

UNIVERSAL CONTENT PRODUCTION — MORE BANG FOR THE BYTE

Dave Bancroft

Thomson, United Kingdom

ABSTRACT

The broadcasting and film-making industries are fragmenting in one sense yet uniting in another. Broadcasting is separating into content creators versus content distributors, while the film industry is embracing a new mix of traditional photo-chemical and cutting-edge digital.

The driving forces? Content creators want to sell the same content to as many buyers as possible; content distributors want to be paid to transport content in any format. Significant economic benefits will accrue through reduced duplication of work and streamlined workflows in this realignment.

Disparate technologies, standards and practices used to obstruct these desires. Now broadcasters are taking steps to make content more universal (metadata, MXF, etc.) and film-makers are adopting the "digital intermediate" approach to make single-inventory masters for theatrical, broadcast and packaged distribution.

BACKGROUND

Television broadcasting has historically used video technology end-to-end; the cinema industry has used photo-chemical film technology end-to-end (Fig.1).

The linkage between these two industries was limited, e.g. telecines played complete films into the television chain, while film recorders made permanent film copies of complete television programmes.

This started to change as post-production rapidly evolved following the introduction of VTRs and electronic editing.

FORCES OF CHANGE

Packaged Media

Packaged media (initially VHS, now DVD) was added as a secondary distribution channel for these industries (Fig. 2). Although packaged media formats can obviously carry both kinds of content, feature films, being less ephemeral, dominate this outlet channel.

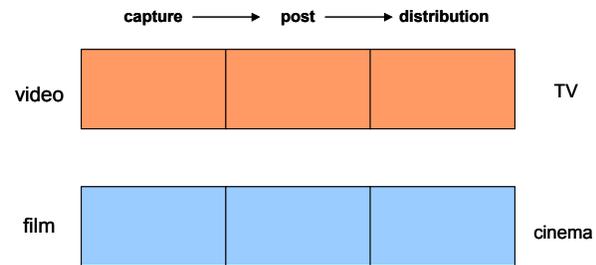


Figure 1 - Legacy Industries

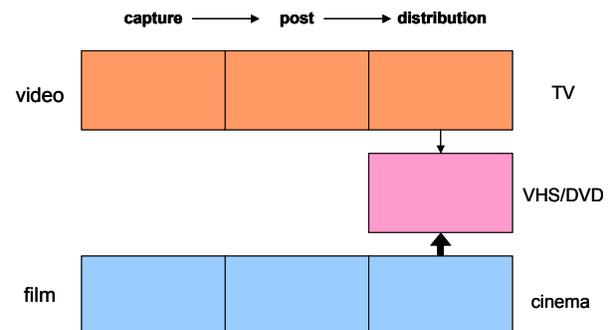


Figure 2 - Addition of Packaged Media

Digitisation of Post-Production

Initially, the television and cinema industries adopted separate techniques as digital post-production entered their respective domains. Television retained video formats similar to acquisition and distribution formats; cinema post-production adopted a data-oriented architecture that has become known as the "Digital Intermediate" workflow (Fig. 3).

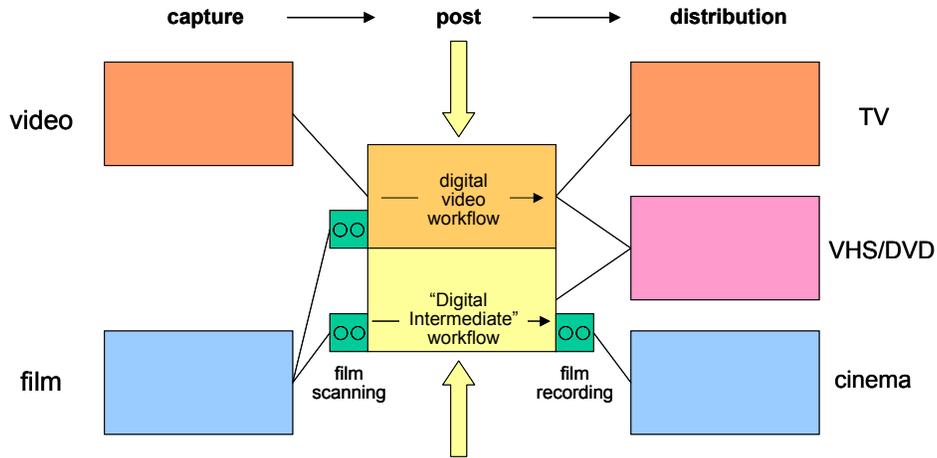


Figure 3 - Digitisation of Post Production

Commonalities in digital techniques have started a trend towards harmonisation of the two. However, a lot of television is live in capture (if not in transmission), requiring "streaming" capabilities not needed in movie production.

Cross-Fertilisation of Television and Cinema Production Techniques

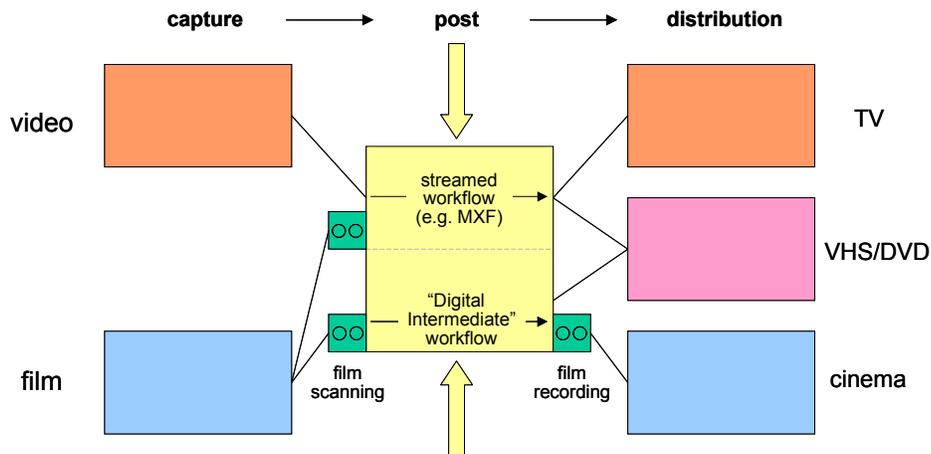


Figure 4 - Cross Fertilization of Production Techniques

The continuing trend towards common IT-based approaches brings less and less difference between architectures for TV and film post-production (Fig. 4). However, the streaming requirement for TV will always be present. Comprehensive content exchange formats such as MXF* embrace both streaming and file transfer capabilities within the same framework,

* Material Exchange Format

but an architecture such as AAF[†] is also needed for more intensive and complex capabilities in file-transfer-only (non-streaming) productions.

Completion of Digital Evolution (long-term future)

(Fig. 5) Capture methods will be distinguished not so much by fundamental technology choices such as electronic or photo-chemical, as by the type of scenes, action or events being captured. The critical difference will always be between “staged” events that can be broken down into manageable chunks and started and stopped entirely under the control of the director, and “live” events for which the moving image medium is mostly an onlooker. This difference dictates factors such as frame rate, exposure range, where and when is image processing performed, etc.

All post-production will probably use IT-based, data-centric architectures, with or without streaming, as required.

Distribution will be characterised as public arena (e.g. cinema) versus home environment. The latter is broken down into streamed and packaged.

The significance of this diagram is not that there is anything really new happening here in terms of process or flow (except in post-production). What matters is that the differences in capture and distribution types will not stray outside the capabilities of a single, homogeneous, digital environment.

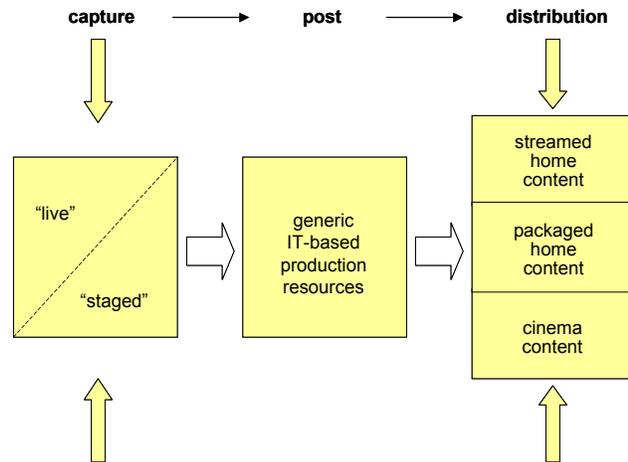


Figure 5 - Completion of Digital Content Evolution

FOCUS ON DIGITAL INTERMEDIATE WORKFLOW

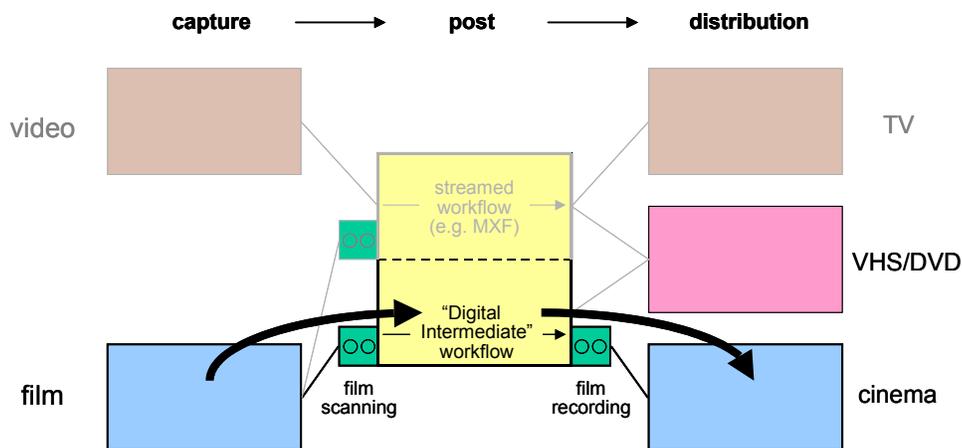


Figure 6 - Focus on Digital Intermediate Workflow

To return to the present time, the Digital Intermediate process and workflow is the focus of this paper (Fig. 6), because it is an area of particularly rapid change. What is meant by the

[†] Advanced Authoring Format

term “intermediate” is that a digital process is in an intermediate position in the production flow between analogue (film) capture and analogue (film) distribution and display.

The key performance issues for the Digital Intermediate workflow are:

- speed
- usability
- co-existence with conventional film capture and distribution during a long and gradual migration period towards digital cinema capture and distribution.

Key enabling technologies are required to deliver these system attributes (Fig.7)

The key technologies are in:

- film scanning (1)
- film recording (2)
- the workflow itself (3)

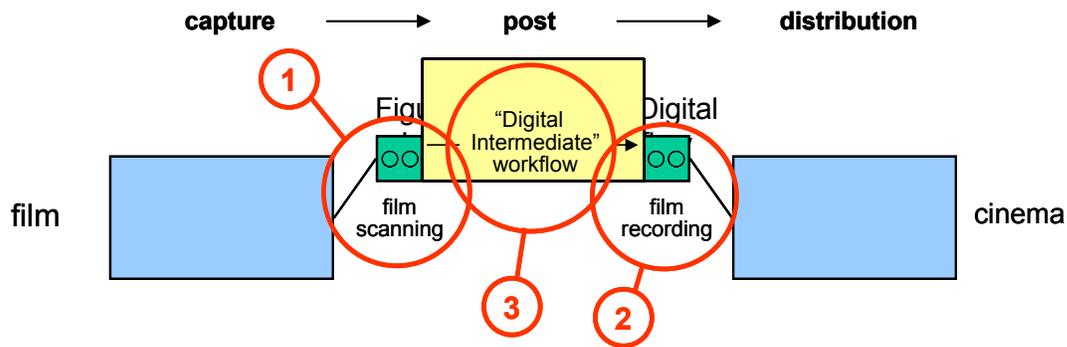


Figure 7 - Key Enabling Technologies

Film Scanning

	Pin registered scanner (1)	Pin registered scanner (2)	Continuous motion scanner
Scanning speed at 4K resolution	0.2 frames/sec	1 frames/sec	8 frames/sec
Time to scan 1.6 million frames - 100 min. feature movie (10:1 shooting ratio), at 4K resolution	109 days	18½ days	2 1/3 days

Table 1 - Effect of Different Film Scanning Speeds

For “Digital Intermediate” to achieve its full promise, every frame of the movie, not just certain scenes, must pass through the digital domain, and must do so “transparently,” so as not to lose any of the information on the original film*. In fact, “every frame” refers to not just

* The Digital Intermediate must achieve the same flexibility as the original film from which it is derived, i.e. function as a

all of the finished content, but includes a lot of material that will not make it into the final version, e.g. the “out-takes,” and also supplementary footage, e.g. background plates, in addition to principal photography.

When this volume of material is combined with the amount of detail required to represent each frame of film from the camera without significant loss, the total burden on the film scanning device is enormous.

The Digital Intermediate process began some years ago with the scanning of isolated scenes and sequences only, and thus relatively slow intermittent-motion pin-registered scanning devices sufficed, but the “total movie” scanning requirement described above has mandated an enormous increase in scanning throughput (Table 1).

Accordingly, new continuous-motion, high-speed, high-resolution film scanning equipment has recently become available, capable of scanning the useful limit of 35mm film resolution (“4K” resolution[†]) at multiple, instead of fractional, frames per second.

Correspondingly higher-speed interfaces and capture servers have had to be introduced to keep up with the enormous data output flow rate of the new scanners (see “Core Technology Requirements,” below).

Film Recording

The film recorder might be thought to represent a potential bottleneck equal to that of the film scanner in the overall DI workflow speed.

However, the burden is considerably lessened by a number of factors:

- The volume of finished film footage is much less than the original capture footage.
- The relative simplicity of the film recorder compared to the film scanner translates into lower cost, meaning “ganged” recorders can be installed to share the total recording workload.
- The excellent stability and repeatability achieved by modern recorders, together with sophisticated calibration techniques, allows a “ganged” technique to reduce the total recording time considerably, even though the frame rate throughput of an individual recorder is much slower than that of the latest film scanners at the other end of the Digital Intermediate process.

Other “tricks,” such as electronic signal inversion and image re-orientation, allow direct recording to an interpositive print, cutting out the time and quality losses of some of the subsequent optical print generations (for limited numbers of release copies).

Some performance indicators from a typical modern film recorder would be:

- maximum film magazine size 2000ft (22 minutes of finished running time per load)
- good stability & calibration enables ganged recording across multiple machines, allowing:
 - (a) make one complete release copy faster by making matched recordings of sequential film reels simultaneously
 - (b) make multiple release print copies faster by making matched recordings of the same film reel simultaneously
- example: 5 ganged recorders can make a print master for a 110-minute movie (5 x 22 mins.) in 25½ hours (4K resolution / 1.85:1 aspect ratio digital input, 2.9 seconds/frame recording time)

universal master suitable for all types of finished output — e.g. HDTV and DVD versions, as well as cinema releases. This is achieved by representing the film information content in full and in a data-oriented format, not a video format.

[†] “4K” resolution means 4096 pixels are resolved across the width of the exposed or selected area on the film.

It is the assurance that recordings from multiple machines will match that mitigates what would otherwise be a bottleneck in the workflow.

Digital Intermediate Workflow

Core Technology Requirements — Interfaces, Processing and Storage

System Interfaces: At 8 frames/second, the latest 4K resolution film scanners are pumping out 459 Megabytes/second (3.7 Gbits/sec). Most of today's "high-speed" interfaces are too slow for this:

- SMPTE 292M 1.485 Gbits/sec¹
- Gigabit Ethernet 1 Gbit/sec²
- Fibre Channel 2 Gbits/sec³

More advanced interfaces are needed:

- Gigabyte System Network® (GSN) 6.25 Gbits/sec⁴
- 10 Gigabit Ethernet⁵
- Future Fibre Channel, 8 Gbits/sec and beyond⁶
- Infiniband™, 2.5 – 30 Gbits/sec⁷

Processing:

New servers and workstations equipped with advanced CPU chips, powerful graphics processors and the latest bus technologies such as PCI-Express⁸ are required to perform process-intensive "number crunching" operations on the very large files arising from the high-resolution images.

Storage:

The storage numbers are also imposing: a single movie frame at 4K resolution, RGB 12-bit, full aperture, is represented by 4096 x 3112 pixels x 36 bits/pixel = 57 MBytes. However, to allow processing operations without rounding errors a wider processing bus may be used — 16 bits per colour, for example, making a total of 48 bits/pixel, hence 76 MBytes per frame.

A movie with a run time of 100 minutes consists of 100 x 60 x 24 = 144,000 such frames. However, some production realities intrude: if it is required to make editorial decisions downstream from capture, then all the takes have to be delivered to post-production. If a shooting ratio of 10:1 average is assumed, that makes 1,440,000 frames. If it is further assumed that some additional material may be shot to be composited with the principal photography as replacement background in some scenes, it seems reasonable to assume a figure of some 1.6 million frames to be pushed

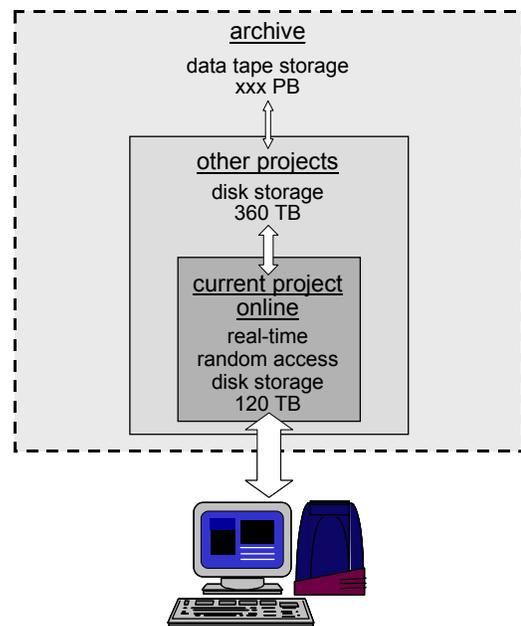


Figure 8 - Storage Hierarchy

through the Digital Intermediate workflow.

This translates into about 120 Terabytes of online storage if access to any part of the movie is required without loading/unloading delay at any workstation (Fig. 8). This would normally be disk storage, with a striped or otherwise paralleled architecture to deliver random access in real time.

If the facility is working on more than one project within the same time frame, it may additionally be necessary to extend the workspace to accommodate multiple projects simultaneously to avoid excessive delays loading and unloading projects. This extended space might not have to be accessible in real time, but it must be relatively quick to swap projects in and out of the real time workspace; it would probably therefore also have to be based on disk storage. Three projects' worth of workspace could be seen therefore as requiring 360 Terabytes of total disk capacity — the Petabyte region is not far away. If less expensive storage were used for this space extension, for example data tape drives, the tape-disk transfer rate would have to be considered a serious workflow pinch point. At a fairly typical 24 Megabytes/second, for example, using a single tape drive to unload one of the 120 TB projects above and replace it with another in the online disk workspace would take nearly 2 months, including the time taken for swapping the 600 200GB data cassettes of one project and the 600 cassettes of the other project! Higher transfer rates can be obtained by ganging tape drives, rather like disk drives, but the solution is then eroding the cost savings compared to disk.

Data tape, however, is well-suited to archival storage, since this is of finished productions that each require far less storage capacity than ongoing projects.

With Moore's Law continuing, the storage volumes mentioned here that seem intimidating today may seem quaintly modest in just two or three years' time!

Rationalization of Storage

These vast storage requirements and demanding access speeds cannot be provided in a piecemeal manner. The extensive and prolonged bulk transfer operations described above would saturate a conventional data network, blocking all other traffic, such as a command to abort a large transfer! A separate Storage Area Network (SAN) will be required (Fig. 9). However, the SAN does require a switched pathway in and out common to all SAN traffic, which could itself be a bottleneck. Modern SANs intended for this application are therefore designed on a "Quality of Service" (QoS) basis, meaning that a new request for data transfer in or out of the SAN will not be granted if it would compromise a transfer already in progress by stealing excessive bandwidth. In other words, the existing "contract" between the first data transfer and the SAN is respected.

Some operations do not fit neatly into this architecture. Continuous motion film scanners as mentioned earlier require a guaranteed synchronous transfer. In other words, the high inertia of the film reels and transport components, together with stability settling times and other factors associated with the film being scanned mean that the machine cannot be constantly stopped and re-started as in a more common asynchronous data transfer. These factors may dictate that the capture storage for the film scanning device be a separate body of storage, possibly a Network Attached Storage (NAS) subsystem. The separate NAS isolates the synchronous transfer completely from other system storage operations.

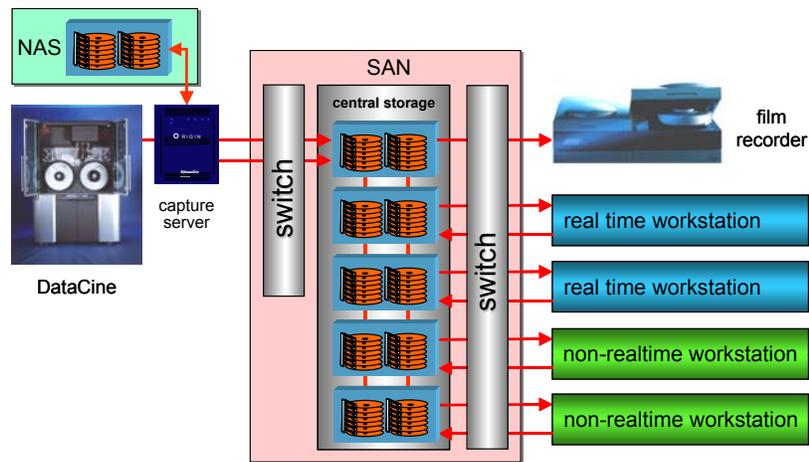


Figure 9 - SAN + NAS storage for Digital Intermediate system

System Level Technology Requirements

Change of workflow model

Traditional sequential workflow:

The analogue medium (film) means generations have to be limited in number to minimize optical losses in image quality and deterioration of originals through excessive handling — this limits the scope for making duplicate copies, which in any case might not be well-matched in tonal, colour or other characteristics.

With this single workprint model, changes made in one operation, e.g. editing, compositing, colour correction, have to be embedded in the content before the next stage is possible. Early digital systems did little to move beyond this model (Fig. 10).

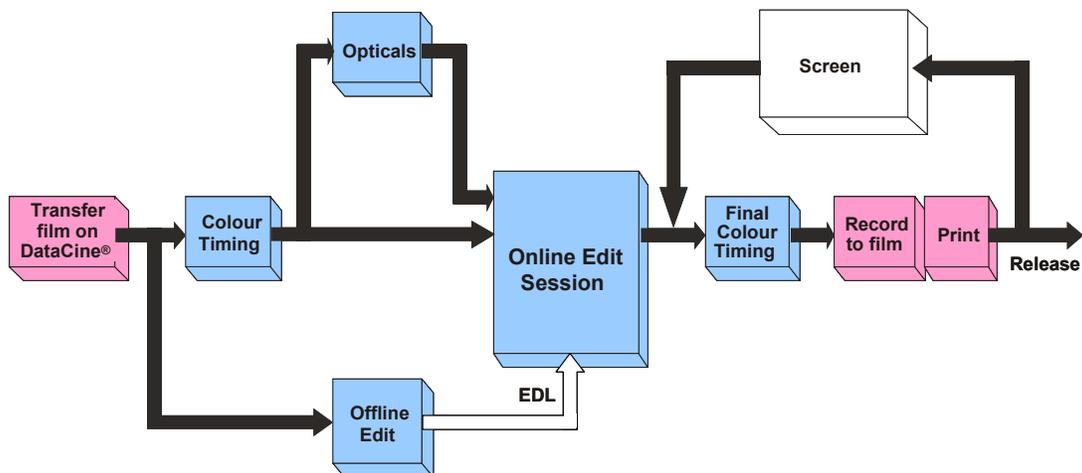


Figure 10 - Sequential workflow

New parallel workflow model:

This is made possible uniquely by digital image representation with the following principles

adopted:

First main principle — use of metadata: View changes you are making to the image but do not necessarily embed the changes into the original image file (i.e. scanned file from DataCine®*). Instead record the list of changes as a set of parameter values and *attach* them to the image file as *metadata*;

Second main principle — think realtime: Movie production can benefit from the realtime, interactive environment that has always been taken for granted in broadcast television production. Selectively applied, this can replace the culture of "guess some changes, send them to the lab and hope it comes back right." An example is real time colour correction on a full cinema-size screen at 2K resolution*;

Third main principle — commutability: Because the original image file is retained in an unmodified state, one type of image operation, e.g. colour correction, can be made independently of another operation, e.g. compositing. Each operator "checks out" the unmodified image file, performs the changes on the monitored image, then "checks in" the unmodified file with the changes expressed in the attached metadata;

Fourth main principle — convolution: To see the effects of multiple change operations, "check out" the unmodified image file and merge all the sets of metadata into a convolved set; use that to control the playout and monitoring of the image file to see how all the changes will look when combined together. The convolved metadata need not be saved because it can be reconstructed at any time and indeed should not be saved while the project is still in progress, since it will be superseded at frequent intervals;

Fifth main principle — rendering: This is the opposite of the first principle. When you *do* want to commit the changes into a permanent version of the production, you create a changed version of the image file by *rendering* or embedding all the production changes into it. This normally occurs at a point in the production flow where there is a change of medium, for example going from digital back to photo-chemical film;

Sixth main principle — recognise exceptions: Some processes are not amenable to the commutability principle. For example, scenes must be colour corrected prior to an optical effect because it is not possible to colour correct the combined image that occurs during, say, a cross-dissolve — each of the "A" and "B" sides must be individually corrected first. This means that an actual rendered version of the modified image file must be created — not just more metadata;

Seventh principle — rationale: The driving force behind the first six principles is economy, meaning economy of time — avoiding waiting for one process to finish before beginning another — and economy of resource usage, which mostly means limiting the proliferation of multiple versions of the huge image file;

Eighth principle — avoiding chaos: Multiple, similar-looking versions of a file made by several different people produce the classic merge problem — how do you prevent the latest changes being applied to the "wrong" version of a file, causing intervening changes to be lost? The principle of embedding changes into a metadata file instead of creating multiple image file versions is the same used in, for example, business document control where multiple authors are involved and in version control for large software development projects farmed out to multiple code module writers. The task of managing the metadata falls to an Asset Management System (AMS). Many such systems have been created

* DataCine®: a machine that, like a telecine, scans film with continuous motion, but which produces an image data output instead of a video output. The name originally came from a product created by Philips Broadcast, now Thomson Broadcast and Media Solutions. The name DataCine® is registered to Thomson.

* "2K" resolution in a display means 2048 pixels are resolved across the width of the image.

primarily with a view to locating completed assets (content), but it is only an extension of the same principle to use the AMS to control the constituent parts of a uncompleted content item (a movie) that is still under construction as a vast collection of scripts, shots,

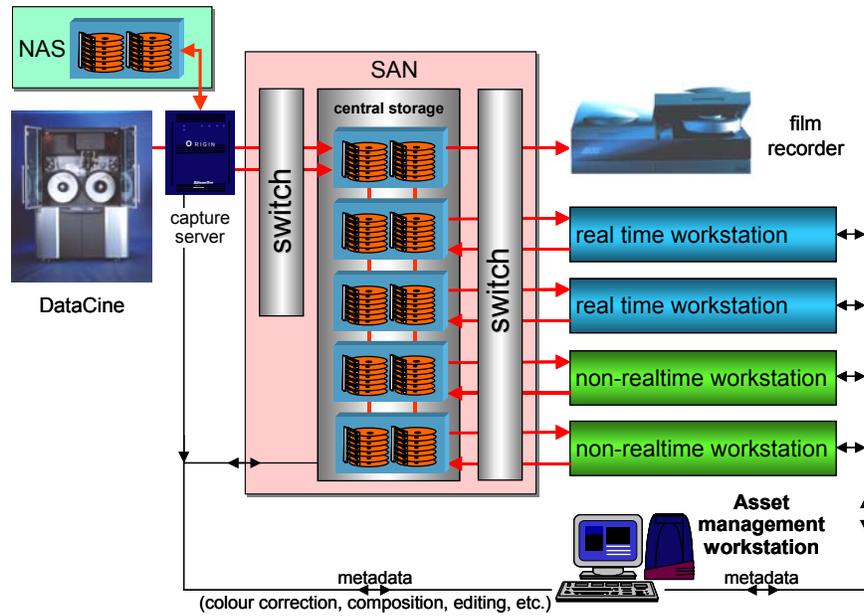


Figure 11 - Parallel workflow

takes and effects, with different sections in varying states of completion.

These principles are shown at work in Figure 11, where a metadata-based asset management system has been added to the SAN/NAS schema shown earlier. The workstations check image data in and out of the SAN for viewing, but check their decisions as metadata in and out of the AMS.

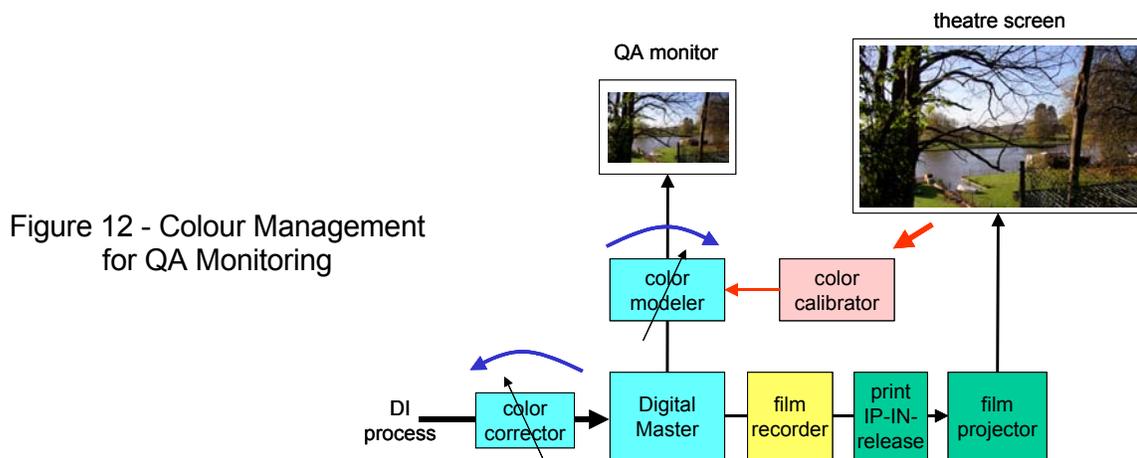


Figure 12 - Colour Management for QA Monitoring

Colour management:

A characteristic of the Digital Intermediate workflow model is that it will must be fully effective as a business during a long migration period from analogue to digital as described earlier.

One of several unique challenges is that of effective colour management.

The term colour management has a large number of meanings. For example it embraces the function of colour correction performed for creative reasons. However, what is particularly meant here is the concept of "WYSIWYG" — what you see is what you get, which is colour management for purely technical reasons. The print industry has already been through this challenge; copy had to be approved for correct appearance on a CRT monitor (additive colour reproduction system), yet be distributed using ink on paper (subtractive reproduction). Making a new print out after each attempt at colour correction on the monitor is not viable. Accurate colour simulation of one display method (printer + inks = reflection print) on another quite different display environment (digital image signal + CRT) has made the process successful where otherwise it would have been impractical.

The same is required in Digital Intermediate post production. The display of the actual photo-chemical film release print via a conventional cinema projector onto a conventional screen has to be measured and characterised in a model and this model then used to produce an extremely accurate simulation on a CRT monitor or digital projector in a screening room. It is made much harder not only by the vastly greater amount of information to be handled compared to the print industry, but also by the very much larger contrast ratio involved (as high as 10,000:1, compared to about 3:1 for reflection prints). Figure 12 shows the basic components of a calibration loop in conjunction with a "modeller" which modifies the way the image looks on the QA monitor to simulate the changes that will have occurred by the time it reaches the theatre screen. The alteration seen on the QA monitor then guides the colorist to adjust the system colour corrector until the image once again looks correct.

The importance of this is that the Digital Intermediate process would otherwise actually slow down production of a feature film. To explain this, consider that in a conventional laboratory process, each time the Director wants to see the effects of changes in a reference master for the release print (the "answer print"), a new test print can be made by noting the changes and feeding them (electrically) to a film printer that has colour balance capabilities (known as "printer lights") so that it can make a new print from the negative. A popular example of such a printer runs at 600 feet/minute, which is nearly seven times the real time speed of 24 frames/second. In comparison, a digital film recorder takes nearly 3 seconds to record just one frame, i.e. 1/72 of real time speed, or 1/480 of the speed of the traditional process. Clearly colour simulation is of critical importance in eliminating the need for test prints at

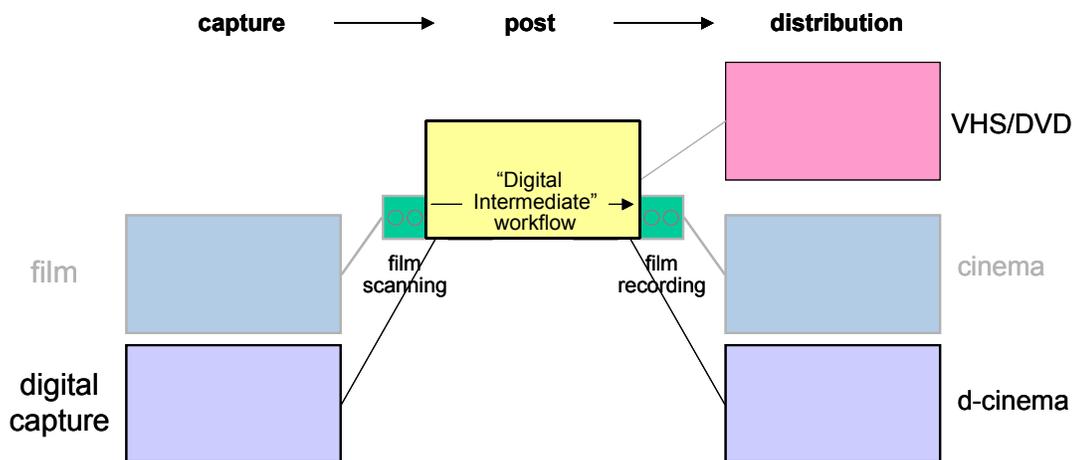


Figure 13 - Digital Intermediate in the Future

such a slow speed.

Future Advances

The previous section mentioned the migration period from digital to analogue. Figure 13 illustrates this graphically by depicting the Digital Intermediate relationship to the capture and distribution stages in the future. Film capture is still present, but has become the minority partner to digital capture*. Likewise, the conventional cinema is still very much present, but digital cinema has now become mainstream. It is fairly apparent that the Digital Intermediate process has always been ready for digital cinema distribution; all that is required is to eliminate the film recording stage.

What is not automatic is that the Digital Intermediate process should also be agnostic to the method of capture, and be able to accept digital capture alongside film capture without modification. This also requires, of course, a certain design philosophy in the digital camera itself, which is another topic, but what is important in this discussion is that the output of the digital camera should not require a separate "video-style" post-production process; it must be possible to feed it into the same process as film camera footage. The reasons for this are manifold:

- digital camera Directors of Photography (DoPs) are likely to have migrated from film cameras and will wish to continue working with the same facility house and post-production team as they did with film
- the facility house and its colorists and effects artists will wish to retain the business of the now-digital DoPs rather than lose it to another facility
- some directors may wish to mix and match film and digital cameras in the same production — a common post-production process is thus essential

The design consideration for the Digital Intermediate process to accommodate digital capture is in principle simple: just as a DataCine® produces, say, DPX image files as a result of scanning film, a digital capture camera's output should result in image files in the same exchange format, with the same image representation (probably logarithmic). In this instance, the DataCine® is then simply replaced by some form of loading station that receives the digital camera's files from the shooting set or location. In some cases conversion from a camera's continuous streaming format to a file container format may be required; this may occur in the camera, in the location recorder, or in the facility's loading station. The bottom line is that the Digital Intermediate process must accommodate the migration to digital capture seamlessly.

Economic Benefits

What has been described so far is the process of merging two cultures. One culture is the television production workflow, which has now successfully absorbed more than two decades of digitisation with enormous benefits in cost reduction; the other culture is the film production workflow, which multiplies the value of the captured footage by deferring decisions about colour correction, apparent exposure and multiple image compositing to later stages, so that the original unmodified footage always remains available for reworking into supplementary deliverables such as broadcast and DVD, in addition to the main theatrical distribution.

It is considered that the economic results of such a merger, when it has been carried to completion, will be:

* correctly termed "digital cinematography"

- feature films are made more quickly, which translates into lower costs
- with digital cinematography as opposed to broadcast-style capture, the place of capture no longer limits international distribution, because using film-related standards removes the 50Hz/60Hz barrier of conventional television; this means more revenues for the same production budget
- production costs for the same quality target drop because a universal digital master eliminates the repetition of costly creative stages to make multiple versions, or, to argue it the other way round, production costs are amortized over a greater total sales value
- archived productions made to universal film-based standards "age" less via obsolescence of standard or format, lengthening revenue-earning lifespan
- easier division of production processes across multiple global locations by means of wide area networking lowers costs of employing best-in-class specialised skills.

CONCLUSION

Feature film production has been the last analogue frontier yet to fall to the digital revolution, but it is now at last succumbing. Of the three components of capture, post production and distribution, so far post production has made the earliest and biggest strides into digitisation via a process now becoming well known as the digital intermediate process. Tremendous advances in processing and interface speed, storage capacity and software sophistication are making this possible.

At the same time, with similar changes having begun much earlier in broadcast television, the two industries now have the potential to share much of the new information technology they are both importing. The result is that there are many new business opportunities arising that will blur some of the distinctions between the two media types, movies and television. Their respective artistic and creative characteristics will remain distinct, but they will be less and less differentiated by technology, to their mutual economic advantage.

COPYRIGHT NOTICE

Copyright held by IBC. First published in the IBC2004 Conference Publication.

REFERENCES

- ¹ Television Standard SMPTE 292M-1998, "Bit-Serial Digital Interface for High-Definition Television Systems," Society of Motion Picture and Television Engineers, <http://www.smpte.org>
- ² IEEE 802.3z, part of 802.x protocol suite, <http://standards.ieee.org/catalog/olis/index.html>
- ³ Fibre Channel Industry Association, <http://www.fibrechannel.org/OVERVIEW/standards.html>
- ⁴ High Performance Networking Forum, <http://www.hnf.org>
- ⁵ 10 Gigabit Ethernet Alliance, <http://www.10gea.org/tech-faqs.htm>
- ⁶ Fibre Channel Industry Association, <http://www.fibrechannel.org/OVERVIEW/Roadmap.html>
- ⁷ Infiniband™ Trade Association, <http://www.infinibandta.org/specs>
- ⁸ Peripheral Component Interconnect Special Interest Group (PCI-SIG), <http://www.pcisig.com/home>