.

a)

System	Clock	Fundamental	Crash knee length	Practical length
NTSC Composite	143 MHz	71.5 MHz	400 m	320 m
PAL Composite	177 MHz	88.5 MHz	360 m	290 m
Component 601	270 MHz	135 MHz	290 m	230 m
Component 16:9	360 MHz	180 MHz	210 m	170 m

CABLE: BICC TM3205, PSF1/2, BELDEN 8281
or any cable with a loss of 8.7 dB/100 m at 100 MHz

b)

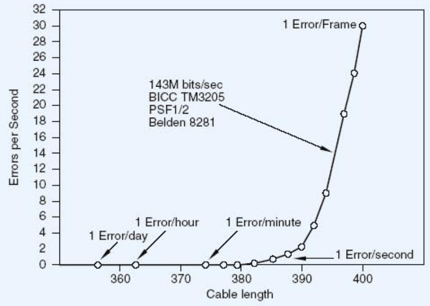

Bit error rates (BER)	NTSC	PAL	270Mb	360Mb
1 bit error/field	4.2×10^{-7}	4.2×10^{-7}	4.2×10^{-7}	4.2×10^{-7}
1 bit error/second	4.2×10^{-9}	4.2×10^{-9}	4.2×10^{-9}	4.2×10^{-9}
1 bit error/minute	4.2×10^{-10}	4.2×10^{-11}	4.2×10^{-11}	4.2×10^{-11}
1 bit error/hour	4.2×10^{-12}	4.2×10^{-12}	4.2×10^{-13}	4.2×10^{-13}
1 bit error/day	4.2×10^{-14}	4.2×10^{-14}	4.2×10^{-14}	4.2×10^{-14}
1 bit error/month	4.2×10^{-15}	4.2×10^{-15}	4.2×10^{-15}	4.2×10^{-16}
1 bit error/year	4.2×10^{-16}	4.2×10^{-16}	4.2×10^{-16}	4.2×10^{-17}

Fig 14 a) the cable lengths which can be used with SDI. Don't exceed them.
b) the error rate of SDI rises dramatically at the crash knee.
Practical systems must stay to the left of the knee to avoid flaky symptoms.

Why is decoding so difficult?

In ideal analogue TV system the composite encoding process tightly frequency interleaves the chroma into the luminance in three dimensions: vertically, horizontally and temporally. The additional vertical frequencies are evident in the two line sequence of NTSC. The additional temporal frequencies are evident in the four-field edit sequence. Real life composite signals are often non-ideal.

Cut edits along the time axis and colour transitions across and down the screen produce high frequencies which fail to interleave in the ideal manner.

Such signals are theoretically and practically impossible to decode correctly because chroma frequencies have actually extended into luminance space and vice versa. Composite video is actually a form of compression because it allows colour pictures in the same bandwidth as monochrome. These occasional un-decodable composite signals are literally compression artefacts.

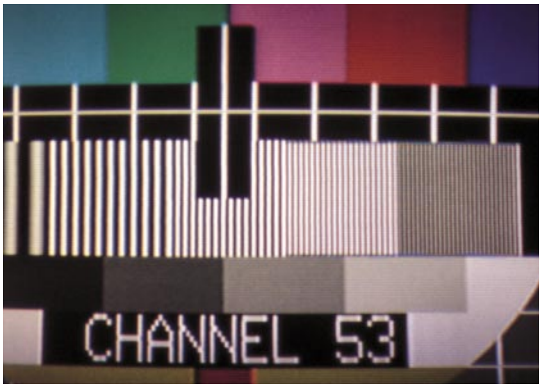

Technicalities aside, the point to grasp is that precise composite decoding is not always possible on real signals. There is no one decoding technique which is ideal. Different decoding techniques have different strengths and weaknesses and so it is not surprising that a wide variety of hardware is available. The art of decoding is to employ the right decoder for the job within the available budget.

Ideal composite video has a three dimensional frequency interleave. ***The ultimate separation of chroma and luminance can only be obtained by filtering in all three dimensions.*** Some decoders do just that. The temporal filtering process, known as a field comb, requires simultaneous access to a number of different input fields and this can only be done with a large quantity of memory and consequently such decoders cannot be made at low cost. Such techniques cannot be used on unstable signals because instabilities destroy the phase relationships of the chroma from field to field. The inevitable delay is unacceptable in some applications. The ideal composite spectrum is only obtained in the absence of rapid motion and a cut edit disturbs the three-dimensional spectrum. It is interpreted by a field based decoder as enormous motion which causes difficulty.

If the three dimensional filtering approach is not appropriate, the next alternative is to use vertical filtering which depends upon the vertical cancellation of chroma from one picture line to the next. By providing a filter with a number of line delays, a vertical filter, known as a line comb, can be made having good performance without causing a substantial delay to the signal.

It is well known that such filters fail on vertical chroma changes, where the phase relationship from line to line is disturbed. This is the equivalent in the vertical axis of the cut edit in the time axis. The result on screen is called "hanging dots" which are caused by residual chroma which has broken through into the luminance.

Traditionally, line comb decoders adapt to this situation by reverting to steep-cut low pass filtering for the luminance signal. This avoids the chroma breakthrough, but results in a low resolution picture. Snell and Wilcox have developed a superior alternative known as a Gate circuit which is used after the line comb. This patented system contains a notch filter, centred on subcarrier frequency, whose depth and width can be varied over a wide range. The notch is controlled by the degree of chroma breakthrough in a separate comb filter. In this way high frequency luminance is not lost unnecessarily.

Y/C separation cannot be performed by a field or line comb in the case of unstable signals such as VCR replay. In these cases the usual filters must be bypassed and all Y/C separation is performed by the Gate alone.

In the real world decoding is only one of the processes which need to be performed. Input signals will not necessarily be genlocked to station reference and a frame synchronising stage may be needed. Composite input signals are predominantly analogue, but composite digital signals from D-2 or D-3 recorders may also need to be handled. Production steps are increasingly carried out using component digital 4:2:2 signals and it becomes logical to include the analogue-to-digital conversion in a decoder. If the conversion to digital is at the input of the decoder, the filtering can all be done digitally. [•]