

Digital Intermediate

A Real World Guide to the DI Process



Steve Shaw
Digital Praxis Ltd.

[Revision - 11 Jan. 06]

Setting the Scene

This document is one of three that continue to evolve and grow as the world of Digital Intermediate becomes more mature and new systems, processes and ideas come to light. The documents are a distillation (although long winded in parts) of the experience gained by Digital Praxis in its work with clients building and operating real-world DI environments.

The hope is that through the information contained herein it will become possible for an average industry player to understand and gain the most from a Digital Intermediate process, whether it be for a cinematic 'film' project or within the realm of tv, commercials or promo's - as will be seen the DI approach is applicable to all projects, film & video.

It is also hoped that this information will be of benefit to end users (cinematographers, directors and producers) as well as post-production facilities building and operating DI environments.

While these documents do include, by necessity, an amount of overtly technical information, they have been written as a pragmatic approach to the requirements of DI based on real-world needs and applications experienced by the author, who has far more years then he is keen to admit of working within the digital film industry.

Many liberties have been taken - not with intellectual property - just with the depth of technical information that is gone in to. The author feels people often spend too much time researching absolutes when this is an imperfect world we live and work in.

At some point good enough is actually good enough - and that is where business pragmatism should overcome technical purity and where this document is aimed. If you are expecting a philosophical debate on the purity of technological approaches to DI you will be disappointed.

However, if you are looking for a realistic approach to the business and technology behind a commercially viable DI environment I hope these documents are of interest and use.

This first document, *Digital Intermediate: A Real World Guide to the DI Process*, is an attempt to discuss to various technical parameters behind the DI process, including cinematography, post-production, & deliverables.

The second document in the series discusses the more creative aspects of the DI process; a Scene-to-Screen process if you will, and is called *Digital Film, Scene-to-screen: a user's guide to Digital Production/Post-Production*.

The third document is a user's view of Quantel's iQ Pablo DI system - and associate periphery equipment - and is called *The iQ-Pablo Digital Intermediate System: A Real World Guide to iQ Pablo and the DI Process*.

It must also be said there is some serious overlap across the 3 documents - for which I make no excuses. If someone was paying me to write this stuff I might be more careful!

As always, I appreciate feedback & comment and can be found as follows.

Steve Shaw
Digital Praxis Ltd.
+44 (0)7765 400 908
steve@digitalpraxis.net
www.digitalpraxis.net

These documents are produced with thanks to all of who companies that have allowed Digital Praxis to be part of their DI operations in one form or another. You know who you are!

Contents

Introduction to the DI process

Digital Intermediate vs. Digital Film

- Digital Film
- Digital Intermediate or Digital Film?

Digital Intermediate vs. Traditional Chemical Labs

- Traditional Chemical Lab Workflow (simplified)
- Typical DI Workflow (simplified)

Digital Intermediate vs. Traditional Telecine

- Dynamic Range Compromise

DI Capabilities

- More than just colour correction

DI Benefits

- Typical DI Environment
- Benefits for the Director
- Benefits for the DoP
- Benefits for the Producer
- Benefits for Production

DI Applications

- Applications

DI Parameters

Dynamic Range

- OCN film contrast vs. Print film
- Digital Bit Depth
 - 10bit Log vs. 13bit Linear
 - Example of Lin & Log Sampling
 - Linear Data in a Log World?
 - Linear Data grading restrictions vs. Log Data
 - Log Data viewable contrast range Print LUT
 - TV Gamma
 - HDRI vs. 10bit Log
 - RAW
 - DCI - XYZ

Resolution

- MTF Measurements
 - ITU horizontal MTF measurements
 - ITU vertical MTF measurements
 - Actual MTF values for the film process
 - Theoretical CCD MTF values
 - Actual Measured Film & CCD MTF values
- Super 2K
- 4K

DI Workflow

- The Calibration Chain
- File Format

Film Scanners

- Quality Control
- Scanning Requirements
- Scan Parameters
 - Cineon/DPX Specification
- Scanning Systems

- Scanners
- Data Telecines
- Film Scanning Procedure

Digital Cinematography

- Technical Considerations
 - Resolution
 - Dynamic Range
 - Extended Dynamic Range
 - The Benefits
 - Compression
 - YUV vs. RGB & Compression
 - Field of Focus
 - Image Framing
 - Shooting Digital
 - Digital Cinematography Rough Shooting Guide

DI Workstations

- Monitors
 - Preparation
 - Reference Material
 - The Grading Environment
 - The Monitor
 - Colour Temp Coordinates
 - Contrast, Brightness & Chroma
 - LUTs
 - Monitor Calibration Procedure Guide
- Digital Projectors
 - Projector Calibration

Film Recorders

- Film Recoding Process
 - DPX or CIN?
 - Controlling the Recording Process

The Film Lab

- Matching the Film Lab
- Printer Lights
 - Printer Lights & DI
- Being Practical

Video Deliverables

Film Based Workflow Descriptions

- DI Lab Workflow Uncut Neg.
- DI Lab Workflow Cut Neg.
- The Future for Dailies?

Data wrangling

Offline & Timecode

- Timecode Conforming

PAL Offline
NTSC Offline
Telecine Data Transfers
Motion effects & VFX

Existing DI Operations

Virtual Telecines
Re-purposed VFX Systems
Distributed Workflow Systems
Shared Area Network Systems (SANs)
Quantel's iQ Pablo

Introduction to the Digital Intermediate process

The move to DI as an alternative to the chemical lab, and more, is now well under way. This is not just because it is a cleaner, quicker and more flexible alternative to the traditional chemical approach, but also because it is a far better fit for the modern media world.

Films are rarely viewed only in the 'film' cinema; a number of cinemas have already gone digital -with more following daily- and many people see films at home via high quality DVD or via digital broadcasting, both SD and increasingly HD.

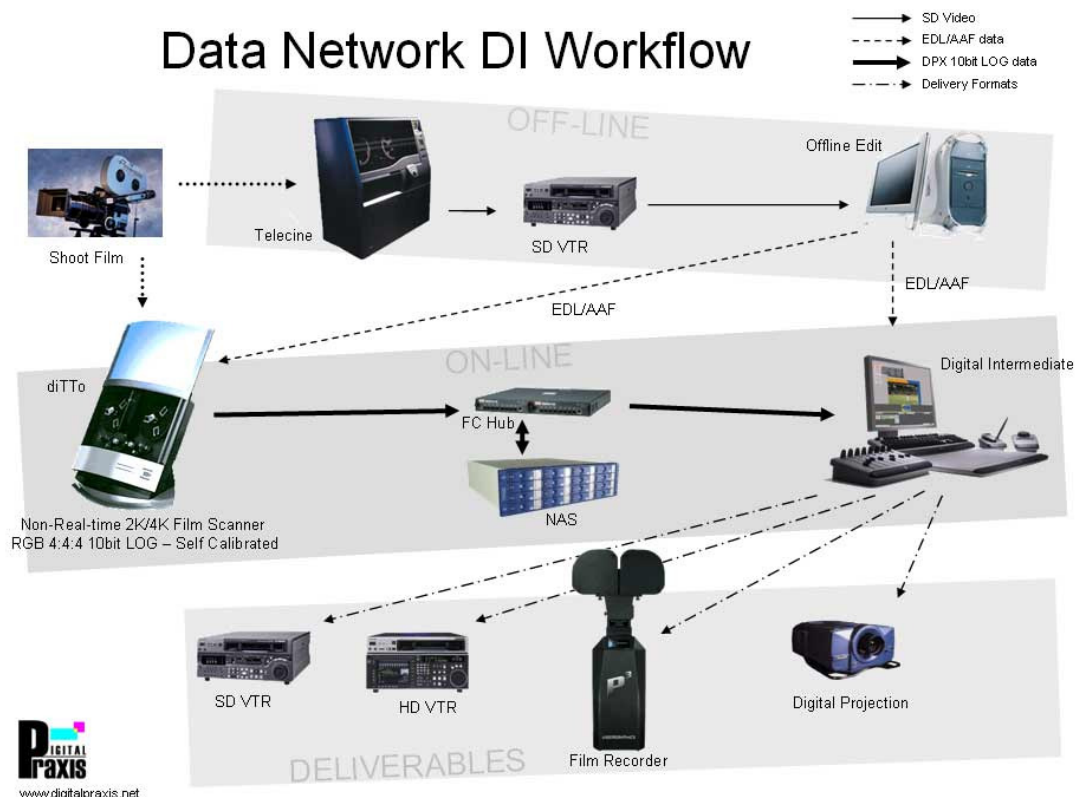
DI creates an edited and graded digital master from which it is easy to make the highest quality film deliverables as well as the highest quality versions for all other media.

In addition, DVD's tend to offer more than a simple home viewing version of a movie, containing additional content such as deleted scenes, director's cuts & extended versions, making-of documentaries, interviews, games, etc. making DI an appealing and effective operation in streamlining DVD authoring requirements.

At the same time a well-designed DI system can offer an easy and cost-effective path to create different versions, such as for airlines, foreign language distribution, etc., including trailers and additional promotional requirements for the movie.

Additionally, the DI workflow betters the traditional video approach to colour correction via telecine, providing both a more flexible and creative environment, as well as providing for all colour correction to be performed after editing - in-context - enabling far better end results. From a business standpoint it is more cost-effective.


In time, all Post-Production will be performed via a DI workflow, regardless of the source material being worked on, or the final deliverable destination.



Digital Intermediate vs. Digital Film

The term Digital Intermediate describes a process that is nominally one third of the operation loosely termed Digital Film. The other two thirds encompass acquisition [the obtaining of the source material through some form of capture] and presentation [the distribution, projection and/or transmission of the final result - also known as D-Cinema or e-Cinema, although D-Cinema tends to be applied to films [cinematic projections] and e-Cinema to alternative content [sports events, concerts, corporate, etc. projections].


Digital Film



DI itself encompasses the post-production stage, using digital processes to perform editing, grading, text, vfx and more, generating a digital master ready for the various distribution & delivery processes.

Note: DI is not just about colour correction, although a lot of 'DI' systems out there are capable of nothing more than colour correction! Be sure you understand a systems capabilities before embarking on a DI project.

Digital Intermediate or Digital Film?

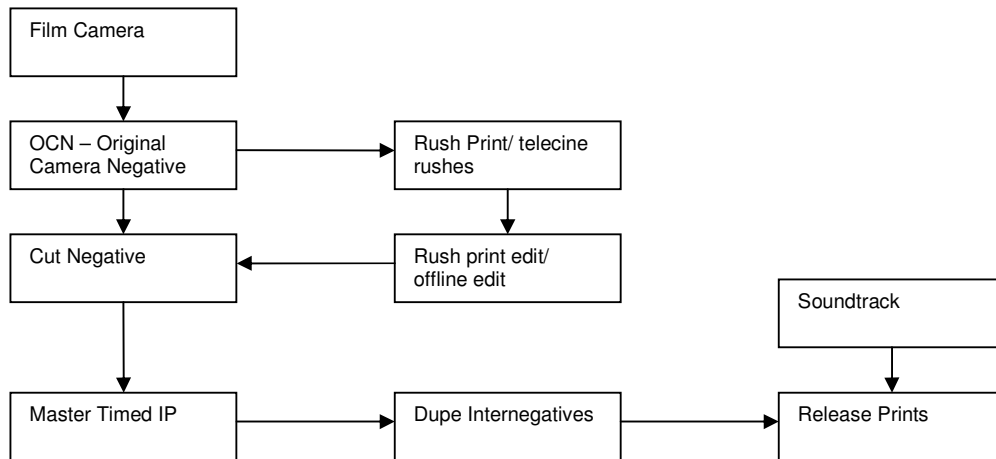


A major point about these three processes is that they are all independent of one another, and inclusion of one process in a project doesn't immediately require the involvement of the other two. Therefore, digitally acquired material can just as easily be projected via traditional celluloid means as via a digital projector and equally as likely film originated material may be shown digitally via a digital projector. The Digital Intermediate [DI] process is therefore applicable for traditional film projects [acquired and presented] as well as those classified as 'digital', through the use of digital capture or projection techniques. The DI process is chosen [or should be] on the benefits it provides for the project in question, not the use of any additional digital technology elsewhere within the project.

Digital Intermediate vs. Traditional Chemical Lab

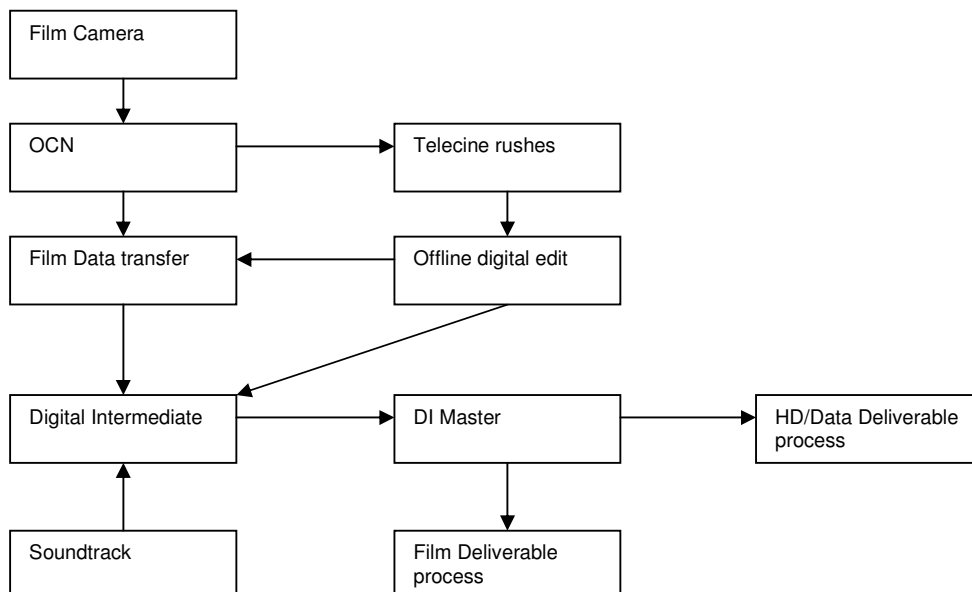
The traditional chemical film post-production chain accepts original camera negative (OCN) and, after copying, editing, grading, optical effects, etc. produces a number of internegatives, via interpositives, from which the release prints are made. An interesting point is that the copying processes incur quality losses, so some detail [a lot!] of the OCN is lost during processing. This is discussed later when talking about resolution.

Traditional Chemical Lab Workflow (simplified)



DI realises a digital form of the chemical lab, producing a timed intermediate from which distribution prints can be generated. The key to realising Digital Intermediate is to edit, grade and add vfx in the digital domain, requiring the film negative to be scanned once, and thereafter no further manipulation of the negative is required. From this point onwards the DI process is lossless, up until the optical generation of release prints.

Typical DI workflow (simplified)



Digital Intermediate vs. Traditional Telecine

The traditional telecine approach to film transfer and grading usually involves working with the OCN, or dupe negative, transferring the film images to video via the creative tools available to the telecine colourist.

While this provides a greater level of creative tools than available via chemical lab grading, enabling multi-layer colour correction, power-windows, selective adjustments, etc, etc, there are inherent problems with the 'telecine room', especially as deliverable requirements move away from standard definition video (SD), into higher resolution image formats including high definition (HD) and 2K, and even 4K data.

Dynamic Range Compromise

At high resolutions the compromises inherent in the telecine approach to colour correction become very apparent – specifically because of limitations in image dynamic range.

This is because of the generation of direct 'deliverable' colourimetry & dynamic range during the video transfer, meaning there is no room for later grading flexibility as all headroom has been removed.

Additionally, the 'telecine' grade is performed on clips out of context; clips are not graded with reference to surrounding shots within the final edit, so any graded clip cannot be seen relative to surrounding shots within the final edit.

The telecine toolset is also designed for 'real-time' video colour correction, rather than the more flexible DI 'data' based approach, which has its own creative compromises as it limits the available real-time processes based on the available hardware - if you need more secondaries than are available, tough!

These issues make the final result a compromise, rather than the ideal end product.

DI Capabilities

One of the biggest misconceptions of Digital Intermediate operations is that they are totally focused on colour correction. This view has been sustained, in part, by system manufacturers themselves as many offerings purporting to be DI systems are capable only of colour correction.

As was said previously *DI is not just about colour correction, although a lot of 'DI' systems out there are capable of nothing more than colour correction.*

More than just Colour Correction

This misnomer of DI systems being 'colour' only can be the biggest restriction for DI operations built around this single concept. The reality is that true DI requires many tools and capabilities if it is to be able to perform creative work in a cost effective way, and as a viable business, including:

- Full editorial capabilities, including edl autoconforming, manual edit changes, wipes, dissolves and custom transitions, re-order edits, replace shots, etc.
- Compare online to offline via visual comparison – i.e. real-time and simultaneous split screen between on-line & off-line.
- Text capabilities for lower thirds, titles and credits.
- Paint and retouch for image repair.
- Ability to work with any image format, including high resolution 2K & 4K data, HD video, SD video, progressive & interlaced, etc.
- Full colour correction including primary correction, secondary, selective control, tracking based correction, etc.
- Unlimited creative processes with no restriction on layers, selectives, windows, etc.
- Immediate playback of changes, grades, effects, etc., for review & comparison.
- Versioning and play out in real-time, including Pan & Scan in any video format.

DI Benefits

As can be seen above, the digital Intermediate process, in its most basic form, can be considered as a replacement for the opto/chemical film lab environment offering digital neg. cutting from offline edit information, optical transition processes such as wipes & dissolves, film grading, text, image repair, etc. aimed at producing the equivalent of a timed IN [internegative], assuming a traditional 'film print' distribution model.

This base level of operation is the heart of any Digital Intermediate Lab, and through its ability to provide enhancements throughout the process can be a strong enough reason to be chosen in preference to the traditional chemical lab.

The additional benefits it offers can make its choice yet more compelling, and range from almost unlimited flexibility, with multiple variations available for assessment immediately, without need to process and project film, to a guarantee of quality regardless of the number or complexity of optical process undergone.

To be able to sit within a DI environment and see, interactively, changes as they are made, is a film Producer's/Director's/DoP's dream. To be able to try a dissolve between two scenes, followed by an alternate wipe to assess the difference is an impossibility outside of the DI environment for film projects.

Additionally, a DI environment is able to provide all deliverables from a single DI master, including: film internegative master, digital projection master, HD and SD video masters, DVD master, trailers and promo material including high and low resolution requirements such as print media and web content.

In general, the things taken for granted for years within the long-form video market have been all but impossible for film based projects, but through the use of DI are now available to all.

And have you ever attempted to sit & work at a traditional film grading or timing station? The operator's screen on a Hazeltine or Colormaster Optical Colour Analyzer system is about 6" wide, often of indeterminate playback speed with variable colourimetry compared with the final print, requiring 'interpretation' by the film timer.

Compare that to grading via a 32" HD film-colourimetry monitor, or better still, a D-Cinema digital projector on a full size theatre screen with real-time 2K or even 4K playback, all fully colour calibrated to the final film deliverable.

Typical DI Environment



The DI suite at FotoKem, LA.

Benefits for the Director

- Ability to see the entire project at an early stage
- Complete flexibility for change after the offline is finished
- Review changes instantly on a large screen
- Greatly improved colour enhancement tools without going through the vfx pipeline
- Simultaneous multiple versions including the 'Director's Cut'

Benefits for the DoP

- Cinematic editorial and grading environment
- Retain complete control over the look and feel throughout post-production
- Sophisticated colour enhancement available on every shot
- Complete colour control over different format deliverables
- Maintain high quality irrespective of the number of 'opticals'

Benefits for the Producer

- Reduce the time taken for post-production
- Reduce the costs associated with late changes
- Quickly produce multiple versions for audience testing
- Quickly produce all the different format deliverables

Benefits for Production

- Reduce total post-production costs and timescales through single mastering for all deliverables
- Easy project tracking and management
- Enhanced quality control for all deliverables
- Improved security capabilities

Beyond such basic benefits the DI environment offers advanced capabilities that have previously been held within the domain of the vfx film.

For some considerable time film effects work has been carried out almost exclusively by digital techniques. However, the only films that have gained access to such tools have been comparatively high budget projects, produced with digital vfx written into the storyline. With the DI environment projects that would traditionally have never budgeted for digital are now able to take advantage of the benefits digital offers, without having to find ways to increase the project budget.

To enable this has taken a major change in the underlying digital technology to enable a digital workflow akin to that found within traditional labs, effectively moving away from the previous distributed workflow environment of digital vfx to a more integrated [hero] environment enjoyed within chemical labs and video long-form post houses.

In simplified form a DI environment needs to be a centralised & self-contained environment enabling all necessary DI functions to be carried out without the need to move operations. This is explained further later in this document, but suffice to say if you are moving data between environments - allowing for the need to relocate complex vfx shots - you are losing time & money.

DI Applications

While to many people the term Digital Intermediate immediately refers to the generation of a full feature film using digital post-production technology this is not an exclusive application, although it is the arena within which DI has initially evolved.

The processes used within the DI 'film' world are just as applicable in the commercial, promo & broadcast world where the prospect of enhanced quality and creativity combined with ultimate flexibility is extremely appealing, regardless of the destination format.

Therefore, the contents of this document can be applied throughout the moving image world, not just for 'film' applications, although it is with this primary market in mind that this document has been written.


Applications

- Feature film post-production and mastering
- Trailer post-production
- Commercials post-production - TV and cinematic
- Music video post-production
- Promo post-production - for cinema and TV
- Feature film restoration
- Broadcast episodic series & mini-series
- Multiple release format single mastering - film, D-Cinema, DVD, HD, SD, etc. from the one DI master

DI Parameters

Having introduced the concept of DI, the process itself requires more understanding. It's not enough to simply describe it as a Digital Lab:

We need to know how image information can be got into the DI system; what the parameters are within which the DI process must be able to operate if it is to function acceptably; how the entire process can be calibrated; what toolset is required to perform the necessary DI functions; and what steps need to be taken to get material out of the system.



If we accept that the DI process aims to digitally master a film in its entire form there are a number of parameters any DI system must be able to meet, assuming major compromise is not acceptable.

- Provide controlled input to output calibration, including working with different colour space, bit depth, resolution & format images.
- Hold an entire film [with handles and variants] in full film resolution [what ever that is - see later].
- Autoconform from an offline edl with ability to compare online to offline for verification.
- Playback the entire film project in real-time, with temps/animatics/previews as necessary, at full resolution.
- Display the film images in true 'film' colourimetry, contrast, gamma, etc.
- Respond immediately to changes for review and sign-off.
- Perform the basic functions necessary for a DI operation to proceed including editing, opticals, colour correction, vfx work, titles, etc.
- Work with accepted technical image parameters to present a final image quality acceptable as 'film' [?].
- Etc.


As is expected, such a list of requirements raises more questions than it answers, with 'what exactly is digital film?' being topmost amongst them.

What follows is a discussion on the main point of interest, evaluating possible options and approaches including technical reasoning as well as realistic business requirements.

Dynamic Range

Original camera negative (OCN) has a dynamic range capable of recording detail both well into shadow [blacks] as well as into highlights [whites]. Such dynamic range enables decisions on contrast, brightness & colour balance, in RGB [CMY] to be made during the post-production [film lab] process, ensuring the DoP's views are seen in the final film print. It is the print film stock that extracts from the wide dynamic range camera negative a realistic dynamic range suitable for human viewing.

OCN film contrast vs. Print film



The above diagram shows the difference between OCN's (Original Camera Negative's) low contrast capture and Print film's high contrast final deliverable. Note that the combination of the two contrast values (0.6 for the negative x 3.0 for the print) results in a final image contrast of above 1 (1.8) which means that the final image has a higher contrast ratio than the original scene! This is required to overcome the less than optimum viewing conditions encountered within a film theatre.

OCN




Print




The reason for capturing a low contrast image is that if in-camera grading is performed it is possible to clip or crush lowlight or highlight detail (under expose or over expose) which it is then impossible to regain in post-production. And it's only in post-production where the real interaction of shots can be seen as it's only here they are viewed in their real final edit context.

The 'print' process is what extracts the 'visible range' from the high-dynamic range OCN by 'moving' the print curve up and down the captured low contrast image data.

Moving the print curve into the shadow detail will brighten, or print up, the image.



Moving the print curve into the highlight detail will darken, or print down, the image.



This process, from negative to print, occurs because at no point during the shooting process can anyone be sure what actual exposure level the film is actually recording - although a good DoP will have a good idea. But by having a wide dynamic range captured by the negative any exposure errors can be corrected in the lab during the grading process.

Digital Bit Depth

As much DI work will originate from film camera negative, until Digital Cinematography becomes the norm, it makes sense that the DI process offers a similar level of dynamic range. This can be through the use of 10 bit Log data or 13 bit [or greater] Lin data to maintain the full dynamic range of OCN film stock, with a good level of granularity [brightness change per digital sample], although 10 bit Lin. may be acceptable in some circumstances due to the 'What-You-See-Is-What-You-Get' nature of the digital DI process – but I wouldn't recommend this approach. The final decision depends on the approach being taken to the DI process and at what point a 'timed digital internegative/positive' is being produced. For flexibility and security this is usually the final stage [colour correction being performed after all editing and vfx work has been completed] requiring wide dynamic range source material.

There are relatively few disagreements on this, allowing for pragmatism and business requirements to shape the needs, with a 10bit Log workflow being the norm.

10bit Log vs. 10bit Linear

There is a lot of confusion regarding the use of Log or Lin data. In truth the question should be what is the dynamic range that is best to use?, and then what is the best method to hold that dynamic range?

As has been discussed already, original camera negative film has very wide latitude, which is another way to say it has the ability to simultaneously see detail deep into shadows and high into bright areas of an image. To experience what this means try the following.

Stand facing a window opening out on to a bright day. If you look out the window you will be unable to see [via your peripheral vision] detail surrounding the bright window - everything will appear dark and detail-less. If you chose to look away from the window [to one side] you will now be able to see detail within the room but outside the window will now appear blown-out, without any detail. This is because the combination of deep shadow and bright sunlight is beyond the ability of the human eye to capture. The eye's dynamic range is too small. Original camera negative film on the other hand can see both simultaneously.

Film also sees (records) illumination (brightness) as logarithmic information, which is also true of the human eye. That is to say that within shadow detail a small change in illumination is easily seen but that the same level of change in bright areas hardly registers.

Imagine being in a dark room and someone across the other side strikes a match. It will be immediately apparent. However, if the room is now filled with sunlight and the same match is struck (ok, so a new one as the first is all used up!) it will be all but invisible...

The problem is that the devices used to turn image data into digital information (CCD's - Charge Coupled Devices) output an analogue voltage signal that is a linear response to the illumination they see. This is then converted into a digital signal that represents the linear analogue signal - linear digital data.

As small changes in shadow detail are easy to see the illumination change per digital sample needs to be very small but as the digital conversion is linear this means that the same number of samples per illumination change will be used in highlight areas too - with the result that many samples in the highlight areas will be redundant as the perceived changes in brightness will be too small for the human eye to see.

If the dynamic range being digitised is not large this is not too much of a problem, which is what happens with normal digital video cameras where 4 to 5 stops of illumination information is captured. With 10 bit data providing 1,023 samples there are enough to represent shadow detail accurately and not too many wasted in highlight areas.


If the dynamic range is the same as original camera negative film - 10 to 11 stops - it will take 13 bit linear data providing a staggering 8,192 samples to match the granularity of each sample in the shadow detail to that of the lower dynamic range video image, with a lot of redundant samples in the highlights. This makes the data file huge in comparison to the 10 bit linear file...

However, if the large dynamic range information is converted into digital Log data only the necessary samples will be maintained, with plenty of samples available both in the shadows and highlights. In this way 10 bit Log data can faithfully represent the full 11 stops of illumination information!


This is the basis for the Kodak specified Cineon (.CIN) digital film file format where each digital sample represents 0.002 density of a camera original negative, giving a D-min to D-max range of 2.0D, which is below the threshold of human sight perception [i.e. you never see digital 'banding' or steps between sample value - LSB - 'least significant bit' changes].

In this way Log information can be viewed as a form of loss less compression and is the ideal data format for digital film DI work.

Example of Lin and Log sampling



Linear sampling with an excess of samples within the highlight detail.



Log sampling with evenly spaced samples throughout the scene exposure range.

The above diagrams show the difference between Lin sampling of an original scene image and Log sampling.

In the linear sample diagram a lot of the digital samples in the highlight region are wasted as the human eye cannot perceive a difference between them.

In the Log sample diagram the human eye can perceive a useful amount of change throughout the exposure range, with no redundant samples.

Linear Data in a Log World?

Additionally, the ability of Log data to hold extended brightness range information can be used as a benefit when working with image data captured via a linear source such as an HD camera.

This is a bit difficult to explain, so bear with me...

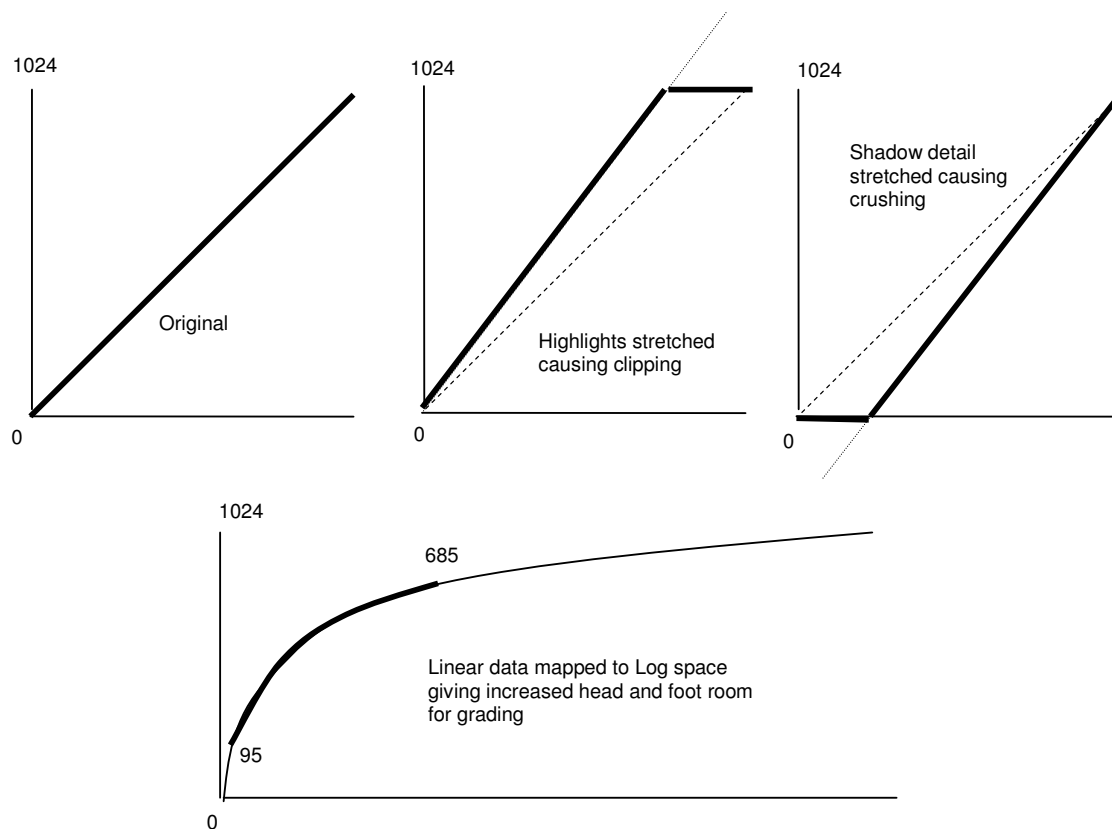
Linear capture devices, such as most HD cameras, all SD cameras and telecines when not set to Log mode, capture the scene/image they are looking at by setting artificial black & white points. Any detail below or above these points is clipped. However, with careful control this clipping can be minimised on all but the widest dynamic range scenes/images. But, all the available data samples are now used up in holding the initial image so if any grading is performed within the DI environment in the same linear space additional clipping will occur as the image is pushed higher or lower (see below).

However, if the linear source data is mapped into Log space the initial linear black 0 (zero) point gets set to 95 and the initial white point (1,023) to 685, providing huge additional head & foot room for grading. Remember that you have also moved from a linear colour space to Log colour space so the new granularity of sample [brightness change per digital sample] has effectively improved throughout the range too.

The result is much better graded images.


Obviously, with today's 'linear' HD cameras such as the Sony f900/950 & Panasonic Varicam, Cine Gamma Curves enable a 'Log' capture profile, improving capture dynamic range. See CvpFileEditor page on www.digitalpraxis.net. And then there's Viper's FilmStream mode for the ultimate in Digital Cinematography, or possibly the Panasonic Genesis, or Arri D20, or even Dalsa Origin, with more true Digital Cinematography cameras to come...

Linear Data grading restrictions vs. Log Data



Even if not going back to film, the use of print D-log E style log to Lin conversion LUTs at the output stage, when recording to the vtr, can maintain more of the new head room than normal simple Log to Lin conversions, where 95 is again clipped to black and 685 to white.

Log Data to viewable contrast range Print LUT (Log-to-Print LUT)



The above diagram shows how the Log to viewable contrast range (film print!) LUT is used to alter the Log data to give the image the correct contrast for human viewing. Note this is not a Log to Lin LUT as that would simply 'clip' the Log data at the ends of the 'Average Visible Exposure Range'. Rather, the Log to 'Print' LUT uses a filmic toe and shoulder characteristic to roll-off the extended dynamic range without clipping or crushing. This is a major reason why 'linear video' looks inferior to film - the blacks are crushed & the whites clipped!


It is now easy to see how 'moving' the Output LUT up and down the Log data curve (right to left on the diagram) will either increase or decrease the perceived exposure of the LUT output, making the image seem brighter or darker overall. This is the basis of all 'film' grading, as was described above, and is mimicked by the View LUT process.

TV Gamma


Note: although the above description discusses Log vs. Lin, in truth what is often referred to as 'linear' is in reality TV gamma, which has a non-linear response which goes some way to dealing with the issues mentioned here for pure linear data – part way between Lin and Log if you like.

TV gamma is an adjustment applied to an image to overcome the fact display CRT's (televisions) don't produce a light intensity output that is proportional to the input signal voltage, with 50% input signal giving only 18% intensity output. As camera systems produce output voltages proportional to scene intensity (linear) a gamma correction is required when displaying the image...

CRT Gamma (1/2.2)





TV Gamma correction (approx 2.5)




The Gamma correction value is approx 2.5, but does vary from system to system, with Macintosh computers being the notable different system with a Gamma correction of 1.8, and is based on printers, not CTR displays... The gamma correction value is also not a 100% inverse of the CTR gamma value as it also compensates for the apparent reduction of contrast which occurs when a TV is viewed in a typically living room environment.

This is much like the 'enhanced' contrast applied by a film print image (see high-contrast Log to Print LUT above, and additional information later).

If an image were to be viewed in true Linear space it would appear very dark, as can be seen with the image on the left. The image on the right has a display (correction) Gamma of 2.5 (0.4) applied for correct viewing.



Additionally, tv video legal levels set white to 940 and black to 65 allowing some headroom but that cannot be seen in a standard viewing environment - i.e. at home - as any values from 0 to 65 & 940 to 1,023 are illegal and are clipped. This is not true of 'data' systems where all levels are valid.



However, to go into this in depth will confuse the issue... feel free to research further at your leisure. Gamma is a very complex issue!

But, what we have defined in this section is a need for a high dynamic range workspace, ideally covering 11 camera stops, as is provided by 10bit Log data.

HDRI vs. 10Bit Log

Yet another 'bit-depth' format has been introduced in the form of 16bit or even 24bit Linear HDRI (High Dynamic Range Imagery).

The problem with this format – introduced for 3D high dynamic range work and compound images from digital stills cameras – is that 16bit files require an unbelievable 65,536 samples per colour, which as we have shown above is a totally unnecessary amount of information, even for the wide dynamic range capable with film OCN. And 24bit gives rise to 16,777,216 samples per colour!

Compare both these numbers to 10bit Log's 1024 samples per colour for an equivalent dynamic range.

For this reason alone the desire to use the HDRI image format for DI work is something I can't understand.

For 3D work where the systems are unable to work with LOG data I can see a benefit, but then I can't understand why 3D systems can't be programmed to work with, and understand, 10bit LOG images?

As we have seen previously, tv images are gamma corrected, which is a form of Log encoding, so what's the problem?

However, if 3D HDRI images are to be used within a DI project it will be necessary to first map the HDRI images into 10bit log space, using a calibrated LUT, taking special care of gamma values.

This must be done carefully as any errors introduced at this stage cannot be corrected later.

I'm sure there is a valid reason for HDRI to exist, but with the proven capabilities of 10bit log imagery, and its better sample distribution throughout the available dynamic range, but I struggle to understand what this is.

In some quarters there seems to be far too much focus on high bit depth numbers, rather than the needs of images in the real world.

As I have repeated often, business pragmatism is the most important feature of digital film work, and at times I think there is too much focus on technical perfection rather than realism.

I've given little time and space over to the HDRI image format as I do not see or understand its use.

Here's a challenge; if you think know the reason HDRI is preferable over 10bit log, tell me - steve@digitalpraxis.net. I'm at a loss about this...

RAW

There is also a growing use of Digital Stills Cameras for digital cinematography, specifically for stop-motion animation and time-laps photography where their comparative low cost, robustness and quality make them ideal.

And most of these stills cameras can output image data in RAW form, exactly as the cameras imaging sensor captured the data.

This format is a usually unfiltered bayer and requires processing before the image can be seen.

This processing is comparatively straight forward, with any number of software programs being available to format the raw data into uncompressed tiff format, or similar, preserving the original quality of the RAW image file.

However, there are two issues to consider here, and we'll look at the second first

The second issue is that the tiff (or similar) file will need to be accurately mapped into 10bit Log space either during the DI process, or before. The point of conversion (during or before the DI process) will be defined by the DI system used,, and actually goes some way to separating good DI systems from the also-rans, specifically because quality will be maintained if format conversions are kept to a minimum, and this obviously means combining the conversion with the DI processing if at all possible.

The first issue is that the conversion of the RAW image data into tiff, or similar, can be done in a number of ways, with a lot of user input if necessary, including camera set white balance, gamma correction, white point setting, etc.

The problem here is that all too often the 'user' can be tempted to alter too much of the RAW information during the conversion, restricting the ability for later DI manipulation. It is imperative that the best possible 'low-contrast, high-dynamic range' image is generated at this stage if the benefit of the initial RAW capture is to be maintained.

And the plethora software conversion programs available also provide a wide range of final output qualities, including final resolution and sharpness, making choice of the right program very important.

Any quality mistakes made here cannot be undone later within the DI process.

With luck, DI systems will evolve the ability to import and work with RAW images files without the need for prior format conversion, as with iQ's ability to work with the native image format generated by the Arri Tornado high-speed camera system.

Until that time, care needs to be taken in the RAW to 10bit Log conversion procedure to maximise the final image quality.

DCI - XYZ

The final specification for Digital Cinema, released by the DCI (Digital Cinema Initiative) has also added a new colour format to the mix, and has also specified a linear (gamma 2.6) 16bit format, rather than a Log format with View LUT application!

From what I can extract from the very large specification document the DCI released – and it is very large! – their reasons for their final specification are honourable, attempting to provide the biggest possible colour space for future development and enhancement, without historical restriction.

However, the best way to think of the DCI specification is as a delivery format, not a post-production one, so use the best practices possible for Digital Intermediate work, and convert to the DCI delivery specification as needed.

For more information on real-world application of the DCI specification see the Rex Beckett/UK Film Council Digital Screen Network document downloadable for the Digital Praxis website.

Resolution

Resolution is slightly different. All cameras are WYSIWYG with regard to image framing. Therefore it is very unlikely that a scene will be shot 'wider' than required to allow for a zoom in post-production [variations in full-frame, academy, 1:85, etc. accepted].

However, as with dynamic range, the traditional chemical film process of generating a final print for viewing does change the resolution of the material from that held on the original camera negative. Therefore, although the original negative may have a resolution near equal to 4K digital pixels [3K is actually more accurate – even though some people talk of 6K] the reality is that such numbers don't allow for optical losses via camera lenses, unstable film transports, etc, discussing only the 'perfect' capture ability of fine grain negative stock - not a real-world situation.

Due to these losses, and further losses throughout the film processing and duplication operation, the final projection print has a resolution more closely represented by 1K pixels. This obviously depends on the number of intermediate stages undergone and the quality of the processes used, but represents a true situation for the average release print film.

To bring a little bit of true life into this, while grading Pinocchio for digital projection [working at 1280x1024 resolution as time was so short we had to cut corners!] the director & main actor, Roberto Benigni, was surprised to find his 'frown lines' visible in the digital projected final when they were not in the film print. As a result we had to 'soften' the focus [blur it!] on the digital final [and I won't tell you how we did this but it wasn't via digital technology – ok, I will tell you; we smeared light engineering oil on the projection room glass; I kid you not!].

Therefore, what is the correct resolution for the DI process? 3K or even 4K may be a true representation of the negative, and therefore desirable for long term archive of material that may require future work, such as restoration, but is easily overkill for DI operations based on the generation of a final result for immediate distribution from a clean, new negative that can be held for future generations if necessary [there are serious doubts on the need to hold all shot camera negative as the goal for any film project is the deliverable, not the source material and all restoration processes aim to restore a deliverable - no one will pay to see a negative! And remember, the original 'view' of the director, DoP, etc, has little to do with the negative and everything to do with the final print where they can recreate what they 'saw' while shooting, which often has little to do with what is actually on the negative].

As the DI process does away with much of the need to perform multiple intermediates its output can be considered qualitatively superior and capable of maintaining resolution through processes that cause losses through the traditional chemical lab. Therefore, scanning the original negative for a 2K DI process will result in a final image of superior quality to the chemical process equivalent, making 2K, or even less, more than acceptable for the DI process, especially if Nyquist over-sampling is used to generate the 2K data for the DI process from a higher resolution scan - super 2K [see later for more information].


Note: for vfx work being intercut with original camera negative [i.e. not going through a DI process] 2K is the realistic minimum as the digitally generated output negative will go through the same dupe processes as the original negative.

MTF Measurements


Scientifically, the above approach can be verified through empiric testing. The MTF [Modulation Transfer Function] of any given process can be measured and qualitative comparisons made showing the various losses through lenses, contact printing, aerial printing, film stocks used, etc. This has been performed a number of times by industry bodies & manufacturing companies alike, and the results are available for all to see.

As an example, in 2001 the ITU undertook a study that concluded perfectly exposed and processed Kodak Vision 200T 35mm film, when duped, printed and projected in a good cinema has a resolution of just 600-750 lines - which equates to about 1.2K pixels for a 1.85 aspect ratio image.

ITU horizontal MTF measurements




ITU vertical MTF measurements




For an understandable approach to film and its characteristics look at the latest Quantel Digital Fact Book - the Digital Film section at the back, where the following information was scientifically evaluated.


Actual MTF values for the film process



Theoretical CCD MTF values



Actual Measured Film & CCD MTF values



Less scientific, but equally important, the number of films originated on digital formats of less than 2K resolution [George Lucas's Star Wars - the prequels - being an obvious example shot in HD at 1920 pixels by 1080 lines, as was Sin City, with Silence Becomes You, Zodiac & Highlander: The Source, Miami Vice, being shot via Viper] shows what is truly possible. We can all point to specific scenes where things don't look right, regardless of the film and capture format used. Quality owes a lot more to technique than technology in the majority of cases!

So we now have a system that is defined technically as 2K resolution [2048 x 1556/1536 for full aperture 35mm film] and 10 bit log [or 13 bit lin!] bit depth. Such a system will produce a result suitable for final film output for celluloid projection or for digital mastering for digital projection. The same data can also be used to generate all additional deliverables, such as video, TV, DVD, etc. making the process even more effective.


Super 2K

Although we have shown 2K to be the ideal resolution for DI work that doesn't preclude the use of higher resolutions where applicable, for example, when using large matte painting background plates for pan & scan within a 2K window.


Also, the use of 4K scanning to produce 2K images via Nyquist sampling will maintain resolution far better than a native 2K scan. This is because physically scanning at the same resolution as the theoretical maximum of the source material is not guaranteed to maintain the full resolution or detail of the original, and is referred to as Super 2K.

Consider this. You have 2048 alternate black & white lines recorded onto an academy OCN film frame. This is probably the maximum resolving power of the film showing its maximum line pairs per millimetre. If you now scan this with a 2048 pixel CCD one of 2 things is likely to happen (ok, one of a varied number of results but the two extremes will show the theory).

If the CCD cells happen to align exactly with each black and white line the digital result will be perfect as shown below.

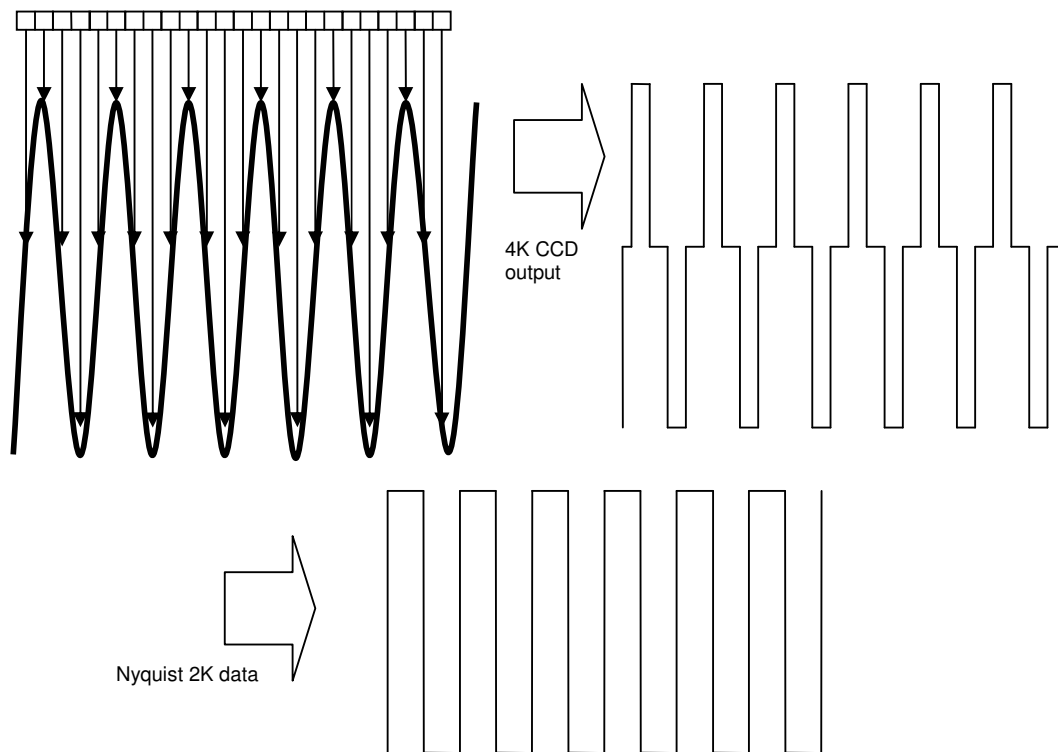


However, if the CCD cells are half a line out they will see half a black line and half a white line at the same time, with the result that the CCD will see near grey!



While this is an excessive example it does show how detail can be lost even if scanning [sampling] at the same maximum resolution of the source image.

If we now double the number of CCD cells to 4096 we are always going to see the underlying image detail and can then Nyquist sample down to the original 2K resolution while maintaining almost all of the frequency detail.



The above diagram shows a near perfect alignment but from it the reader can extrapolate what would be seen if the CCD cells were aligned differently – there would always be an output from the CCD that showed a good representation of the waveform from which a high quality 2K image can be generated.

It is worth saying though that most scanning devices use CCD cells (or flying spots!) that are not directly equal the image area they are sampling, resulting in better detail representation than one might expect. This tends to make the argument for Nyquist sampling less of an issue than it may at first appear. It's always best to perform empiric testing for one's self and define your best methodology.

Equally important from a business perspective, a 2K image requires about 12Mb of data per 10 bit log RGB frame. A 4K image requires about 48Mb of data, a 1K image 3Mb and a video SD image 1Mb. This shows the level of technology investment required to manipulate 4K vs. 2K vs. 1K, which has a direct impact on the cost of film production and therefore the profitability of any given film project based on the number of bums paying to sit on seats to see the film... Joe Public is unlikely to be willing to pay more to see a 4K processed film vs. a HD film, except for an occasional IMAX style outing! And how many IMAX theatres have succeeded in being financially successful? Make what you will of this lack of success!

4K

Having gone to some length to define the real resolution of traditional film, and show 2K (and especially Super 2K) to be the ideal resolution for DI work, the growth in 4K work is inevitable. A number of film studios are already requesting 4K DI; the DCI specification was initially 4K only, conceding 2K only after the rest of the world (non-USA) refused to make 4K a guaranteed requirement.

And the simple fact that 4K is a bigger number than 2K will ensure some clients will not be happy with anything else – exactly as we have seen with the introduction of such huge bit depth numbers as with HDRI, be they needed or not.

Regardless of the true technological requirements for digital film, 4K will happen – be ready for it!

DI Workflow

The actual process of performing a DI project goes well beyond the actual DI component of the operation, and requires a full understanding of the whole DI workflow chain if mistakes are to be avoided.

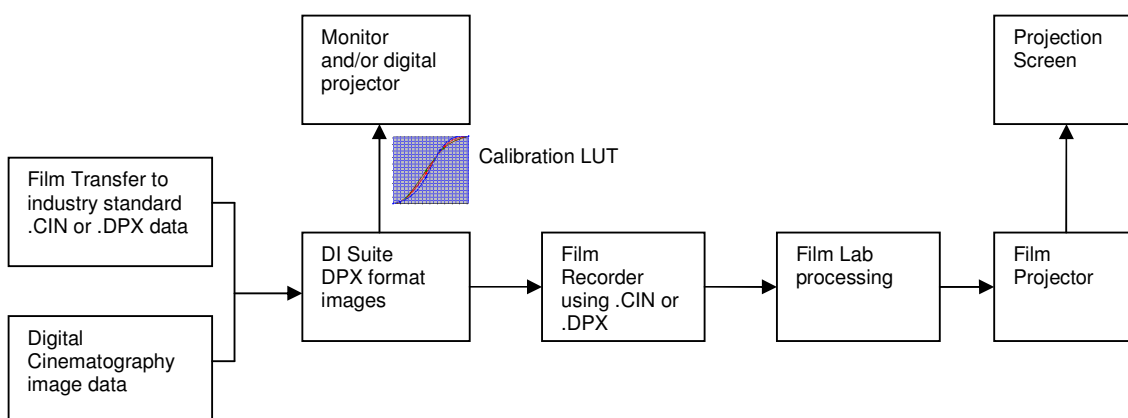
This section of the document will therefore act as a guide through the various stages, explaining the technicalities and requirements as best as possible for each stage.

The Calibration Chain

One of the 'black arts' of digital film, and an area that needs to be understood if the process is to work correctly, is the calibration of the entire DI chain, from input to output.

However, with a little logical thought this is actually a relatively simple process, and nowhere near as complicated as some people like to make out.

In reality, the only requirement is to 'match' the digital viewing monitor(s) to the final delivery format, which today means print film. The rest of the chain can be set to maintain a unity throughput, unless, and where, creative changes are made.



To make sense of this, what's needed is an understanding of the different requirements for film based photography vs. digital cinematography at the start of the chain, with an understanding of the rest of the chain from the DI workstation to the deliverables.

What it is worth pointing out up-front is that the digital part of the DI chain is very easy to calibrate. It is the chemical lab processing at the end of the chain that is more difficult to control and as the digital grading process aims to match the final chemical print colourimetry you can already see the main difficulty in end to end calibration.

Each part of the DI process can best be considered in isolation as by maintaining calibration through each part it is possible to predict the effect of any change. This is especially important if film scanning, DI, vfx & film recording are to be performed in separate locations. What is to be avoided, at all costs, is a calibration process where any error at the input stage is countered for within the output stage, or vice versa.

File Format

Ignoring the few DI operations that have their own data format and are therefore isolated islands, all DI systems are [should be] based around the Kodak developed Cineon [.CIN] file format, since ratified by SMPTE as the DPX [.DPX] format.

This format, in .CIN form, is a 10bit log format that maps each digital sample to a equivalent negative film density of 0.002d per sample, giving a total density range of 1,024 samples X 0.002 density = 2.048 d-min to d-max, which is beyond the exposure range of most scenes captured on negative film. This results in the .CIN [DPX] format capturing a digital 'clone' of the original negative, including all shadow & highlight data.

The .DPX format differs only in that it can hold Log or Lin data, depending on its header flag setting [.CIN is always Log], and also supports 4:4:4 and 4:2:2 YCbCr with both Rec.601 and Rec.709 luma encoding, and can even be used to hold YCbCr encoding as the underlying format, but with Log encoding on top! This results in fewer actual colors, but with the dynamic range of log, and it fits through a single HD SDI link. This is explored in detail later as it's a very cost-effective way of performing DI on low-budget features.

For DI work it is Log we are interested in [although it is worth noting that as of writing this document most 'scanners' do not set the log flag for .DPX headers, requiring post scan setting through a computer script or similar].

With this choice of file format each stage in the DI process need only to adhere to the format specifications for full end-to-end calibration to be maintained.

We have previously mentioned the likes of HDRI formats, but still, 10bit Log DPX meets all the necessary imaging requirements for digital film. It really is the ideal format.

Film Scanners

All film scanners are either factory calibrated to the .CIN/.DPX format [diTTo, Imagica, NorthLight, ArriScan, dataMill, etc.] or able to be set to match the format via some user intervention [data telecines such as the DSX, Millennium, Spirit, etc.].

Note: if using a telecine based scanner [DSX, Millennium, Spirit, etc.] care is needed to ensure accurate calibration as they are not 'factory calibrated'. However, the goal of any DI environment is to grade the film to gain a specific look, meaning that small calibration inaccuracies during scanning can be acceptable. However, changes/errors in contrast/dynamic range should be avoided as much as possible.

Quality Control

Controlling the quality of input scanning has a direct impact on the final image quality output from any DI system as, although the DI process will not further degrade an image or sequence, it cannot either improve on quality as measured by resolution and image detail.

Care taken in scanning will be rewarded by improved final quality and a general increase in ease of workflow and image manipulation ability.

Scanning Requirements

The general approach to Digital Intermediate means that every frame will be affected in some way by alterations in colour and contrast and possibly size and position. Unlike the more traditional use of digital for VFX shots, work performed within a DI system will not be intercut with original camera negative. For this reason the requirement for total perfection in scanning is lessened. Remember this is a real-world approach.

However, this doesn't mean no or limited care is required. Any errors made with images being cropped, crushed, clipped or otherwise lost cannot later be recovered. All the necessary information must be scanned and available for the DI Lab to do its work.

Therefore, the important parameters for scanning can be defined as follows:

Scan Parameters

- Image geometry
 - correct aspect ratio - circles not oval
 - correct image size - side to side correct
 - no image cropping - no missing image area
- Image dynamic range
 - blacks not crushed - black point correct
 - whites not clipped - white point correct
 - gamma nominal - grey point correct
- Image resolution
 - all image detail retained - no missed samples
 - image sharpness natural - electronic aperture correction minimal
 - if available 2K image resolution from 4K scans - Nyquist sampling to maintain MTF

If these parameters are maintained the scanned data will contain all the necessary information from the original camera negative enabling the DI system to perform image colour and size manipulations without restriction, as if working directly from the negative.

Cineon/DPX Specification

For optimum set-up the scanned data should match the Kodak Cineon log density map for film density to digital data. This will ensure a unity status as most digital film systems work with, and understand, film data mapped this way, making the DI process ostensibly invisible other than the specific creative work performed.

Cineon .CIN/DPX Density Mapped data

- D-Min to D-Max of 2.048
- Mapped to 10 bit Log data:
 - D-Min = 95
 - 2% black = 180
 - Digital LAD = 445
 - 18% grey = 470
 - 90% white = 685
- Granularity of 0.002 density per sample [LSB]

The active image area should also be mapped to the available 2048 [or 4096 if Nyquist sampling to 2K or performing full 4K scanning] horizontal pixels rather than the full frame area. For example, if an academy mask has been used during shooting the academy area should be mapped to 2048 pixels not the full frame aperture, which would result in an active image of 1828 pixels per line for the academy area. This maximises the resolution transfer to digital and provides a unity workflow regardless of image source. The one requirement is to be able to map the resulting 2048 image pixels to an academy film area when film recording [see Film Recording].

However, having said that, fixed sensor scanners will not allow for this re-sizing unless they include an optical zoom capability. And often this optical zoom will introduce image degradation through poor optical quality lenses. As always, it's best to try first and assess the best way to generate the final result.

Scanning Systems

There are 2 basic options for film scanning for DI work.

- Traditional film Scanners
- Data capable Telecines

Scanners

Film Scanners have been around for some years and are very accurate devices for transferring film image information into a digital form. They tend to be based on pin registration transports, are very image stable and capable of high resolutions but have tended to be slow, with one frame taking multiple seconds to scan.

For this reason Scanners have not been ideal for DI use where entire films must be transferred cost effectively and quickly. However, they are getting faster & will become the norm for DI scanning in the near future – check out Cintel's diTTo to see what's becoming possible with cost effective, fast scanners!

As scanners improve in speed, and reduce in cost, they, combined with the benefits of a DI workflow, will see the demise of telecines, or more specifically the telecine grading room, and make fast, dumb film scanners the normal film transfer device, with all grading done in post in a DI based workflow.

Data Telecines

The modern Telecine is capable of data scanning up to 2K and even 4K resolutions. As they can transfer a film frame in a fraction of a second their use for DI scanning is presently ideal.

However, they are not without their problems.

Firstly, as they do not use pin registration gates they tend to not be as stable as the original taking camera. This lack of stability is not huge [they are far more stable than any film projector!] and any instability is rarely visible on non-vfx projects, but can be enough to make vfx work difficult at 2K resolution, although tracking tools can be used in DI post-production systems to overcome this.

Note that pin reg. gates are available but slow scanning speeds dramatically - it would be unrealistic to consider 4K data scanning via a data telecine without pin reg.

Additionally, and more importantly, the use of a capstan, sprocket and edge guidance for film stability can cause image distortion around film splices, especially overlapped, tape or badly registered ones. This requires long 'tails' [1ft to 3ft] if the problem is to be avoided within active image. For modern OCN stock this is less of an issue as splices tend to be weld-joined butt splices.

Speed of transfer is another area that requires consideration with data telecines. Depending on the network connection [Ethernet, SAN fibre or HSDL] and transfer engine used [the device that translated the raw telecine data into Kodak Cineon density mapped log data - probably .DPX format but possibly .CIN] transfer speeds at 2K can range from 4 frames per second to real-time. However, for some systems the faster the transfer speed the lower the likely quality of the image as the image sensors do not have sufficient time to accurately measure the true film density resulting in image noise and a perception of lower resolution.

The ideal scanning speed for quality and throughput for most systems is approx. 4 frames per second with present data telecines, although testing of new material is always recommended as the average density of any given project [the likely over or under exposing of scenes by a given DoP] will affect the light transference during scanning and hence the quality for any given scanning speed.

It is also worth noting that the use of a sprocket to locate the vertical position of the film can show problems during 'slow' data scans that are not evident during 'normal' real-time video scanning. This is because the sprocket is a loose fit in the film perf., relying on its close contact with the trailing edge of the perf. to maintain the correct position information. When shuttling back and forward the sprocket obviously 'swaps ends' of the perf., taking up its normal trailing position when scanning commences due to inertia. However, when scanning commences at slower speed film-to-roller friction can prevent the sprocket taking up its correct trailing edge position after a fast rewind. This can cause vertical instability in the scan until correct alignment is attained; to prevent this instability, a longer than normal pre-roll time is required.

Film Scanning Procedure

For most DI applications the latest generation telecines will be used to perform 2K data scans, acting as data scanners rather than traditional telecines.

This requires the data scanner to be set-up for log scanning based on the Kodak Cineon density data map, making the data scanner a 'clone' of a traditional 'film scanner'. In this way the DI film path adheres to a known standard throughout, making it a true 'intermediate' process.

For most data scanners the use of TAF is an ideal initial set-up, looking to set 0% black [d-Min - frame bar black rather than scene black] to a digital value of 95, 90% white to 685 and LAD grey to 445 [or 18% grey to 470].

After TAF set-up, replace with the film to be transferred & re-check all settings. Accuracy here will ensure quality throughout the DI process.

However, there is no reason to stick rigidly to these values as a **small** amount of 'pre-grading' can assist the DI post production process, taking into account possible shot colour bias, lighting errors, etc. This is especially true for heavily overexposed film, where an offset from d-Min is desirable to maximise the sample range.

It is also important to correctly set masking for specific film stocks, enabling the data telecine to correctly understand the stock's cross colour component and provide an as accurate as possible RGB separation.

If these suggestions are followed, the data used for the DI process will enable the best possible result to be obtained.

Digital Cinematography

When using images captured by a digital camera for digital projection, or for video deliverables only [not for film projections] the options are slightly more varied as there is no real need to match print film colourimetry. The images can be graded as seen fit, using a monitor/projector format and colour gamut balanced to the final viewing requirements - D-Cinema or home tv.

Having said that, there is presently little requirement for digital only deliverables and so the needs/options are limited. However, this is the beginning of what will inevitably become the norm over time and thought needs to be taken with the needs and requirements as investments made today in DI technology will inevitably need to work for tomorrow's total digital cinema market.

The biggest issue with D-Cinema is that digital projectors can display a wide colour gamut. Therefore, from a very simple perspective, to what colour gamut should they be calibrated?

- to match print film?
- to match SMPTE phosphor specifications
- or to match something else entirely?

These are questions that can't be answered here, and probably will result in varied specifications globally. However, today's DI technology is such that whatever specification is chosen there are likely to be few limitations to its implementation within a top-end DI environment.

The US based DCI (Digital Cinema Initiative) has recently released its specification for D-Cinema – see www.dclmovies.com – which as a delivery format can be converted to after the DI process has been performed.

For additional, and relevant, information see the DCI explanation document written by the UK's Rex Becket for UK Post, and available as a download from www.digitalpraxis.net.

Technical Considerations

Technically, there are really two key decisions to be taken for Digital Cinematography:

Firstly, what type of digital camera to use? [We will assume HD or greater resolution here only.]

- traditional HD cameras?
- or newer digital cinematography cameras?

Traditional HD cameras include those from Sony or Panasonic, along with variations from the likes of Panavision. Such cameras can record linear YUV HD images in compressed format to camera mounted tapes or in less compressed format to tethered vtrs, or uncompressed to disc recorders.

Newer digital cinematography cameras are only just coming to market and include offerings from Sony, Thomson, Arri and Dalsa, with others on the way. They differ from HD cameras in that they tend to be RGB based with low compression but tend to require tethering to disc recorders or special vtrs. They also tend to capture dynamic ranges greater than HD cameras and can even provide 'log' outputs.

The choice depends on cinematic requirement & budget, but even with HD cameras there is no excuse today to shoot Linear Video! See later.

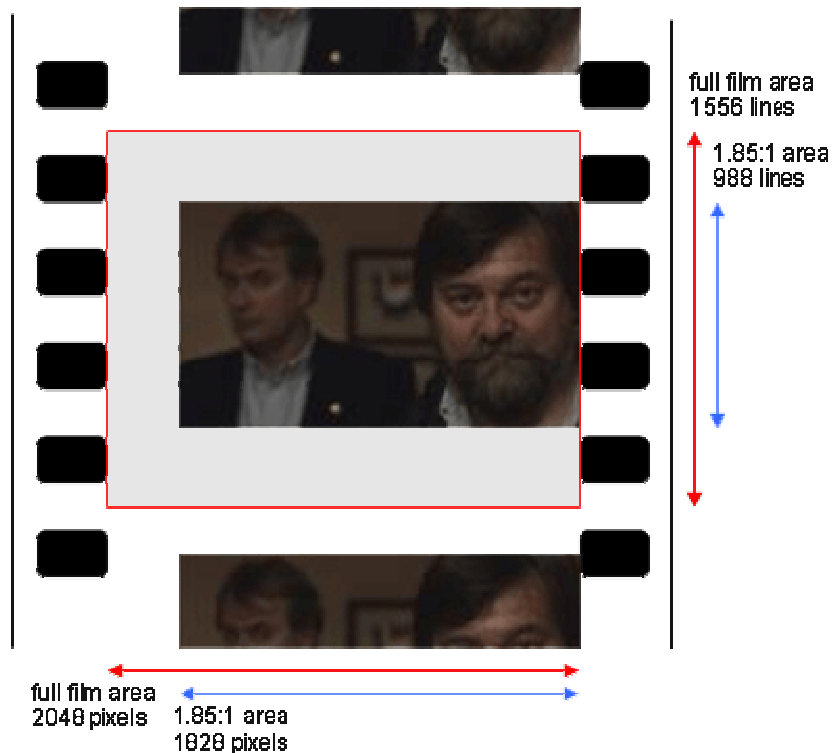
But what is obvious is that with the introduction of ever improved high-resolution digital cameras the source material for a Digital Intermediate system cannot be expected to only be film. HD and digital cinematography cameras are evolving in quality, ease of operation and capability every day, and are already considered to be easily the equal of 35mm film – the author of this document included.

Using HD digital cameras as the capture medium for 'film' projects is growing in popularity, not least because of their WYSIWYG [What-You-See-Is-What-You-Get] operation, lighter equipment, lab-less shooting and potential for lower production costs.


Resolution

The resolution of an HD frame is also amazingly close to that of a film frame, accepting 2K as a realistic acceptable digital resolution for film scanning. A full aperture [open gate] film frame is represented by 2048 pixels by 1556/1536 lines, an academy frame by 1828x1332 and a 1.85 frame by 1828x988 while a 16:9 HD frame [a 1.778 frame] is made up of 1920 pixels by 1080 lines. In a full 2K film scan, the 1.85 film aperture area is actually 1828x988 pixels, and a 16:9 1.778 image area would be 1828x1028. So in reality, which is best for resolution?

1.85 film frame image area within a full frame scan



16:9 HD frame image area



It is also worth noting that an HD camera is capturing directly from the live scene, not a pre-recorded version as with scanning from film. This has the result that maintaining detail through the use of

Nyquist sampling [over sampling] is provided by default as the source image [the scene to be captured] is totally analogue [it has no grain structure as with film] and therefore has an infinite resolution. Therefore capturing at any sample resolution will provide an image as perfect as that sample resolution allows, maximising the 'detail holding capability' of the capture resolution.

Obviously I am simplifying the issues here, for example ignoring lens attenuation, etc. but as they affect film scanning too the same holds true. An image captured via a digital camera can, in the real world, be considered to provide a digital image of greater resolution than an image scanned from a film source with the same digital capture resolution. If you like, the image source [the scene] has not been pre-filtered.

Dynamic Range

However, most HD cameras have what is viewed as a major limitation exactly because of their WYSIWYG operation; namely that they do not capture any extended dynamic range [additional black or white information]. This lack of range means that during the post-production process it is not possible to greatly alter the image appearance [contrast especially] without blacks turning grey or whites looking flat and clipped.

This is where the extended dynamic range of film has an advantage as large post-production changes [grading or timing specifically] are possible without danger of image degradation, enabling the exact view of the DoP to be generated, almost regardless of the actual exposure used during shooting.

Video Dynamic Range



Film Dynamic Range



Note: The average dynamic range of a film negative is approx. 11 stops compared to a 'linear' digital camera, which captures approx. 4 to 5 stops.

Extended Dynamic Range

Having said the above, as with all rules, they're there to be broken and there are two approaches to improve the dynamic range of digital [HD] cameras.

The first technique uses the fact that the camera image sensor [CCD-Charge Coupled Device] is capable of sensing a wider dynamic range than the camera is nominally set to output, which is based on providing an immediately useable 'TV' gamma [contrast] linear image. Therefore, it is possible to adjust the camera gamma output so that a wider dynamic range [lower contrast] output is provided. This output is obviously not immediately broadcast useable but is a closer approximation to the capture methodology of film and provides additional headroom for post-production grading/timing.

With this approach the likely extension in dynamic range is approx. 2 to 3 stops and can be applied to most, if not all, digital HD cameras. [Refer to manufacturer information on the possibilities]

The second, and possibly more interesting approach, is for the manufacturer to provide an output from the camera in a form more closely matched to film. The first HD camera to provide this was Thomson's Viper, although more cameras from alternative manufacturers such as Sony & Dalsa are already appearing with such skills.

The Viper uses a CCD with a 12 bit front end [*CCDs actually capture analogue information and use an A-to-D to output digital information*] and this linear data is immediately translated into 10 bit Log data and output via Dual HD-SDI bnc cables to an uncompressed disc recorder in RGB colour space. The available dynamic range is in the region of 8 to 9 stops, far better than traditional 'linear' HD cameras, but a bit lower than film negative.

As an example of the future direction of digital camera systems the Viper shows serious promise.

Film Dynamic Range



Viper Dynamic Range



If you compare the above two images the Viper DR is considerably closer to film than the HD Video image above. With grading the final image can be set to look as desired, while the video dynamic range image is restricted in available headroom.

Note: It actually takes 13 bits of Linear digital resolution to accurately represent a negative film's dynamic range of 2.048 with a 0.002 granularity as specified by the Kodak/Cineon digital log file format. But then, film needs a great dynamic range as you can never be sure what the film has captured until after processing.

The benefits

With a large dynamic range the grading process can shift colour, contrast, gamma and brightness without fear of producing clipping or crushing. In the chemical film domain this is assisted by [print] film's natural transfer characteristics that produce a 'soft-start' and 'roll-off' to density information - print film's D-log-E curve.

Video cameras use CCD technology that is inherently linear in its response, causing the familiar white clipping and black crushing associated with video images, hence the desirability to work in log space in post production to minimise such crushing & clipping, enabling 'soft-start' and 'roll-off' to be added as part of the digital intermediate process; either electronically for final output back to a digital format or as part of the process of being written back to celluloid film. As the use of wide dynamic range cameras, such as the Thomson Viper or Dalsa data cameras, become more common the whole idea of DI will expand – especially with the take-up of digital projection systems too.

Compression

A second and equally worrisome trait of standard HD cameras is that they record image information via a 4:2:2 YUV data stream onto compressed tape based recording mediums. If the 'camcorder' mode is used with such cameras [the self contained tape cassette recording capability] the level of compression is exceptionally high and often reduced to 8 bit information and possibly 3:1:1 colour

detail as well as reduced resolution [pixels per line]. The alternative is to record on to a 'studio' vtr tape deck where less compression is used and 10 bit 4:2:2 information is recorded.

Different HD camera systems use different levels of compression and bit depths so it is advisable to run quality evaluation checks under likely shooting conditions prior to any project.

Obviously the output from such cameras can be recorded directly to a disc recorder [as with the Viper camera] without the application of any compression but as with the use of a studio vtr the camera will be tethered to the recording medium via a bnc cable. This can reduce the benefit of using digital capture in the first place. *At least a film camera always carries its recording medium with it and never uses compression – unless the grain size can be described as a compression format...?*

Newer HD cameras and vtrs are offering RGB 4:4:4 capture and reduced levels of compression such as Sony's latest CineAlta F950 camera and SRW tape deck, and show a real step forward for digital capture technology, along with the Viper & Dalsa cameras.

YUV vs RGB & Compression

It is worth stating that although compression is a problem for some film work, especially where vfx shots are involved, such as with keying and extensive grading, the use of YUV vs. RGB is a lot harder to differentiate, even if YUV is used 4:2:2 rather than 4:4:4. The reasons are complex but suffice to say the human eye is more attuned to luminance than colour.

Note: YUV colour space is actually larger [wider colour gamut] than RGB as it can display colours that are illegal in RGB and can therefore represent more total source colour range. Not always a good thing but don't be fooled by people who suggest RGB is better by default – it's not true! Although YUV colour space does have more redundant locations within it, so is not as efficient at holding colour information...

Compression is a problem as the coding used introduces errors [or artefacts] that tend to be different frame to frame, even if the image content is superficially identical, especially around edge detail. This causes 'noise' that becomes visible when level dependent processes such as blue/green screen keying & selective or secondary colour correction are used.

However, for more basic story telling where post work is limited to basic grading or timing, compression need not be a problem & often the benefits of lightweight electronic capture direct to an in-camera tape for dynamic story telling can outweigh the possible problems of compression. Creativity is always of prime importance.

Field of Focus

HD cameras use arrays of image sensors [CCDs] that are smaller in area than a 35mm film frame. This results in a larger field of focus, depth of focus or depth of field compared to an equivalent 35mm camera [very similar to 16mm?]. As focus is used extensively during filming to help tell the story, by leading the eye to the point of interest, having a larger field of focus can detract from the perceived quality of the final image. Therefore it is common to shoot HD with as open an aperture as possible, using ND [neutral density] filters on the lens to 'stop down' the camera, maintaining as 'filmic' an image as possible.

Using lower F or T stops [aperture wide open] also causes the focus to be slightly 'soft' giving a more filmic image.

Running low F or T stops also means there is increased danger of highlight blow-out. As stated before, use ND filters & take great care with scene lighting.

Back focus is also something that needs to be checked as this can 'move' during camera shipment and use, especially with high temperature changes, with an obviously detrimental effect on image focus.

Image Framing

One of the additional benefits of a film cine camera is its use of an optical viewfinder. Most [all?] HD cameras use electronic viewfinders and as a result only show the image area they are actually capturing. It is therefore imperative that boom mic's and the like are kept out of frame. It also means that the DoP cannot see 'around frame' and may miss the best possible camera framing for action as it enters frame...

Shooting Digital

When shooting HD for a DI project there are a number of guide suggestions that will ensure the best possible end result. As with any image based creative process the quality of the initial capture defines the maximum final quality attainable. Care and time taken during capture will save time and cost in post-production.

Garbage-in, Garbage-out.

Quality comes from the beginning and shooting on the highest quality format available ensures any compromises made are minimal. Obviously creative considerations come first, and if a production style demands grainy, less stable imagery to put across the emotion of the story then that consideration outweighs technical quality. But creativity aside, there are some simple rules-of-thumb to follow for best quality.

- For effects based projects, especially those using invisible effects to enhance a story line, such as period drama, compression, especially as part of HD shooting, is to be avoided. The compression process introduces low-level artefacts that while not obvious to the naked eye under normal viewing conditions become obvious when performing effects work such as chroma keying and possibly during grading if large changes are involved.
- Full bandwidth RGB/YUV image information [4:4:4] is preferable to half chroma bandwidth YUV [4:2:2 or even worse] for similar reasons given above, although less of a problem in the real world.
- The availability of a large dynamic range is always desirable to maximise the potential of the grading process. And as the grading process is at the heart of any film project, chemical or digital, its importance is difficult to overstate. Film has the widest dynamic range available [a possible 11 to 12 stops] with new generation digital cameras, such as Thomson Viper & Dalsa data cameras, having an equally impressive range [a nominal 8 to 9 stops for the Viper & 11 stops for Dalsa] with standard video HD cameras being WYSIWYG to standard tv monitor viewable dynamic ranges [a nominal 5 stops] - unless you set them differently - and you should!
- If possible, use the 'Cine Gamma' modes that HD cameras are beginning to come equipped with as a pre-set setting. These settings map the image data to a film OCN characteristic, to give a low contrast, high dynamic range image.
- Expose as if using Reversal film stock, taking special care with highlights.
- Even when using standard 'linear' HD or telecined data it is preferable to work in log colour space during post production.
- Shoot 24fps (23.98fps). One of the major reasons for this exists within PAL countries where 25fps is the video standard. In such countries it is common to perform 'tape-to-film' where 25fps video is recorded to film and hence plays at a speed decrease of 4.1%. The big problem with this is changing the speed of the audio with the necessary re-pitch. As the audio now has to run slower new audio samples must be inserted into the audio stream, with the result that phasing errors will almost definitely become apparent between stereo channels. Starting at 24fps and converting to 25fps for PAL video deliverables, where audio samples are removed, is a lot easier.

Digital Cinematography Rough Shooting Guide

The following guide points are just that - rough guides not god given rules. They will result in an image suitable for DI manipulation on iQ and will generate an image that will be low contrast and 'grey' when viewed on/under normal HD monitor/conditions. The biggest worry in following these guide points will be the introduction of digital noise [video grain?] if exposure levels drop too low. As with film capture, lighting plays a vitally important role both in quality of the final image and setting the correct creative mood and need as much care as with film capture.

- Check back focus – and re-check it regularly!
- Shoot 24psf - maintain a 24 fps filmic image without interlace artefacts - also vital for audio phasing for 24fps film/25fps video masters.
- Use 1/48 shutter speed - mimic nominal 'filmic' motion blur
- Use ND filters rather than stopping down - maintain a filmic depth-of-field
- Expose for highlight detail - treat HD capture as a Reversal film stock
- Take care with scene lighting - avoid highlight blow-outs
- Set camera for low contrast, extended dynamic range capture - mimic film DR (*see camera manufacturer's instructions for set-up – if available use the cameras Cine Gamma mode.*)
- Avoid black crushing - mimic film toe characteristics
- Avoid white clipping - mimic film shoulder characteristics
- Aim to grade or time in post-production not in-camera - retain flexibility
- Watch scene framing - there is no safe area or over scan area in an HD image

If followed, the result will be the most flexible use of HD capture for an iQ based DI post-production project, without the restrictions usually associated with HD vs. film.

DI Workstations

The central DI workstation where editing, colour grading, vfx, etc. is to be performed can both be easier and more difficult to deal with.

Easier, as if performing no grading function the workstation must simply pass data through itself with no change in the image [i.e. work as a transparent 'unity' device]. This is always a good first check for any system purporting to offer DI workstation capabilities - is it transparent? It's amazing how many workstations and grading systems are not transparent when set to default 'zero' settings.

Difficult, because once grading begins the workstation's viewing environment must display an image identical to the final projected film image. This is the hardest part of any DI system calibration procedure but is also where top-end DI environments separate themselves from the run of the mill. It is also closely associated with chemical lab calibration [and hence film recorder] as discussed in the following sections.

As a general rule it is easier to calibrate a digital projector than a monitor as a projector has a wider colour gamut compared to monitor phosphors and LCD displays as it uses a white light source and optical filters, very similar to a film projector with print film.

Monitor Choice


The choice of monitor is equally important, and flat screen LCD, TFT, etc. monitors are improving in quality and dynamic range and are starting to compete with CRT monitors as the primary choice for a grading monitor.

However, although computer flat screen monitors are RGB input devices they are manufactured to lower quality levels than grade one broadcast monitors. While the fluctuation in colour monitor to monitor can be overcome as part of the calibration procedure a lack of set-up stability cannot. More importantly any monitor's ability to track black levels throughout varying dynamic range inputs is vitally important for grading work to maintain an accurate presentation of shadow detail. This is an area where high quality broadcast monitors cannot be bettered.

Therefore, for accurate colour calibration, stability and contrast tracking the best monitors to use are those that are the norm in any top-end telecine suite - grade one broadcast monitors. Do not be lead down the 'computer monitors are RGB and therefore better' path. You will live to regret that.

It is also important to understand that the colour gamut of a monitor [the ability to display a range of colours] is different from that of film and therefore there are colours that film can display that a monitor can't and colours a monitor can display that film cannot.

This is shown in the following [rough] CIE chromaticity diagram.



However, TFT/LCD monitors do tend have better colourimetry than traditional CRT monitors (on which the above CIE diagram is based) as do digital projectors, so can be matched more accurately to projected film [see later].

Preparation

It is important to understand what the intention is when setting up a monitor for digital film grading. Although a video device it must be approached with a film mindset.

The basic requirement is to ensure that what is seen on the monitor matches, as accurately as possible, what will later be seen as a projected film print.

This has a number of requirements before we get anywhere near output LUT calibration.

1. The use of known reference images in both digital & film print forms.
2. Access to a film projector.
3. A correct colour temperature light box [if possible with variable illumination capability - or a selection of ND filters to match the projector's light output].
4. Someone with 'good' eyes [i.e. a colourist].
5. A Grade One HD monitor [used in preference to RGB computer monitors as it is a known quantity & of high quality].
6. Very dark viewing conditions with limited light contamination. [This is to match a film theatre's viewing conditions].

Reference Material

Using digital reference frames [which must include a good grey scale, colour swatches, flesh tone & a LAD patch] output the raw data files to film via the usual [well calibrated] film recorder & process through the chosen [well calibrated again!] film lab & print.

Measure the Status M negative density of the LAD AIM patch to prove the accuracy of the film output device/neg. bath and Status A density of the print LAD patch, to prove print bath accuracy, and ensure they are accurate to Kodak specified tolerances. Nothing further can be done until an accurate print exists!

Digital LAD Negative AIM Density - 445 patch

Film Printing: LAD 445

Status M:	R	G	B
OCN (5245) above D-Min	0.67	0.72	0.69
D-Min av.	0.21	0.60	0.98
Total:	0.88	1.32	1.67
IN (5242) above D-Min	0.87	0.93	0.91
IP (5242) above D-Min	1.02	1.09	1.08
D-Min av.	0.07	0.57	0.67
IN Total:	0.94	1.50	1.58
IP Total:	1.09	1.66	1.75

Av. Film Recorder Printer Lights:

OCN	25	25	25
IN	30	37	25

Status A:

Vision Print	R	G	B
	1.09	1.06	1.03

The Grading Environment

Careful consideration must be given to the illumination conditions within any room to be used for film grading. Normal telecine rooms use lighting conditions more appropriate for projects to be seen under general living condition in normal homes. Film projection occurs in more controlled environments within a theatre and the ambient illumination is many times darker than home viewing of TV. For this reason the film grading environment must be similar. Any light contamination will reduce the contrast range & shadow detail of the grading monitor resulting in inaccurate grades being performed. If in doubt shut out all light except for the usual green 'Emergency Exit' that is always to the lower left of any cinema screen!

The Monitor

The main grading monitor needs to be a Grade One HD monitor for reasons outlined above & the Cine-tal series meets the requirements well.

The monitor should be set-up as per manufacturers instructions as a starting point, making sure this is performed accurately with careful attention to black detail with a Pluge test generator in a correctly set-up grading room as defined above.

The monitor colour temperature should be set to 5500K (D55) to better match film projector colour temperature, which is 5400K (-100/+300), hence the use of 5500K for the monitor *[note: if possible one alternate monitor input should be kept to normal HD standard D65 for video deliverable viewing]*.

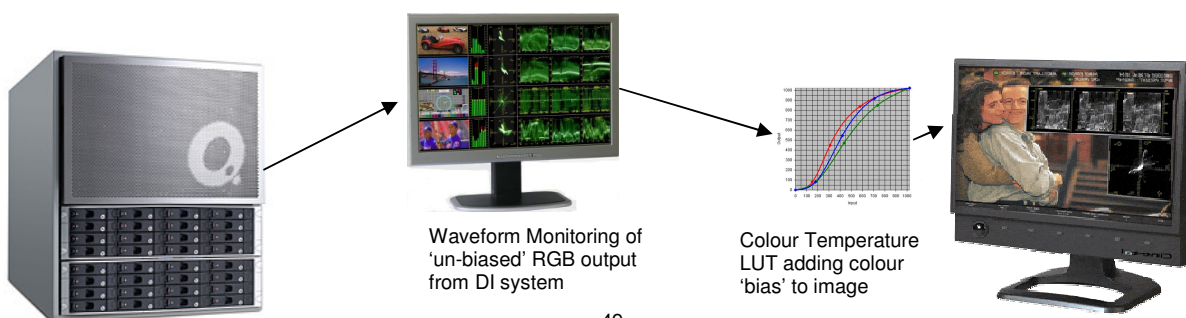
As standard monitor defaults are D65 & D95 this 'film calibration' may need to be set manually using the following coordinates and a monitor probe.

Colour Temp. Coordinates

Temp	X coordinate	Y coordinate
D50	x0.3457	y0.3585
D55	x0.3324	y0.3474
D65	x0.3127	y0.3290
D95	x0.2848	y0.2932

This is done to enable the output to be colour balanced for 'film' viewing without the need to bias the output via the output Lut printer light controls. This enables an attached HD waveform monitor to correctly show grading levels.

With newer monitors becoming available, View LUTs can often be used to directly set-up the monitor, at the monitoring point. It is important that this is done at the monitor, not at the output of the DI system as this allows for accurate waveform monitoring of an unbiased image.



Contrast, Brightness & Chroma

With the correct colour temperature set the next step is to match black and white points via a projected or light box displayed printed grey scale & saturation via selected printed colour frames.

Note: To use the following calibration procedure requires the use of a LOG to Print basic 2D LUT, which can be generated via the Digital Praxis View LUT Builder.

Standard film projector output levels are 16ft lamberts for an open gate [no film in the gate], as measured at the screen. With base film [d-min] in the gate the output drops to 12 to 14ft lamberts. Therefore the light box should be matched to 16ft lamberts with no film present and the grading monitor to 12 to 14ft lamberts with digital peak white [1023].

With the same grey scale film frame on a previously checked light box [checked via a direct comparison with identical film in a projection theatre & matched to colour temperature & light output] and a Log to Print View LUT used to display the digital Log image, the monitor output can be compared and brightness adjusted to set blacks & contrast to set whites. Note that if the previous 'video' default set-up has been performed correctly blacks [brightness] should need little adjustment. If you are making major changes go back and re-check the video Pluge signal alignment. White [contrast] can also be checked by measuring the ft lamberts output when displaying a full white frame, as described above.

The contrast setting should also be checked with real images, such as Marcie, as grey scale ramps & steps can be difficult to gauge. Note that it is the main white point that is being set as contrast [gamma] is set via the DI system's output LUT [see later].

Note also that if any ND filters are being used [see below] to limit the light output from the light box black detail may appear more crushed. This needs to be taken into account during monitor calibration.

Using a frame with good colour detail [Macbeth is good] chroma can also be set to give a similar level of saturation, normally requiring the monitor output to be de-saturated from its default.

It will be necessary to use a selection of different frames and repeat the calibration procedure to gain the best overall result.

At this point the monitor set-up is complete & displays a Log image Via a Log to Print LUT, with colour temp set to 5500K, with the correct contrast, black & white point and average colour saturation. It is now down to an accurately generated & calibrated full display LUT to finish the calibration.

Colour Calibration LUTs

With the monitor set-up correctly a fully calibrated LUT is now required for accurate 'print' calibration. This can be performed as empirical process, generating a LUT via a 2D LUT Builder, such as that available from Digital Praxis, or via a specific 3D calibration program such as Kodak's KDM (Kodak Display Manager), or FilmLight's TrueLight.

Either way, the process required is very similar, with the goal being to make the displayed digital image match the final film projected image.

However, things are not as 'ideal' as one would like, with the concept of an 'accurate film print' being one of the major issues. See later.

Monitor Calibration Procedure Guide

- Obtain a perfect 'Status A' print of calibration reference frames. These MUST be as accurate as possible.
[Grey scale, LAD, Macbeth, Marcie, etc.]
- If using light box calibrate to film projector standards using optical filters as necessary.
[Colour temp to 5400K & Illumination to 16ft/Lamberts]
- Set grading room light levels to match film theatre.
[Darker than telecine room!]
- Calibrate monitor to default TV specifications as per monitor instructions & probe.
[Check accurate black setup via 'pluge' signal]
- Calibrate monitor to film projector colour temp & illumination levels.
[Use probe to set colour temp to D54 (D55) & peak white illumination to 12 to 14ft/Lamberts - re-check black setup]
- Display ref. grey scale image from DI system & turn on Log to Print View LUT.
- Place 'perfect' Status A grey scale print frame on light box as reference.
- Check 'contrast' for accurate black & white point on monitor image compared to print frame.
- Display ref. Macbeth image from DI system & place 'perfect' Status A Macbeth print frame on light box as reference.
- Set monitor 'chroma' control to match average colour saturation of ref. print frame.
- Check & compare black & white levels and adjust monitor 'brightness & contrast' controls as necessary.
[Note: Adjustment of monitor 'brightness & contrast' should be minimal at this stage.
Large change requirements show errors in earlier set-up stages]
- Display ref. Marcie image from DI system & place 'perfect' Status A Marcie print frame on light box as reference.
- Re-check settings.
- Display ref. LAD image from DI system & place 'perfect' Status A LAD print frame on light box as reference.
- Re-check settings.
- Repeat last stages with various calibration reference images to verify monitor calibration.
- Generate Colour Calibrated LUT using specific LUT Builder instructions.
- **MAKE SURE ALL FILM PRINTS ARE PRINTED ACCURATELY TO AIM.**
[R.1.09, G1.06, B1.03]

Repeat calibration procedure on regular basis to maintain accurate calibration.

Digital Projectors

As mentioned previously, a digital projector can be far better colour gamut matched to the final film print image. This is because the light source is a wide spectrum white light, matching that of a film projector, and the RGB light paths are controlled via optical filters that are/can be matched the print film dyes. The use of digital lookup tables [LUTs] can further be used to match film's cross-colour contamination [the fact that within film there is no such thing as pure red, green or blue - or pure cyan, magenta or yellow for that matter].

As a result a digitally projected image can better match a final film result compared to any presently available monitor. Additionally, grading on the big screen has direct benefit as the image is being seen in it's final form - large - which also helps assess editing as on the small screen [tv] fast edits are a lot more acceptable than when viewed large. Grading via a digital projector helps bring editing and colour together for a better final result.

Note: Although a digital projector can better match the final 'film' image, as projected by traditional print film, that is not to say or suggest that film has a better or preferable colour gamut than a monitor or TV. Just that the two are different. If you refer to the chromaticity diagram previously you will see that neither colour gamut overlaps all of the other. In fact some of the more serious major digital film companies would like to see a projection film stock with a colour gamut matched to monitor phosphor. That would make all image formats immediately image compatible!

Also, any projector is only ever going to be as good as the image it is presented with, and there have been a lot of poor projections due to inferior source material. Sending a projector a 'TV' image [linear, with crushed blacks and clipped whites] will always result in the projector looking poor. The image needs to be presented in a 'filmic' form, which means as log source material via an accurately calibrated digital colour look up table [LUT].

Projector Calibration

The procedure to present a digital projector with a high quality image is similar to that outlined previously for monitors. The major difference is that theoretically for print film matched projections the basic colour gamut matching can be performed via the projector, probably using optical filters to match the film dyes. There has also been a lot of digital colour manipulation tables created for digital projectors with P7v2 being one of the more widely known.

Experience with digital projectors being matched to film projectors [DLP type digital projector from Barco & Christie, film projector from Kiniton] has shown that superb results can be obtained with minimal colourimetry distortion via LUTs.

Obviously, as with monitor calibration, the use of accurate 3D LUTs, as generated via the likes of Kodak's KDM (Kodak Display Manager) or FilmLight's TrueLight, will generate the most accurate calibration.

Film Recorders

As DI systems expect digital film image data as Cineon density mapped DPX information [assuming the project is to be worked in Log image colour space] it can be viewed as a unity process, adding only its creative processes to the raw image data. i.e. Log-in, Log-out.

With monitoring set-up & calibrated as described above, any output will be matched to the industry .DPX or .CIN log data standard & therefore a known & understood entity. A well calibrated film recorder will therefore output a known 'digital film negative' matching the expected intermediate Status M density for a digital 445 LAD patch as specified by Kodak and assuming a well balanced and maintained film processing lab.

From this negative an accurate print can then be struck, matching the Status A density values for the LAD patch.

In essence, the DI film out process aims to produce a neutral negative equivalent to a perfect OCN. Obviously, as the film recorder will output to intermediate stock [5242] the actual printer lights used will be different from this, but the theory still holds true.

As with film scanners, film recorders [Lasergraphics, ArriLaser, Celco, etc.] are also factory calibrated to the .CIN/.DPX format. However, this calibration is based on Kodak's perfect 'film lab' processing [as is scanner & workstation calibration] and therefore will only be as accurate as the film lab itself. The difference between scanner calibration and workstation monitor & recorder calibration, as previously mentioned, is grading occurs after scanning and is therefore less critical, while what is seen on the grading workstation is required to be accurately reproduced on the final film print.

From this we can see that the greatest variable within the DI process is the film lab, and the requirement to match the calibration of the workstation monitor and film recorder to the specific lab to be used for final film processing.

Film Recoding Process

Once all DI work has been completed the digital film frames will be sent to a digital film recorder. Each frame will be exposed in turn, building a new digitally mastered intermediate negative [or intermediate positive, although this is less common but may have the benefit of reducing the number of dupe processes undergone when striking the final print deliverable - at least until the speed of film recorders enables multiple internegative to be produced cost effectively for the direct production of projection prints].

DPX or CIN?

One question when outputting digital film frames for film recording is the format to use.

Although DPX tends to be the standard for DI post-production there is a major danger with this format when recording out to film. The DPX format can be Linear or Log, specified by a 'flag' within the header of the image file. However, most film recorders don't read this flag, and assume DPX to be linear! This can cause major issues with the film-out process.

If working Log, as one should, I recommend outputting Cineon .CIN files as these can only be Log, not Linear.

Speed of film recording varies dependent on film recorder make & model as well as data format and associated processing such as sharpening. Expect 1 to 4 seconds per frame.

Controlling the Recording Process

As a guide to control of the film recording and processing stage, a selection of reference frames should be exposed at the head & tail of each and every output. These images should include the

digital reference frames used to calibrate the DI grading monitor [see monitor set-up above] and any reference frames required by the film recorder used. This may include a 21 step grey scale densitometric strip, but as a minimum must include at least one 445 and 470 LAD to enable Status M measurement of the negative [intermediate] and Status A measurement of the print to be made.

Status A checks should be made using the reference frames at each end of the output print (head/tail) to ensure accurate & stable film recording. If in any doubt measure the negative too. If the film recorder is not recording stable data the whole DI process is useless.

The main cause of film recorder instability tends to be temperature drift within the recorder's environment. And a density change of only 0.02 in status M (negative), which is very small in density terms, equates to 0.07 in status A (print) and is one printer light, which is very visible. And it tends to be 'laser' film recorders, such as the ArriLaser, that are most susceptible to density fluctuation due to temperature variations.

For a good 'digital output negative' the Status M AIM values should be:

LAD Intermediate AIM Density - 445 patch

(Kodak values – other film stock will have different aim values)

Status M:	R	G	B
IN (5242) above D-Min	0.87	0.93	0.91
IP (5242) above D-Min	1.02	1.09	1.08
D-Min av.	0.07	0.57	0.67
IN Total:	0.94	1.50	1.58
IP Total:	1.09	1.66	1.75

With these values attained a print struck from the negative should require the following Printer Lights:

LAD Based Printer Lights

(Kodak vision & premier print stock)

IN/IP Printer Lights	R-30	G-37	B-25
----------------------	------	------	------

These values are due to the use of Intermediate film stock with its salmon base colouring, and are directly relative to accurately exposed OCN's perfect 25 across [R-25 G-25 B-25] printer lights, and will produce a Print with the following Status A AIM values:

LAD Print AIM Density - 445 patch

Status A:

Vision Print	R-1.09	G-1.06	B-1.03
--------------	--------	--------	--------

If a film recorder is being used to record out to Original Camera Negative (OCN) the Status M AIM values should be:

LAD OCN AIM Density – 445 patch

Status M:	R	G	B
OCN (5245) above D-Min	0.67	0.72	0.69
D-Min av.	0.21	0.60	0.98
Total:	0.88	1.32	1.67

LAD Based Printer Lights

(Kodak vision & premier print stock)

OCN Printer Lights	R-25	G-25	B-25
--------------------	------	------	------

The Film Lab

It is worth saying a word or two about the chemical lab process as having a basic understanding of the way a lab works will aid in obtaining the correct final result.

A original camera negative records a scene as density within the film stock's emulsion. After processing this results in the areas of the film that were exposed to bright information in the original scene being dense [dark] within the film frame and areas of the frame that were exposed to dark information being less dense [bright]. The film also records the inverse of the original colour. Hence, the initial film recording of the image is a negative copy of the original scene.

This is also true of the internegative produced as the output from the DI process. The positive image seen on the grading monitor is recorded [usually!] as a negative image on film.

When this negative film image is copied [duped] light is shone through the processed negative onto a new, unexposed, piece of film stock. This stock is similar to the original camera negative in that it records brightness & colour information in the same way [the more light the denser the film becomes] but as it is 'seeing' a negative image from the original camera negative it records a positive representation of the original scene.

Printing a projection print from a negative [OCN or dupe] follows the same procedure.

Note: intermediate stock has a longer 'linear' [less toe and shoulder region] response smaller grain structure [it's a very slow stock as a lot of light can be used in the dupe process] and has a gamma of 1 so that it distorts as little as possible the characteristic of the OCN. For a better understanding of the film chemical process [without getting too technical] review the Quantel Digital Fact Book - film section at the rear.

The level of exposure 'seen' by the new film stock is controlled by varying the amount of red, green and blue light shone through the original processed negative. This is called 'timing' and uses values called 'printer lights' to define the amount of light being used. This timing process can be used when making a dupe [intermediate positive or negative] or when making a final projection print.

For a well exposed original camera negative printer light values of R25 G25 B25 produce a balanced and mid level exposed print, often described as 25 across [duh!].

In the DI world a 'perfect' digital neg. will print the equivalent of 25 across with printer light values of R30 G37 B25.

Matching the Film Lab

As the film lab has the final say in the colourimetry of the output from the DI lab it is obviously the reference to which the digital image must be balanced. If the film lab is 100% calibrated to Kodak aims then calibrating the digital DI lab to .CIN specifications will result in a 100% accurate DI process. Therefore, ideally the lab should be perfect, allowing the DI environment [scanner, workstation monitoring & recorder] to be calibrated throughout to default .CIN specifications.

However, labs are very rarely 100% calibrated due to the nature of the film processing procedure. And what if the client requires a film stock other than Kodak to be used? And additionally there are two lab processes that need to be accounted for - negative and positive process baths.

Because of these variables adjustment of either the workstation monitors/projector or film recorder, or both, is often required to counter lab/film stock differences.

Ideally, the film recorder & negative film bath should be calibrated to output a perfect negative [internegative or interpositive] after processing. From this a positive print can be struck. The best method to calibrate the film recorder & negative film bath combination is to output via the recorder a 'sensimetric' film strip with a range of grey values throughout the full density range.

Film labs are calibrated by Kodak using 21 step densitometric strips, which is not really enough for accurate calibration. Therefore, generating a 128 step strip [a single grey frame output for every 8th digital sample] will provide more than enough accuracy for recorder calibration. The procedure is to output the strip, process the negative and then measure each step [frame] with a Status M densitometer. The read values should be compared to Kodak's .CIN aims and the film recorder output Luts adjusted to bring inaccuracies back into line. If this process is carried out multiple times per day excellent accuracy can be attained. As a minimum it should be done after a new negative stock is placed in the recorder [new batch number], different stock type [Kodak, Fuji, etc.] or when the lab changes the composition of the bath.

One of the benefits of calibrating the film recorder and negative processing bath to output 'perfect' negatives is that the workstation monitor/projector can be calibrated to 'default perfection' by matching to perfect Status A reference print frames [i.e. perfect .CIN calibration]. This enables such a film recorder/negative bath duo to output any .CIN/.DPX image and generate an almost perfect internegative, from which a near perfect print can be produced, assuming the workstation used to grade the material was accurately calibrated for .CIN/.DPX image files.

With the recorder & negative bath calibrated the final calibration process is the positive bath & workstation monitor. This can be performed either empirically or scientifically - but both methods require an on-aim print, i.e. the most accurate print possible.

What is different about the print process bath is that it can be/is controlled via timers/graders adjusting printer lights. This allows some correction for bath inaccuracies so it is theoretically possible to 'always' gain an accurate Status A print from an accurate Status M negative. Also, unlike the negative bath/recorder calibration, print calibration requires 'real images' with colour and detail.

Printer Lights

As printer lights are used to expose a dupe film copy or print they can control the relative density of the new film by varying the amount of light that is made available.

For example, increasing the amount of source printer light will make the copy denser overall - darker to look at assuming a print [positive] is being made but effectively lighter if a negative is being copied from a positive dupe - confused? Let's see if we can explain better.

If a negative is made very dense [over exposed] when capturing the original scene, when printed it will produce a light print for a fixed printer light value as the dense neg. will limit the exposure of the print [under exposing it] making the print mimic the relative brightness the negative 'saw', i.e. bright! Therefore, to make the print denser and hence darker requires more light to be shone through the dense negative to further expose the print stock. Better now?

Printer Lights & DI

For DI work we tend to assume the new digital negative output via the film recorder is a fixed entity as little can be done to alter it other than re-outputting the material again and re-processing. [Obviously, this assumes a well calibrated film recorder]. Therefore when discussing printer lights within the DI world it is usually in the context of striking a projection print from the output digital film negative - the first answer print and proof of the DI process.

Obviously the DI process and the processing lab are striving to produce a perfect final print so that the film when projected matches perfectly the image(s) seen on the DI grading monitor [see above]. However, as it is a chemical process the lab cannot realistically be expected to attain perfection and it is usual to expect a tolerance of +/- 1 printer light per colour per lab process. This may not sound much but a variation of one printer light is a difference that can easily be seen, especially when a print is reviewed immediately after viewing the 'perfect' DI image on the LUT calibrated DI grading monitor.

Therefore the print's Status A densities must be checked prior to any screening and an assessment made of the colour/brightness variation to be expected. If the variation is great a new print should be requested, with the printer light values adjusted to counter the error.

Additionally, if the error is slight - within acceptable tolerances - the Status A density measurements can be used to alter output calibration LUTs for the grading monitor to mimic the print error. This allows the DI data to be reviewed as a match to the print, with the same error introduced. This is a very powerful and accurate way to prove the DI process allowing for the lab variations.

LAD Print AIM Density - 445 patch

(Kodak values – other film stock will have different aim values)

- Status M (above D-Min) - R-0.87, G-0.93, B-0.91
[5242 intermediate IN stock]
- Status M (above D-Min) - R-1.02, G-1.09, B-1.08
[5242 intermediate IP stock]
- Status A - R-1.09, G-1.06, B-1.03
[vision print stock]

One printer light is [roughly] equal to 0.07 density as measured by a densitometer [Status A]. *[7 printer lights also equal one stop approx]*. Therefore, if the Status A density reading of the 445 LAD patch is out by 0.07 the lab can adjust a reprint by one point [one printer light] to correct.

It also means that the calibration LUT can be adjusted to match the grading monitor to the print, balancing to the actual print values, rather than simply working with the ideal [perfect print] values used during the digital grading process.

Obviously if all printer lights are out by an equal value the overall colour balance will be maintained - the image will simply be lighter or darker depending on the direction of the error.

If one or two printer lights are out, or all lights are out in different directions, the print will have a colour bias as well as being potentially brighter or darker. This is explained in the following chart.

Printer Light Effect - on positive print

Printer Light	-ve	0	+ve
	Red	25	Cyan
Colour	Green	25	Magenta
	Blue	25	Yellow
Density	Lighter	Nominal	Darker

As can be seen above, increasing the printer lights will make the print more dense and more cyan, magenta or yellow. Decreasing the lights will make the print lighter and more red, green or blue.

Therefore, if the Status A measurement showed red nominal, green and blue both less dense, the print will be cyan and brighter than expected.

Note: the print will also appear less saturated and lower in contrast as a direct result of being brighter. This is more an optical trick of the eye than a physical reality.

Printer lights are also not direct acting controls - adjusting red for example will also affect, to a lesser extent, green and blue.

This is definitely more of an art than a science but understanding the basics will be of great help in gauging what to expect from the DI process.

Being Practical

However, in the real world it is understood that film labs the world over are all working to 'acceptable tolerances'. This is based on each film process bath being accurate to within +/- 1 printer light or 0.07

density of the final print. This is a tolerance that is visible to the human eye should various prints be compared one after another, but is a realistic understanding of the difficulties involved in keeping a chemical bath balanced. If you want proof of this watch a film in, say, Los Angeles, and then watch it again in London. Each print will have come from a different lab, although possibly from internegatives produced at the same lab at the same time, showing the variation possible from print baths alone.

Although difficult to see in printed form, if you are reviewing the document on a PC monitor you should be able to see a marked difference between the following three prints, all of which are accurate within Lab tolerances...



Because of this many realists in the DI world accept that no matter how tight they can hold their workstation monitor, film recorder and film lab calibrations, when it comes to the final deliverable as seen by Joe Public, good enough is good enough.

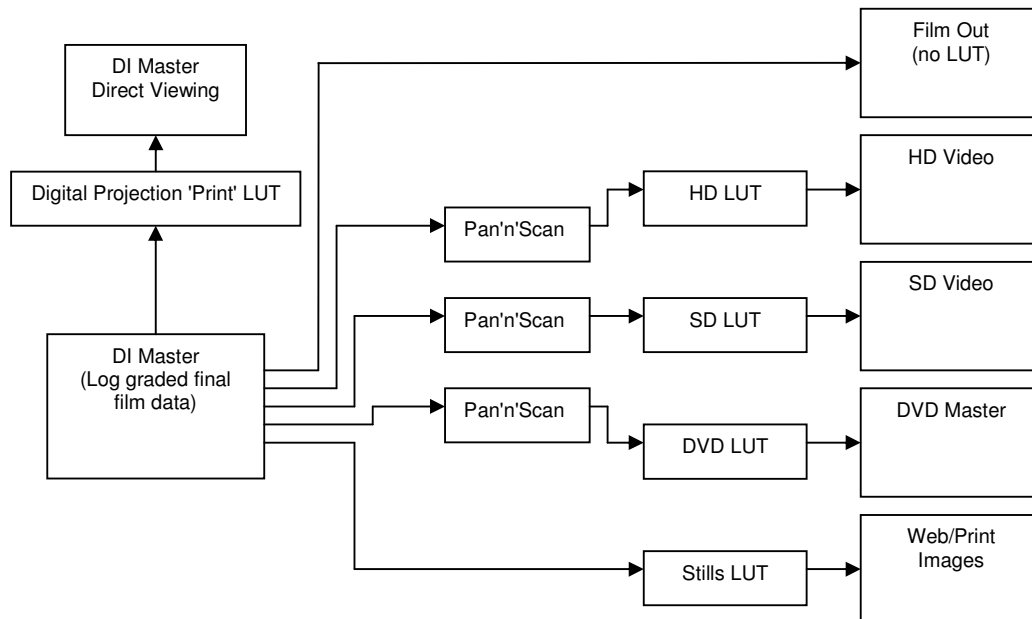
Therefore, a DI environment with each stage in the process calibrated to the .CIN specification combined with a film processing lab that is within Kodak tolerance [negative and positive] will often turnout results that are more than acceptable.

We should remember that all film shot in-camera and processed without the benefit of a DI environment is at the mercy of the film lab. And for the past 50 years plus that has been good enough. DI really shouldn't be considered any different from a business perspective.

Video Deliverables

Assuming a colour set-up calibration has been performed for film finals the monitoring environment needs to be readjusted for video deliverables. This will probably include colour gamut, colour temperature and contrast/brightness as well as viewing conditions - cinema tends to be a lot darker than home TV viewing conditions.

The DI master will then need to be globally re-corrected to match the original 'film' colourimetry. With practice it is possible to define a global LUT to perform this operation, but it is wise to never rely fully on such a one re-grade approach as the different viewing conditions of cinema vs. TV can often show unexpected differences that require more than a global change.



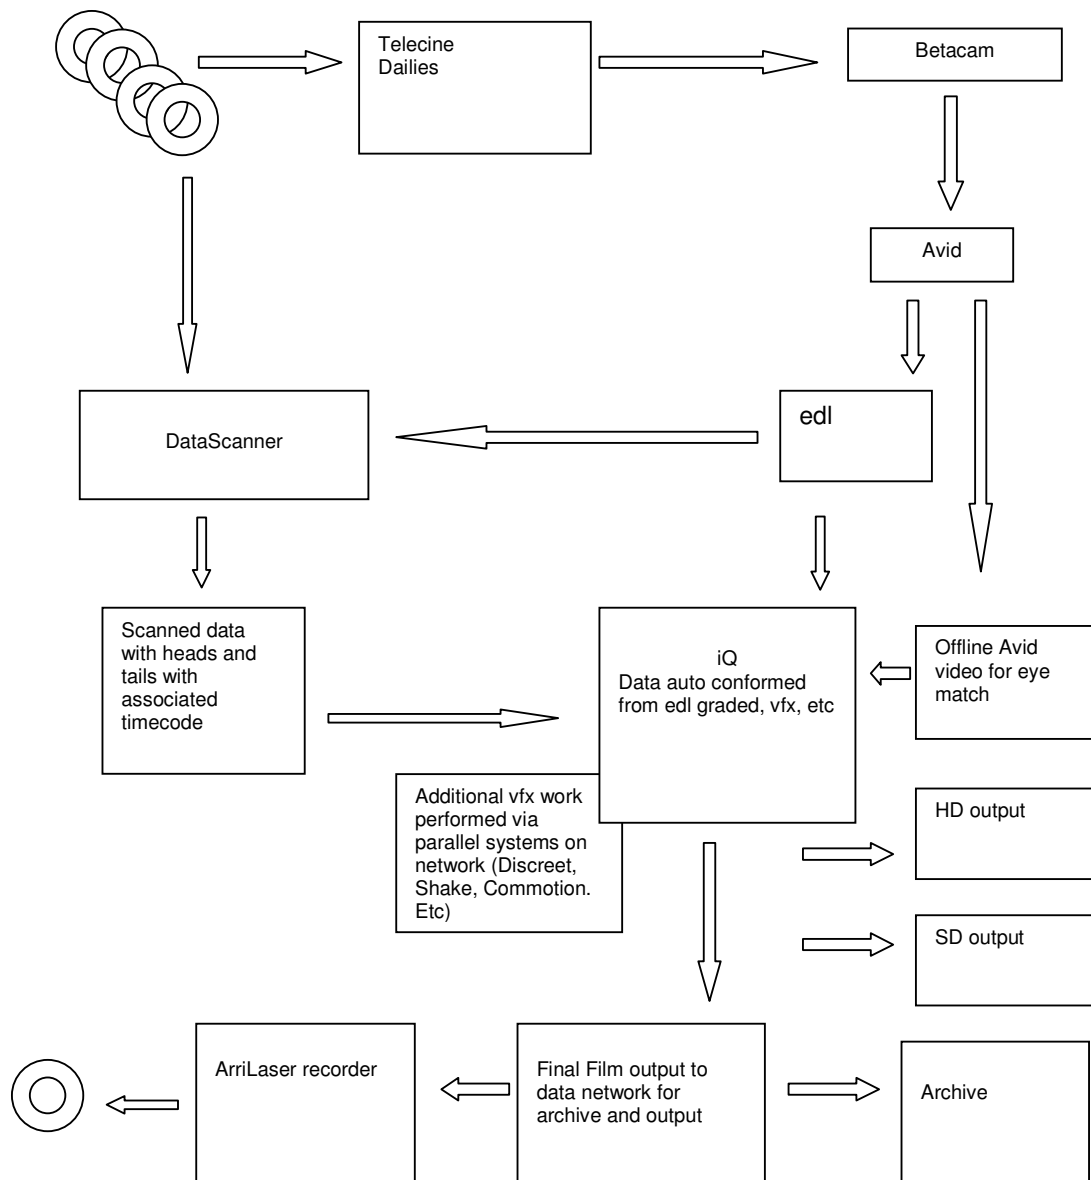
Film Based Workflow Descriptions

The following diagram shows one approach to DI work from film-sourced material where the original camera negative is never cut.

Original camera negative is telecined as 'dailies' transfers, usually as 2000ft reels of spliced camera rolls. Each roll is 'hole punched' to mark the start frame & the video transfer performed with timecode starting at 1:00 hour, 2 hours, etc for each reel. The transfer is performed 'frame-to-frame' which means running the telecine at 25fps for PAL output to ensure one full film frame is transferred to each video frame. For NTSC the telecine is run at 24fps with 3:2 applied to generate 30fps video. However, working at 30fps will cause problems later as the offline video copy will not be suitable for use as a split screen reference to the online unless the 3:2 is undone. This is also why for PAL operations the telecine is not run at 24fps as an extra field is introduced every 12 frames.

From this video transfer the offline is performed, generating an edl & low-resolution video cutting copy.

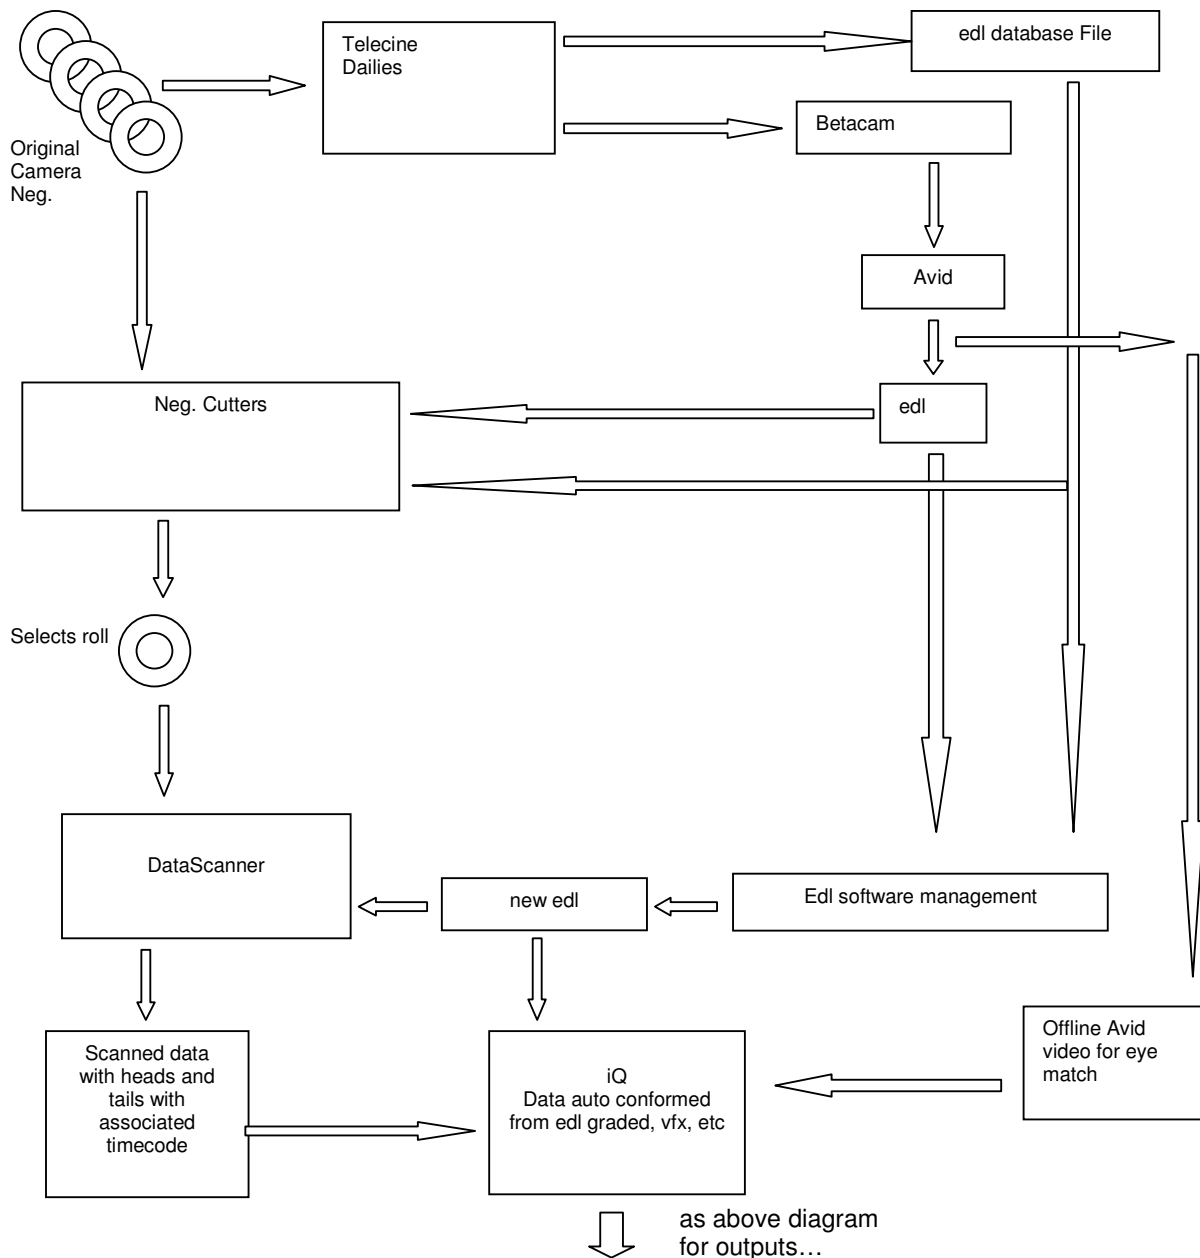
DI lab workflow (Uncut Neg.)



In this workflow model the selected takes, as defined by the offline edl are scanned at 2K from the original, uncut camera rolls. This enables the offline edl to directly relate to the 2K scan requirements, making the online scanning easy to complete and conform. It also means the original camera negative is not cut, minimising the chance of damage & dust pickup. The downside is that the selected takes must be extracted from the complete rolls, taking longer at the scanning stage. However, scanning can proceed before a full edit lockdown & less time and money is spent neg. cutting and the likelihood of film damage is lessened.

The other alternative is to produce a negative 'pull list' from the edl keycode information and neg.-cut the original camera negative into selects rolls. These can be the selected scenes from the offline edl, with handles (tails), or the whole selected take, flash to flash (clapper board to clapper board).

DI lab workflow (Neg. Cut)



It is worth noting that it is usually cheaper to neg.-cut just the scene plus handles, rather than flash to flash, but this can cause the image to 'jump' during 2K transfer if a data telecine is used without a pin registered gate. Although the splice is within the tails frames, as it moves over the rollers of the telecine, especially the capstan, it causes the length of the film run to alter, pulling on the frames presently being scanned in the gate, distorting them visibly. To avoid this phenomenon requires tails of about 3 feet or the use of a pin reg. gate.

Producing a selects roll(s) also means the original offline edl no longer references the source material correctly. It must be updated to reflect the new location of the scenes, and this should be done as part of the neg. cut or someone is going to have to manually track and translate the keycodes back to timecode!

Although there is more 2K transfer work involved in scanning [film lacing and shuttling specifically] the original, uncut, neg. rolls, there is a lot to recommend that approach.

Whatever method is used, when all the material is loaded into the DI system the edl is used to autoconformed each reel automatically.

Additionally, many DI systems are evolving the ability to work with film keycode, conforming the digital online to the film editor's cut list information.

From this point the online work commences, with edit alterations, vfx work, colour grading, until the final 2K master is complete, using all the DI tools available, as well as networking shots to additional systems within the DI environment. Care must be taken when doing this to ensure no colour, bit depth or resolution changes occur due to restrictions within the networked systems. Not all digital systems can work transparently with 10 bit Log data.

Once complete output can commence via a film recorder. Assuming the job has been performed 10 bit Log, via calibrated Log to print monitor display LUTs, there will need to be secondary colour control performed for any video (HD or SD) outputs.

Although one might assume that it would be a simple procedure to play-out the 2K master final via the Log to print Luts with HD or SD selected as the destination, this is not correct. All work performed for film output is done in viewing conditions best matched to cinema [very dark environment] while for video output light needs to be reintroduced into the room as for a normal telecine session. This changes drastically the perception of the image on the screen and requires a different grade style. It may be a simple global grade, but will probably require the odd shot to be lifted or have its contrast altered to take into account the changes in viewing conditions & preserve the feel of the 'film' original.

The Future for Dailies?

The above workflow descriptions assume present dailies operation with all OCN being telecined at video resolution for the offline selection process and the selected takes being scanned at 2K/4K for the online.

However, as storage capacity increases and costs decrease the goal is to be able to transfer all OCN at high resolution, requiring that the negative only ever be 'used' once and immediately archived. This protects the OCN from damage and reduces the chance of dust and dirt on the data transfer as the neg. is never handled for cutting.

Some operations are already using HD for this process but the use of HD tape for storage means the image information is compromised by the use of compression and a reduction in colour component.

This is the future for dailies though, with low or zero compression tape systems and cheap large storage discs becoming available. When this occurs the line between offline and online will blur and merge even further.

Data Wrangling

Data Wrangling is the fine art of moving stuff around - in the background. It enables a Data Monkey [someone who manages the location, movement and tracking of data - images in this case] to make sure the right images/audio are in the right place at the right time.

It's a very important part of DI as there will always need to be some data movement, even if it's just from scanner [data telecine] to the DI workstation and from the workstation to film recorder [an ideal scenario].

More likely it will involve moving individual shots to/from other, dedicated vfx systems (Discreet, Avid, Apple, Piranha, etc.) while the DI workstation continues with the major role of online editing and colour grading. Therefore, someone needs to know where the vfx shots are, what stage they are at, which is the latest version, etc.

It's a key part of DI and really needs to be a full background task, leaving the DI workstation operator to concentrate on what is important - working interactively and creatively with the client. As far as the DI operator is concerned image & audio material should be just where it needs to be, exactly when it's needed, without his/her involvement.

Offline and Timecode

There is a lot of confusion on the present offline/online process for film DI applications due to the differing frame rates of video and film. However, it is not as difficult as it first appears as described in the following text edited from an original document by Thomas Urbye of The Moving Picture Company... which needs more to work to make it easier to read!

The growing use of keycode lists rather than 'video' timecode eds is also helping further for film based shoots – obviously being of no benefit for Digital Cinematography work where keycodes don't exist!

Timecode conforming

Neg. rolls will usually be telecined each day by a local post-house or the film lab as dailies, ostensibly to ensure the film contains the expected images and no re-shoots are required. Best to know this before moving location or striking a set and it is easier to view a video rush than a film rush, although film prints are often also struck at this stage.

Each reel will be approx. 20mins duration, and will be made up of a number of spliced camera rolls. A hole is normally punched in the reel as a reference point, usually nominal frame one, for the following video transfer and is placed exactly on the 'hour' of the Beta tape that the reel is being telecined to, with each roll using a different hour as roll reference [as much as that is possible]. From there the keycode reader will keep a sync of what timecode frame = what keycode frame and this data is held within the keycode reader's database [a .fLEX file].

Because SD video has to work at 25fps [PAL] or 29.97fps [30]fps [NTSC], the telecine for PAL transfers has to run at a speed increase of 4.1% thus achieving a frame for frame transfer, and for NTSC the telecine runs at 23.98fps, and insert 3:2 to make up the missing frames create 29.97fps.

For PAL it's important to realise that running the telecine at 25fps makes one 24fps film frame equal one 25fps video frame with the result that the video version will run faster [have a shorter duration] than the film version but will have exactly the same number of frames.

The beta tape (which is purely a transfer medium to get material into the Avid) will have a log floppy disk (.fLEX file) supplied with it that will contain the keycode/timecode sync relationship (the database). This means that once the log file is loaded into the Avid (via ALE) it can maintain the timecode/keycode relationship internally (keeping them in phase) which will benefit the lab and VFX department etc.

PAL offline

An Avid film composer gives you two options when working PAL offline [from a 25fps 'frame to frame' transfer]. You can either edit at 24 or 25fps. Creatively speaking, editing at 24fps matches the play-out frame rate at the cinema [obviously!] but play-out to a monitor happens via the insertion of an extra field every 12 frames to make 25fps for play-out [SD monitors can't playback 24fps]. Working at 25fps means no fields have to be inserted but the project 'runs fast' and audio sync can become an issue as it is likely it was recorded on set at 24fps and will require a 4.1% speed increase.

Creatively speaking, working at 24p is always preferable as that is the playout speed at the cinema. With 24'p' the Avid combines both fields of a frame to create one progressive frame on the media drives (which is not a problem as the source material is film and thus progressive), this will let you edit at a frame play out rate of true 24 progressive, playout to a monitor can then happen at 24fps and durations will match those seen in the cinema (more on this later). When working in DI it is always advisable to work in a 24p Avid Project rather than 25fps project with matchback.

With a 24p project it is possible to work at a playout rate of 23.98, 24, 25 or 29.97 to your monitor (the Avid either inserts a field every 12th frame (24), increases playback by 4.1% (25) or inserts 3:2 (29.97)). With PAL this means that when the finished cut is laid back down to SD tape from the Avid at 25fps, the transfer is frame to frame which is good for the online, whereas if the tape playout is set to

24p then the Avid will insert a extra field on every 12th frame (so that a layout to tape is still possible and total duration matches that of the film rolls). This has the nasty side effect of producing a 'judder' on tape (noticeable on pans), but is used primarily for the post sound facility so 24fps sync is possible in the sound mix. If layout is to NTSC 29.97 (3:2 inserted), then the films duration will be correct and thus can be used for the sound mix, but when brought into a DI system such as Quantel's iQ the 3:2 will need to be removed so that unwanted fields are removed and only progressive frames remain.

NTSC offline

If the edl has been created on a NTSC 24p Avid (if tapes were 29.97 with 3:2) you will have a 24fps edl (or at the very least a 29.97 edl where the 3:2 can be removed using an edl converter – though this is less desirable). The processes are the same as described above.

If you're confused, it is important to remember that the timecode of the original beta tape used for transfer is now irrelevant. The only reference that matters is the NEW 24fps timecode that Avid forcibly applied to your footage and the original keycode (this helps as a perfect reference to the original film frame).

Telecine Data Transfers

The 24fps edl is loaded into the scanner and the punch hole set at 01:00:00:00 etc. and the frames are scanned and given a 24fps timecode, which is stored in the DPX file's header, and we have thus long forgotten that we ever went to something other than 24fps for the dailies/offline transfer.

However, as described above, there are 2 methods for data transfer for the online - uncut & cut neg.

Uncut neg. is very straight forward as it is just a high resolution version of the previous offline transfer. Each roll and its punch hole must be matched up and set to its hour timecode [01:00:00:00 for roll one, 02:00:00:00 for roll 2, etc.] and as you transfer, the keycode reader will take the selected shots it needs from your edl using the telecine's pull list option and transfer them as 2K. As each frame is scanned, the timecode (and keycode if necessary) from the edl will be added to the header information for that DPX frame

However, it is often the case that original camera negative [OCN] is cut into 'selects' with a couple of feet of film for 'handles', or cut 'flash to flash'. This has the benefit of reducing the amount of physical film storage needed - but does mean that the negative has been handled more, although with enough handles duration this shouldn't cause a problem.

The original edls from the Avid will be converted by the neg. cutter to represent the new selects roll and thus will not affect the procedures described here.

Regardless of the method used, cut or uncut neg., as each frame is scanned, the timecode from the edl will be added to the header information for that .DPX frame, the frame rate will be irrelevant [although the actual scan speed will usually be 4 to 6 fps] as the timecode is purely acting as a frame reference marker. These .DPX frames can be loaded into the DI system via the network and conformed from the original edl. The offline tape can be played alongside simultaneously so that the operator can check the on-line conform for frame accuracy (assuming the DI system in use has this capability).

Motion Effects

It is always important to remember that simple motion effects can often look terrible if used at film resolution. Firstly, slow downs, if absolutely necessary should only be done at 50% of original speed. What will happen is that every frame will be duplicated once, and thus this can have the disadvantage of producing a 'strobe/judder' effect on the movement in the shot. Any odd number, such as 33.3% or 66.6% should definitely be avoided, as this will cause every third frame to be duplicated etc. this produces a terrible judder, which may look acceptable at compressed standard definition resolutions, but at 2k will be less than desirable. Speed ups will also suffer from the same problems, the speed should only be doubled, tripled etc. so that frames are dropped in a regular pattern, doing a 133% or

166% speed up would mean that every third or fourth frame will be dropped, thus producing the same kind of judder as experienced in slow downs.

As a general rule, only use motion effects if the shot is locked off, and there is no movement on the screen. If smoother motion effects are required then this must be done as a proper visual effect using proper motion interpolation.

In all cases tests at 2k are recommended as image softening is likely – its best to shoot slowdowns in camera!

If any shots are sped up using profile stretching at the offline stage, (where a curve representing speed and time is used) the values will not be transferred through an EDL, these effects can only be recreated as a visual effect, and thus will incur time and cost implications.

It is very possible that the online machine will use different frames to create motion effects than the Avid, it is thus advisable if ramping is desired OR the choice of frames is important that the events be 'frame cut' in the final EDL (though it is suggested that a separate basic EDL for the scanning stage be provided so the required frames exist in one shot rather than lots of individual shots with only one frame inside!)

Online facilities generally accept only CMX3600 edls. This format can only hold information for standard dissolves and standard motion effects (M2). This means that only normal dissolves should be used, these exclude 'Dip to Colour', 'Film Dissolves' etc. If dipping to black, this should be done as a dissolve to a frame of timeline filler, then dissolve back up (particularly used in trailers!), if this has not been done it is common that the online facility will ask the editor to re-supply the edl correctly.

Note: If dipping to a colour other than black then 'Dip to Colour' can be used, but this will not conform as it will appear as AX in the edl (Avid effect), so a comment such as 'Dip to White, centred on cut' should be added to the edl. Comments should be supplied with the edls; information like 'Resize here – zoom in', 'Pan across' etc. are useful (colour corrections are not!). Comments should be used for important information only, which are relevant to the online editor and the online conform.

It is standard practice that the edls supplied will be per final film release reel length, with separate edls for any extra video layers that is has been necessary to add, i.e. titles etc. These edls will be sent to the neg-cutter if a selects reel is being created.

VFX – offline to online

There are a number of ways to deal with VFX shots and the workflow is usually slightly different for each project, but as a rule here is a way of dealing with the issues correctly to keep everything as smooth as possible in the online and help in making sure that the feature is delivered on time.

It is common that temp VFX shots are constantly supplied back to the offline editor to see 'in situ' amongst the other shots.

This is where careful work has to be carried out by the editor or his/her assistant.

VFX shots can come either on tape or as individual QuickTime/sequential frames. When working with QuickTime/SF the editor would import each sequence and apply the full title of the VFX shot (matching what will be supplied by the VFX house to the Digital Lab) as the Clip Name (moving old information like slate/take to Comments).

QuickTimes that are supplied by the VFX facility must have the same number of frames as the shot that will arrive for on-line. If slates or extra black are being placed in the QuickTime sequences with the editor using this as a source rush clip in the edit, this will not match back to what will be supplied by the facility. It is the responsibility of the VFX facility to send QuickTime's that match identically the full resolution 2k shot; same frames, same handles, same duration – otherwise reels will not auto-conform and unnecessary hours will be lost eye-matching.

Each VFX shot will need to take on its own reel name and source timecode of 00:00:00:00 (which is the standard timecode that Avid applies to imported Quicktimes etc.) When the temp VFX are supplied to the offline editor these should be placed in V1 (with the original going above on V2), the clip name should be entered exactly as what will be supplied by the VFX facility to the Digital Lab.

E.g.

003 AX V C 00:00:00:00 00:00:03:04 10:00:01:21 10:00:05:00
FROM CLIP NAME: 030_DOG_VFX_V8

The online DI system then looks for AX events and then does this:

003 O_030_DOG_VFX_V8 V C 00:00:00:00 00:00:03:04 10:00:01:21 10:00:05:00
FROM CLIP NAME: O_030_DOG_VFX_V8

This simple process allows us to make sure we always have the correct VFX version number in the edit. Again, for this process to work, the VFX facility must keep the naming conventions (including case) and number of frames identical between the Quicktimes/folders with low res sequential frames and the highres versions, otherwise confusion regarding shot numbers and revisions can become a big issue.

Existing DI Operations

Having defined digital film as realistically requiring no more than 2K working resolution and as being 10 bit log, and having shown how simple end-to-end calibration can be, we can look at what the operational requirements for a DI environment mean for the physical system.

Fairly obviously, any DI environment needs to be able to ingest a standard offline edl and conform the online to it [both image & audio, timecode & keycode], regardless where the online material has come from [scanner, digital camera, telecine/data scanner] making frame accurate edits for image and audio. For future offline/online compatibility AAF [Advanced Authoring Format] conforming will be required and should therefore be available in any of today's DI systems. Once conformed, the DI system needs to be able to review the entire project, on demand, at full 2K resolution, with temp shots, animatics and previzualisation scenes as well as up to 8 channel audio included, making very specific demands on the technology.

The DI system also needs a full colour grading toolset offering both chemical grading/timing capability as well as more video orientated colour correction features. Ideally, these capabilities need to operate in conjunction with the system's editing toolset to maintain flexibility.

Also required are vfx tools, both for effects compositing and pan & scan composition work.

And to cap it all, these tools need to be available within one system, limiting the need to move data and environment and incur the associated additional cost and expanded timeframe this entails.

The traditional digital approach of distributed workflow environments means that shots are spread across many, low[ish] power/cost, platforms where specific vfx functions are carried out, with no central environment where the whole project resides. Additional dedicated platforms are used to review shots, usually in proxy [lower resolution] form with other platforms used for editing and colour grading, and yet more for background rendering at full & final resolution. There is rarely an occasion where all material can be found co-located in final & intermediate form, other than at the end of the project when the material is recorded back to film or to the digital master. This makes review of work in progress almost impossible, especially seeing any given shot in context.

New generation DI systems have to be able to hold the entire project, including active edits, work in progress, temps, etc, along with the low-resolution offline [Avid/FCP cut] for reference. It will still transfer out specific shots to dedicated vfx systems for complex and time consuming work that is not cost effective to perform within the DI 'hub', but this is as part of a controlled workflow approach rather than as a necessity demanded by the restrictions of the systems in use. At any given moment the client [DoP, Producer, Director, etc.] can review the whole project at full resolution in its latest form, interactively and on demand with all associated 'tails' and alternate versions. If a cut needs to be seen as a dissolve or a dissolve as a wipe, or its duration lengthened to confirm a scene, it can be reviewed, changed, and reviewed again. And with the correct image colour temperature and colourimetry, contrast, gamma, etc.

Such capabilities obviously demand high disc capacity, powerful real-time processing and an interactive and intuitive user interface. It also demands flexibility and scalability: things that have previously been incompatible. You either went for power and capacity and bought dedicated black box systems, or flexibility and scalability and bought standard PC systems. But times have changed.

Virtual Telecines

Although the DI concept has been with us for some time, the only near contenders to true DI workflow have been virtual telecines. These have been the only systems to offer a semblance of integrated edit conforming with optical transitions and the more obvious colour grading. Their weakness is that they are all based on comparatively old digital technology & methodology – often servers and disc storage with the front-end processing from a telecine combined with hardware based colour graders such as from da Vinci and Pandora. This results in a 'disc to disc' workflow with the image data passing through the hardware, including colour corrector, as it moves from disc to disc. This is a very linear approach when compared to alternative systems and has associated restrictions that within the video world resulted in the dominance of non-linear systems.

The use of a telecine front-end can be a specific problem because such hardware expects data to come directly from a film scan, forcing specific scans & not allowing multiple grade passes through the system due to colour distortion and gamma compounding. In essence such systems are nothing more than their name - 'virtual telecines' - a phrase they were once proud of until the true restrictive nature of such workflow became apparent in the face of better competition.

Re-purposed VFX Systems

Outside of virtual telecines a number of companies have attempted to re-purpose vfx equipment for use within the DI environment. While some companies have had success, it has been through the use of very restrictive working practices, practices that will not stand up against the new generation DI operations.

To look at this further, one of the fundamental benefits of a DI environment should be its ability to retain flexibility throughout. To make a vfx based environment work for DI use companies are forced to grade all shots as they are scanned into linear [low dynamic range] colour space [10 or even 8 bit]. This prevents [or at least restricts] further colour manipulation during or after the editing process, greatly restricting future options. This alone will guarantee the demise of such an operation in the face of new generation DI systems and associated possibilities.

By comparison virtual telecines provides a more flexible environment, with post-process colour grading being a key part of their operation, but limitations exist due to being 'virtual telecines' with no vfx capabilities [& limited editing], requiring material to be exported to additional systems for even the most basic of vfx work, combined with the colour issues listed above. Between the two [vfx systems and virtual telecines] there could be the basis for an interesting DI environment but at such cost and slow throughput [workflow] that the cost per frame [project] would negate any perceived benefit of going the digital route.

Distributed Workflow Systems

Other players coming to the market are those that wholly embraced the ideas of distributed workflow as a methodology for DI, coming, as they do, from a multi-workstation, computer based environment background. Such systems can in part be considered 'virtual telecines', but unlike those based on dedicated hardware, without true real-time 2K operation, with associated editing & vfx systems as part of the mix. As such they expected post facilities to use multiple systems for data i/o, editing & review combined with vfx systems for vfx work, necessitating major data transfers between environments that has a direct and detrimental impact on the foreground operation of the systems involved [they slow-down, a lot!] with resulting cost implications. The problem with this approach is that it follows the historic workflow methodology of existing vfx operations, which as previously stated cannot offer the business model [price] necessary for the major DI market and clients.

Shared Area Networks (SANs)

An alternative that is growing in interest [and is closely aligned to Distributed Workflow systems] is to set-up a SAN with a number of workstations all accessing the same data via 'virtual memory' type operation [no copying of the material with each workstation treating the SAN storage as its own local storage]. The idea is that multiple workstations can work on a single job without the need to copy material and therefore throughput is sped up.

However, without easy access to the meta data that contains history information [invisible copies if you like] it is next to impossible to undo work so making errors can become a major problem, forcing a return to the source material. Because of this, job management [data wrangling] becomes a major part of the process and a high level of care needs to be taken, as no single location [workstation] is ever aware of the true state of the project.

Even more important, it is impossible to guarantee 'quality of service' making it difficult to provide true real-time playback for high-resolution images without serious caching of data – and that means playback durations will be restricted.

Quantel's iQ Pablo

As a DI system, in my eyes as an operator with multiple DI projects to my name, iQ has no better and it's hard to find any peers. It is a system I have a lot of first-hand experience of, mainly positive [well, it is software based!] while also having hands-on experience of the alternate system types mentioned above.

iQ is a system based around a standard industrial PC for flexibility but with a dedicated image processing engine and dedicated disc storage for power, capacity and scalability. With a standard toolset designed for DI operations and access to any third party software development [sparks and plugins] it offers, for the first time, a truly integrated DI environment.

From the previous description of the alternatives it becomes quickly apparent the benefits iQ offers. It is a single, integrated DI environment with built-in audio & vfx capabilities and the ability to hold entire projects and variants there of. With client attendance being of vital importance, the interactivity of iQ stands out even further, enabling DoP, producer, editor and Director to remain interactively involved throughout the DI process. Working in 10 bit Log or up to 16 bit Lin depth, with 16 bit DR [Dynamic Rounding] processing per colour at full 2K [or 4K!] resolution, with the ability to playback entire projects, complete with uncommitted edits, in real-time it is a truly powerful DI environment.

A growing number of digital film facilities are basing their DI operation around iQ, having evaluated the alternatives under real working conditions, integrating into system wide networks to scanners, film recorders and additional vfx systems, building operations capable of performing multiple Digital Intermediate film projects annually.

For further information see www.digitalpraxis.net

Steve Shaw
Digital Praxis Ltd.
+44(0)7765 400 908
steve@digitalpraxis.net
www.digitalpraxis.net