High Dynamic Range Data Centric Workflow System

Lucian Ion, Alan Lasky, John Coghill

This paper originally presented at the SMPTE Technical Conference and Exhibit, New York, November 2005

Introduction

The DALSA Origin camera is an 8 megapixel (MP) single CCD sensor motion picture production camera capable of capturing images containing in excess of 12 stops of exposure latitude and outputting 16bit per pixel linear RAW data files directly to non-linear recording media. This paper examines the key workflow elements required to move through the entire production and post production process in such a way as to not compromise image quality while also offering options for integration with today's state of the art systems. The key elements to be covered are;

- On set record, review and look management tools
- Data back-up, transport and verification logistics
- Post production considerations
- Global data management considerations

The primary philosophy underlying the design of the Origin camera is to provide a tool that enables cinematographers to create images with the look and feel dictated by the creative story telling process as opposed to the technology used to capture the images. This approach drives the requirements of a host of workflow solutions that are fundamentally different from high definition digital video based solutions that have historically been designed to trade-off image quality and flexibility to maintain lower bit rates imposed by the technical limitations of linear tape based systems.

On Set Record, Review & Look Management

Taking a system perspective, the primary question is "*how do you ensure the creative intent was captured and how does that intent get conveyed throughout the pipeline*?" The output from any system can only be as good as the input and the set is where it all begins.

RAW Linear Format

The fundamental distinguishing feature of this system is that images are captured in 16 bit per pixel RAW linear format with each frame generating a 16.7MB High Dynamic Range (HDR) DPX file. This approach preserves all of the inherent latitude and detail of the scene without applying any irreversible pre-formatting or image processing decisions up front. There are several advantages to recording in this format.

One advantage of this approach is that gamma and gain decisions are not predetermined by the recording media or display devices being used. This preserves the ability to map different sections of the exposure range in a layered fashion. This is important because capturing all of the available scene information defines the input to the creative process of applying different treatments to various areas of the image to help shape the mood and maximize scene impact. Figures 1-3 are a simple illustration of this principle.

In Figure 1, the "normal" exposure layer, it would appear that the highlight detail has been lost in favor of gaining shadow detail. However, an "overexposure layer" (+2Stops) can be mapped (Figure 2) to pull out the scene highlights then merged in the final HDR Composite Layer (2 layers merged with mask—Figure 3) and mapped for the display device in use to ensure the full range of the scene is conveyed as intended.

Another significant advantage of the RAW format is that color LUTs are not baked in. This approach enables the use of any number of the various emerging on-set grading tools to embed color decisions lists (CDL) into the metadata header of the DPX image file. Given the wide range of on-set viewing conditions and display devices (and the practical considerations of time availability in most production situations) this approach allows the on-set decisions to be used as a baseline reference that can then be later refined by the DP and colorist during grading sessions.

The final and possibly most important aspect of the 4K RAW approach is that it ensures a future proof production investment. Image science and technology evolve at a very rapid pace and the value of a production asset grows over time. Future display devices and distribution media will certainly be capable of much higher image quality over a broader range of distribution avenues so it is critical to ensure that original materials can be repurposed for these uses. Archiving the RAW source material complete with the metadata captured throughout the production chain will ensure future presentations are consistent with the original artistic intent.

Figure 1. "Normal" exposure generated from RAW

Figure 2. "Overexposure" layer (+2 stops) generated from the same RAW data

Figure 3. "Normal" and "overexposure" merged with mask in final HDR composite layer, then mapped for display device to capture full range of scene.

There is of course a price to pay for all of this flexibility. RAW HDR images are not formatted for any particular playback or display device so new hardware and software tools are required to visualize, measure and asses the image content. While these generally operate within the realm of standard computer hardware and software requirements they are not familiar tools to production technicians and engineers and will require standardization (particularly in the area of metadata handling) before wide-scale adoption will occur. Also, at this point less than real-time image processing time needs to be factored into the overall deliverables pipeline however as technology inevitably marches on, faster, more efficient tools will become available for real time (or faster) collaborative production.

Recording

For the purposes of this paper the focus is on *what* gets recorded as opposed to the hardware it is recorded on. The hardware is simply a container for capturing images and associated metadata and will certainly change over time so the most important consideration is to design a data centric approach that is compatible with many sources of production data and can migrate to future formats.

A data centric approach enables seamless integration of image data with metadata from a variety of sources such as digital audio recorders, color grading, compositing and editing tools. Key scene metadata is automatically integrated with the image data stream by the camera and as the image data moves through various production stages more information is added to the metadata package. To ensure reliability and availability there are two parallel streams for metadata encapsulation:

- "technical" metadata captured in real-time during production directly into the file header of each image frame using open standard DPX format and…
- clip and project associated metadata files captured into an industry standard XML format

There are multiple fields supported in both the DPX header and XML clip metadata streams that provide important information throughout the pipeline. Logistical information such as scene/shot/roll information, camera serial number, production team and timecode/timestamps as well as on-set technical variables such as lighting & filter notes, camera exposure rating and frame rate are important to various downstream personnel. Additional information such as file size, bit depth, and color space mapping are also key to various post production decisions.

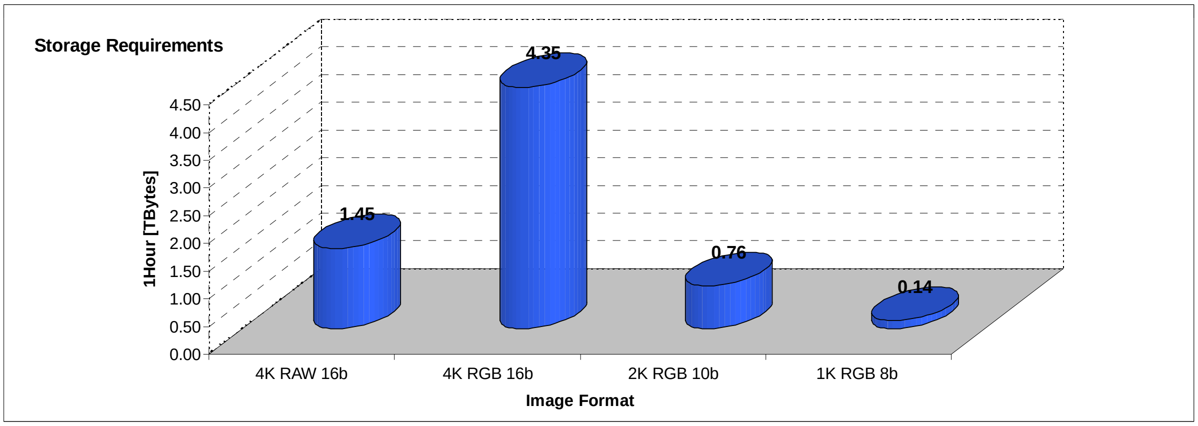

Figure 4. Storage requirements for 1 hour of digital footage

The Clip XML metadata file supports additional information fields such as the number of files in a take, take timing, links to audio files and wedge test results along with dailies review information regarding circle takes, play lists, CDL and DP's comments. This metadata stream allows external applications very fast access to clip related information.

The Project XML metadata file represents a mobile database of information for a complete production. It can be developed starting as far back as the script/story board editor and correlates all the information related to the material captured from dailies review and editorial decisions through to various stages of the processed material such as final grading and conform. Metadata not only improves production efficiency today but helps to ensure any future repurposing of the source material respects the creative intent of the original intellectual property.

One final consideration in favor of recording RAW image data is the fact that it introduces an opportunity for lossless data reduction to the pipeline. As illustrated in the figure below, uncompressed RAW 4K 16 bit data is only 33% of the size of 4K 16bit RGB files but perhaps more interesting is the fact that it is only twice the size of 2K 10bit RGB files. Arguably the benefits of 4K 16 bit HDR images outweigh the additional storage requirements but even then conversion from RAW to RGB files should happen only as and when required to maximize workflow efficiency, but more on this point in the relevant sections that follow.

Given the requirements outlined above, recording to a non-linear media (HDD, Flash, SDRAM) is the only practical approach that can be considered at this time. In addition to the file based metadata considerations, at 24fps the data rate to the recorder is 420MB/sec (RAW, uncompressed) over a quad Infiniband fiber cable and currently available tape based systems are simply not capable of satisfying this bandwidth requirement.

On Set Visualization

In addition to an increased creative toolbox, the primary benefits promised by emerging digital production tools are gains in workflow efficiency through immediacy of feedback. Verifying that the intent of the creative team was captured before striking sets has significant benefit…but only if there is an accurate way to do so! Using a "what you see is *almost* what you get" approach imparts significant risks in today's fast paced, high cost production environment.

The DALSA "Visualize" tool set minimizes these risks and increases confidence in the end result through a simple, intuitive feature called the on-set "Digital Wedge" system. Using a standard Pentium M 1.6GHz laptop connected to the Origin recorder, a full resolution, full dynamic range single frame image can be rendered from the RAW data on the recording device in 15 - 20 seconds. The image is "soft mapped" (see below) for viewing on a laptop display allowing for over and underexposure LUT's to be loaded for exploration of highlight and shadow areas (see Figure 5). Even more important, it allows the DP to instantaneously measure areas of interest to check exposure values, obtain global and line profile histograms as well as color balance information as captured in the RAW source data. This procedure will ensure that the maximum visual information from the set is captured into the RAW source image.

The other significant benefit of this tool is that it overcomes the typical on-set limitations related to suboptimal viewing conditions and multiple display device technologies each with their own color gamut and display ranges. This is achieved through exposure based "soft mapping" of the image data. Unlike high-def camera systems that optimize images for calibrated high-def monitors with baked in display LUT decisions, soft mapping allows a representative sample of the RAW image data to be examined and

X : 1838, Y : 982, Red : 71, Green : 70, Blue : 67, Luminosity : 69

Figure 5. Visualize screen shot

manipulated on any reasonable quality display device that includes a calibration profile. The RAW image is effectively read through a "translator" mapped to the monitor in use leaving the source image data unmodified and retaining the full scene detail and latitude. By capturing the desired "reference look" in the image metadata the same look can be recreated later during the Color Restoration process.

Similar to conducting film stock tests, this tool should primarily be used during pre-production testing to familiarize the DP with the operating parameters of the camera system under a range of shooting conditions. Then, during live production it can be used primarily for capturing creative intent metadata between shots and only occasionally checking exposure ranges in extreme conditions. In either situation, the main advantage is that it will return deterministic wedge measurements in real time.

On-set Grading

A number of on-set data centric color grading tools have been released recently that are designed to enable the DP to make primary color grading adjustments on set. These tools emulate the color space of the display device and generate a CDL as part of the XML metadata in order to convey the desired look to post. These grading tools work on either digital wedges or proxy sequences generated from the RAW 4K data. By manipulating simple gamma, lift, gain and primary color balance variables the DP can embed these look management parameters into the metadata of the reference frame(s) that in turn gets attached to the RAW image sequence for the purpose of conveying creative intent.

Figure 6 is a simplified illustration of the relationships between the various on-set tools that can be utilized on both low and high resolution sample images to convey various production decisions on to the numerous off-set facilities involved in the project.

The primary point to note (and the focus of the following sections) is that the 4K image data should be maintained in its RAW format as far into the process as feasible. In fact, an argument will be made that the 4K RAW format is also the ideal data archival format.

Data Back-up, Transport and Verification Logistics

Three criteria are used to determine the optimal implementation requirements for data redundancy and transport logistics between the set and the various downstream facilities. These are time constraints, cost considerations and of course data integrity of the material. Each of these variables will have a different weighting based on the demands of any particular project but this paper will only touch on the primary considerations.

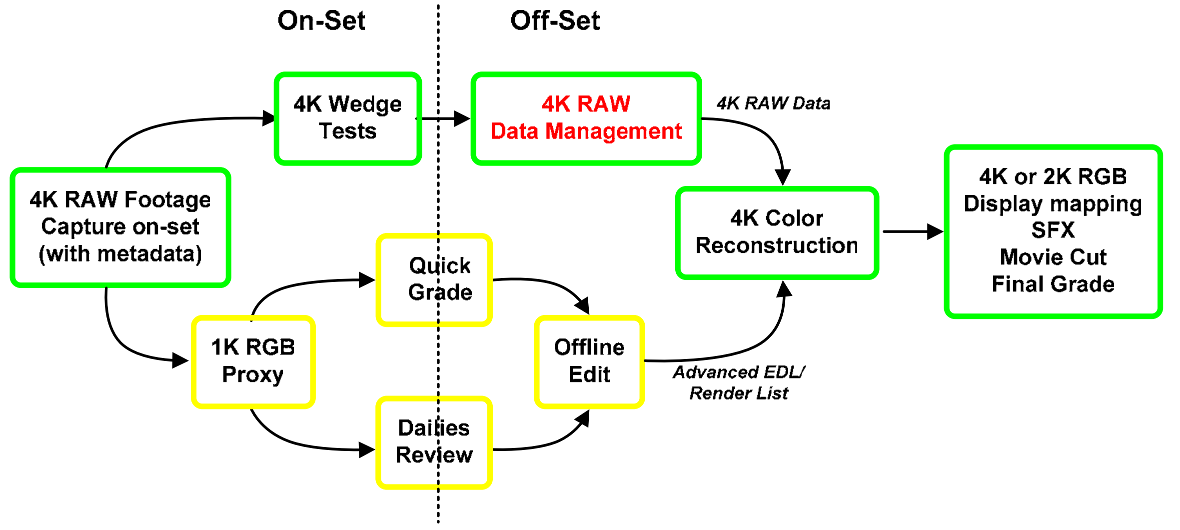

Figure 6. RAW storage and processing workflow

Figure 7 illustrates the data transfer process from the set into the downstream facility. There are two primary steps:

- data offload from the on-set recorder to a Field Transport Magazine (FTM)
- data ingest into the storage array at the facility

The Field Transport Magazine (FTM) is used both as an on-set backup data storage system and a transport container to get image data from the set to the facility. Depending on transportation logistics it may be suitable to utilize lower capacity FTMs that make the round trip to the facility once or possibly twice per day or if distance dictates the FTM may need to be sized to support two or three shooting days worth of 4K RAW footage. In either case the FTM is RAID 5 configured to provide additional data security.

In order to facilitate an efficient, cost effective workflow the speed of transfer is a key consideration throughout the complete data transport cycle. As with the Origin camera, the data interface of choice for the FTM is Infiniband as it currently offers guaranteed quality of service and the highest sustained data throughput at an excellent performance/cost ratio. Alternate solutions with lower performance such as dual and quad Gigabit Ethernet links along with the emerging 10GbE solutions can also be considered.

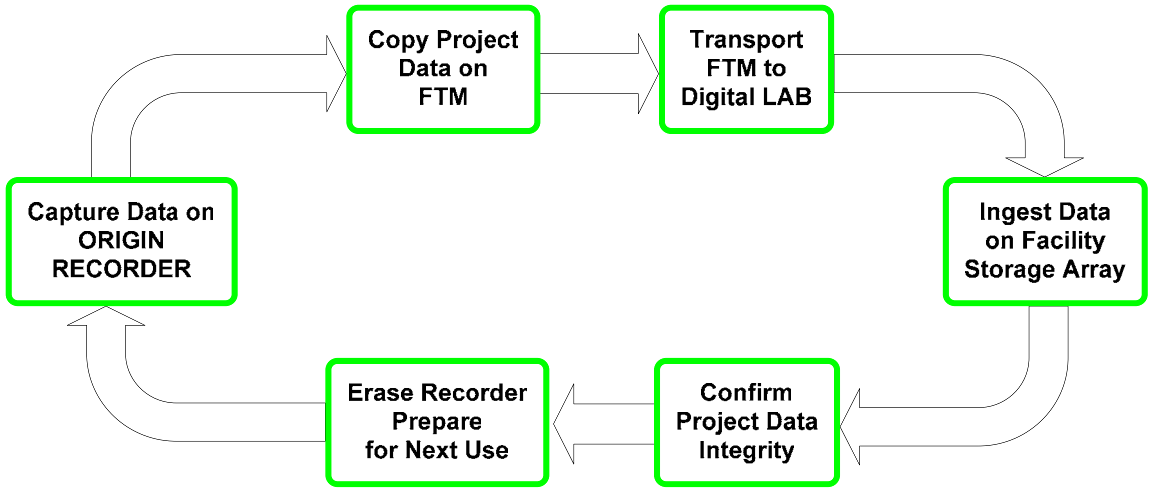

Figure 7. Data Transfer Logistics

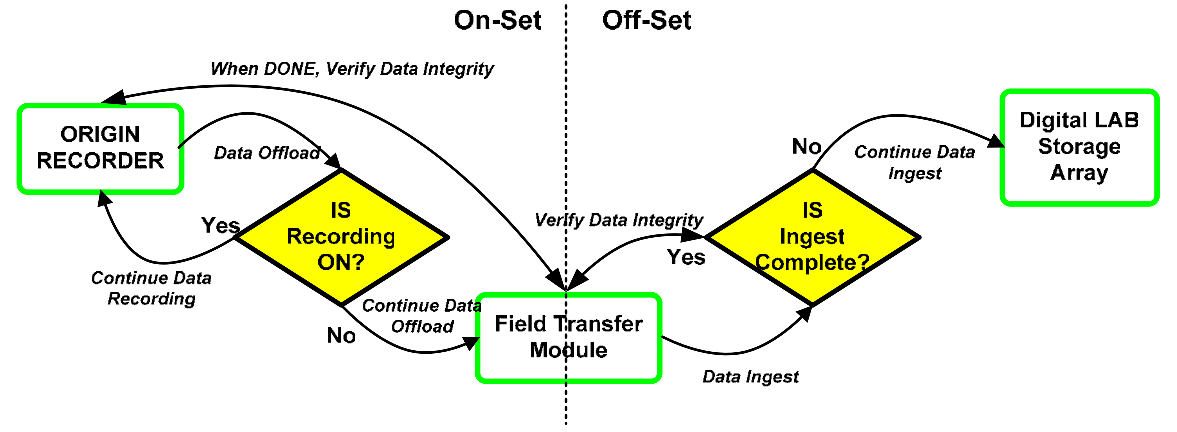

Figure 8. Offload Manager Software

The final key element of the data transfer logistics process is called “Offload Manager Software” (OMS). This application resides on the FTM and the server managing the ingest process at the facility. Its purpose is to provide a reliable data transfer tool for managing the process of transferring data from the recorder to the FTM and from the FTM into the facility.

On-set, the OMS monitors the recorder status and performs an automatic background transfer from the recorder to FTM each time no recording activity is detected thereby providing automatic data redundancy protection. By tracking start/stop points of project session transfers, providing file statistics and performing MD5 verification this tool ensures continuity with previous transfers, confirms complete data transfers and ensures data integrity throughout the pipeline. At the facility ingest point it correlates these file statistics and MD5 results with the previously recorded results to ensure no dropped/missing or corrupted frames have resulted from the ingest process. Figure 4 below provides a high level outline of the operation of this application.

The following outlines four options for transferring image data from the FTM to a downstream facility but first a couple of definitions. Data transfer rates are sustained rates unless otherwise noted.

Receiving Server: A server located at the facility running the DALSA OMS or Synchronizing FTP software used to manage and verify the data transferred from the FTM to storage in the post facility.

SRP: SCSI Remote Protocol over InfiniBand. Storage on the FTM appears as SCSI disks to the server receiving the data. Data transfer rates above 500MB/s can be achieved with InfiniBand if the receiving server has the bandwidth.

IPoIB: TCP/IP protocol running over InfiniBand. Data transfer rates around 200MB/s can be achieved.

Syncronizing FTP: FTP client capable of syncronizing the data on the FTM with the receiving server. Data transfer rate of ~70MB/s over one GbE link and 110MB/s over two links bonded together can be achieved.

Implementation Options:

Option #1: Dedicated Receiving Server running Linux and the DALSA OMS. The transfer will run over InfiniBand between the FTM and the receiving server with multiple options for connecting the receiving server to the facility network such as:

- direct access to SAN (Fiber Channel, etc.)
- GigE or 10 GigE (NFS or Samba)

Option #2: InfiniBand card installed on an existing facility Linux receiving server that has connectivity to the facility storage network. The transfer will run over InfiniBand between the FTM and receiving server using SRP or IPoIB depending on the preference.

Option #3: FTM exposing its storage (NFS, Samba) over TCP/IP with the DALSA OMS or Synchronizing FTP running on a receiving server inside the facility network. Data transfer can be achieved by any of the following (in ascending order of speed);

- Dual GbE
- Quad GbE
- 10 GbE

Option #4: FTM exposing its storage as iSCSI with the DALSA OMS running on a receiving server inside the facility network.

Options #1 and #2 are the most desirable from a performance perspective.

Note: when using InfiniBand SRP, the FTM machine exports its hard disks to the host OS via SCSI protocol. The traffic consists of SCSI commands in a similar fashion to iSCSI with the difference being the underlying protocol (TCP/IP for iSCSI and SRP for InfiniBand).

Regardless of the option(s) utilized to ingest data into the facility this is just the first step leading to a number of decisions that should be well considered, planned and tested in advance to ensure a high quality result in line with the needs and expectations of all stakeholders. This is the focus of the next section.

Post Production Considerations

Earlier sections of this paper used the term facility to denote what could more accurately be termed as a "digital lab". This term represents the organizational entity responsible for the processing and preparation of image data elements required by the various steps in post production. It could operate as a dedicated remote facility for digital processing and data storage or as an internal service element of a more fully integrated post-production or VFX facility. Regardless of its form the digital lab's primary functions are to receive and store (short and/or long term) the 4K RAW image data and provide the following key deliverables/services:

- production dailies (DVD, videotape etc) & editorial files (AVID, FC Pro) with flex files
- wedge test frames with embedded metadata for CDL as created on set
- XML metadata database containing production work in process (WIP) information
- rendering and conforming of final materials
- archival of original source materials and final elements

One of the key decisions to be made early in the project planning process is whether the project will be mastered in either "2K" or "4K" as specified by the recently released DCI specification and what the full suite of release formats is likely to be. For example, a feature production aiming for both 35mm and DCI compliant theatrical releases as well as eventual home DVD distribution will have different image quality and budget considerations than a made for TV movie.

Depending on the mastering decision the RAW data from the Origin camera will then be converted into either 4K RGB or "Super 2K" RGB files. The Super 2K format is compatible with lower bandwidth tape or HDD based "hi-def" pipelines but provides superior results. A 2K image down-converted from a 4K source contains much more image detail than an image originated from a 2K source and when properly mapped, the inherently high exposure latitude of the Origin camera can be maintained in the 2K master. There is an extra processing step required to go from 4K to 2K RGB. The RAW data is first converted to full 4K then down to 2K which takes extra processing time but not extra storage. Any extra cost associated with this additional processing *may* be offset in the short term by competitive pricing considerations driven by the wider availability of 2K pipelines but with sufficient capacity in the render farm this extra time will be inconsequential. Once this decision has been made all the material required for conforming, adding VFX elements and grading can be rendered and delivered in the most appropriate format.

RAW to RGB Processing

Much like film negative RAW files require processing to generate the high quality RGB files from which the image sequences required for the various downstream editorial and post production processes are derived. Also like film, the process of converting RAW image data to high quality images requires both an in-depth understanding of the performance requirements and design of the image capture media and the "secret sauce" required to extract optimum results from the media. There are numerous misconceptions regarding the quality attributes of images derived from RAW images especially when captured using single sensor imagers utilizing what is known as a Bayer color filter pattern array (CFA).

Figure 9 illustrates the basic arrangement of the CFA in a Bayer pattern sensor.

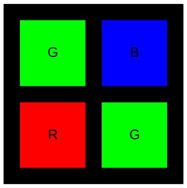

Figure 9. Basic 4 Pixel Cell of Bayer Pattern

While a Bayer CFA imager does produce a sub-sampled version of the image when compared to a three-chip architecture of an equivalent number of pixels per sensor, two common misconceptions are:

- each pixel in the array *only* captures photons from either the red, green or blue portion of the spectrum requiring interpolation to fill in "missing" spectral content in each pixel therefore "true RGB" is not possible
- the fundamental "resolution" of the system is determined *only* by the number of green pixels at 50% of the overall pixel count and therefore the MTF is limited along with the ability to detect fine detail.

A quick review of the characteristics of the imaging system employed in the DALSA Origin camera will help explain why these misconceptions are not true.

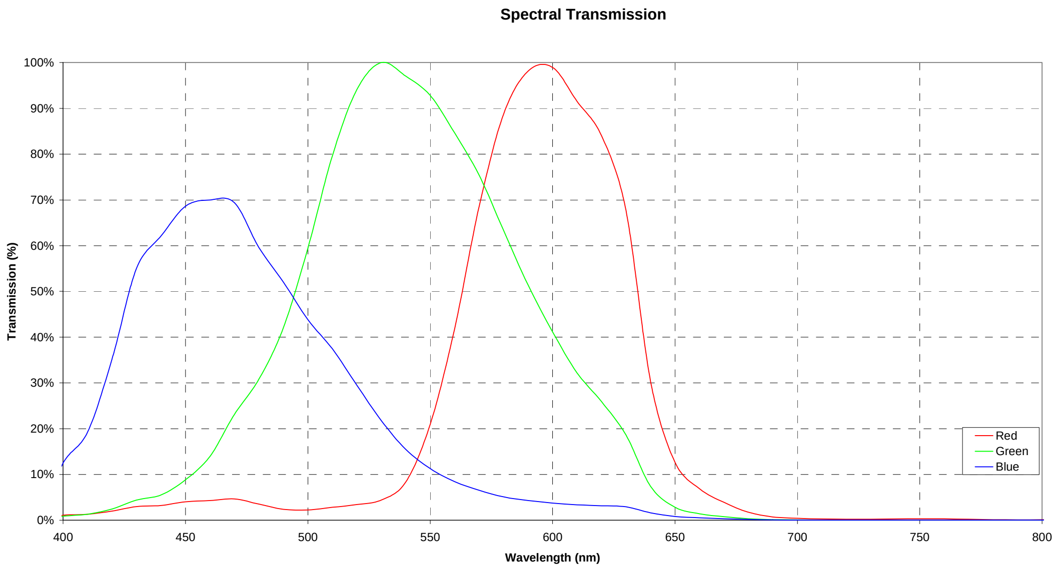

Figure 10. Spectral Response of Origin Sensor CFA

From a spectral transmission perspective a very precise color filter design has been employed to ensure a large color gamut that closely resembles the capabilities of the human vision system. The native spectral sensitivity of the Origin camera is presented in Figure 10.

From the plot above it is apparent that each pixel in the array receives photons from <u>all 3</u> primary color components not just the dominant primary of the local pixel filter. This is achieved by way of the response overlap between all three filters and the "secret sauce" is the foundry process controls required to make the area under these curves consistent, particularly the cross over points. This is critical as it enables the design of the color reconstruction algorithm to produce repeatable, accurate results. The reconstruction algorithm also takes advantage of other subtle pixel performance attributes derived from the photo-gate architecture of the device to give a more "organic" look to the final images when compared to ILT photo-diode architecture devices.

To ensure repeatable, consistent results from camera to camera as well as compatibility with other source materials, precise color calibration is employed to map the output of each camera to the defined CIE XYZ color space with unique coefficients stored within each camera. From that point, mapping to the color gamut of the intended display or distribution media is achieved through adjustments to color matrix values and LUTs. Figure 11 provides a schematic overview of this process.

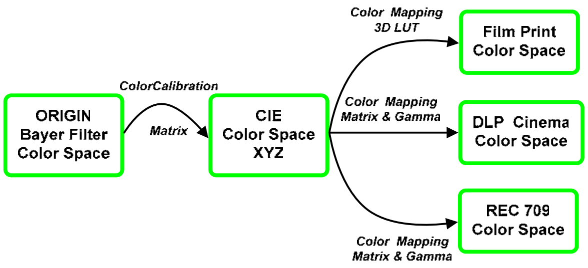

Figure 11. Origin Color Mapping Diagram

The other critical element of any imaging system is the ability to resolve fine detail. A commonly made mistake is to confuse "resolution" with "resolving power". <u>Resolution</u> is defined as the number of pixels in the final image which in the case of Origin is 8MP consisting of 4046 horizontal by 2048 vertical active imaging pixels. <u>Resolving power</u> is determined by the Modulation Transfer Function (MTF) of the system and is a little more complex subject.

The resolving limit of a sampled imaging system is defined by Nyquist's sampling theorem to a maximum number equal to ½ of the sampling frequency of the imager. This point is presented on the plot below as 1x Nyquist. For a 4K digital imaging system 1x Nyquist translates to 4046/2 =2028 line <u>pairs</u> per picture width. The design of the Origin camera and reconstruction algorithm allows for MTF numbers of approximately 80% of this maximum theoretical value, or better than 1600 line pairs. Figure 12 plots the MTF of the Origin camera.

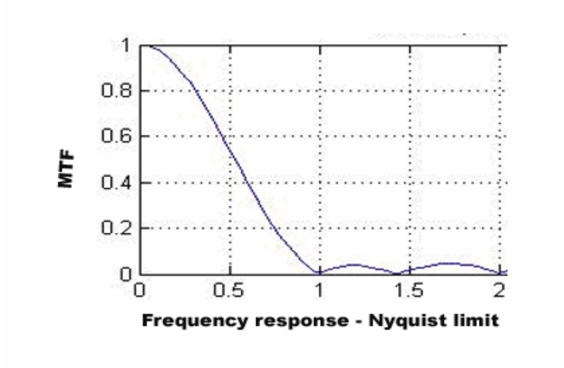

Figure 12. Origin Camera MTF

What this chart means is that if you image a resolution target, or any other highly structured object such as a fence, the system will lose its ability to distinguish between alternating lines as they approach a frequency of 2000 line pairs per picture width.

Fortunately there are other fine "organic" details that are of more interest to a DP such as hair or tree branches that can be accurately resolved without introducing artefacts. The method is counterintuitive in that the image is actually first "blurred" by an optical pre-filter placed immediately in front of the sensor which spreads the incoming optical signal across a known greater number of sampling elements (pixels). This enables the reconstruction algorithm to look far beyond its immediate neighbourhood to detect image content frequency and edge information and make intelligent decisions about reconstruction of the scene content. This is also made possible in part by the very high dynamic range and low noise performance (4DN in 16 bits) of the internal camera architecture.

With all of the above considerations factored in, an extremely complex multi-pass rendering process has been developed by DALSA to produce high quality, high resolution RGB images from the source RAW image data. The choice to perform reconstruction in offline software rather than in-camera video hardware allows for much better results to be achieved, albeit at the cost of processing power.

Given this is a critical function of the digital lab, DALSA is in the process of testing and certifying various facilities to ensure that regardless of where the image data is handled the initial rendering process will be of consistent quality from location to location. Once the RAW data has been rendered into 16 bit linear RGB files a number of other important image processing considerations come in to play.

Downstream Processing Considerations

Another important aspect to consider is all of the various sources of visual elements that will be used to produce the final piece. Reality is that for a significant portion of feature production source elements from digital capture, film capture, VFX plates and CGI will all need to coexist and intercut seamlessly.

Regardless of the source of the material Digital Intermediate (DI) workflows have quickly become the "standard" within the industry. However there are differing opinions regarding the suitability of using legacy 10 bit "log" space systems and new state of the art file based, linear image processing systems. This has lead to confusion and even controversy over what the differences are and which approach produces the best result. While this is a complex subject, there are a few simple concepts that illustrate the nature of the problem and point us in the right direction regarding decisions to be made both on-set and during post production to achieve the best possible results.

Regardless of the source, each image element has unique characteristics for representing the scene dynamic range, exposure latitude and color gamut. Additionally, every display media (film or electronic) has different dynamic range and color gamut characteristics. Given that each processing step remaps both the available dynamic range and color gamut to suit the given display device used at each step of the process, the final results are fundamentally affected by every single processing step performed throughout the chain. Therefore, any small error introduced during capture (i.e. mapping images to tape in HD systems) or early processing step will be multiplied in all subsequent steps.

In all cases, the precision of each operation is the common denominator that determines the quality of the end result and the most important distinguishing characteristic between log and linear processing is precision. Linear processing allows for a uniform distribution of precision within the available rage (10bit, 12bit, 16bit, 32 bit…) while log processing assigns more precision to certain areas of the range than others.

In log based systems subjective criteria (display/capture range characteristics relative to human perception range) are used to determine where precision is maximized and where trade-offs are made.

Figure 13 illustrates output bit mapping commonly used in today's DI process.

When Log values are assigned to the range early in the process it can have a negative impact on the final results even when the processing system uses state of the art 32/64bit per channel processing. This is because between the multitude of intermediate

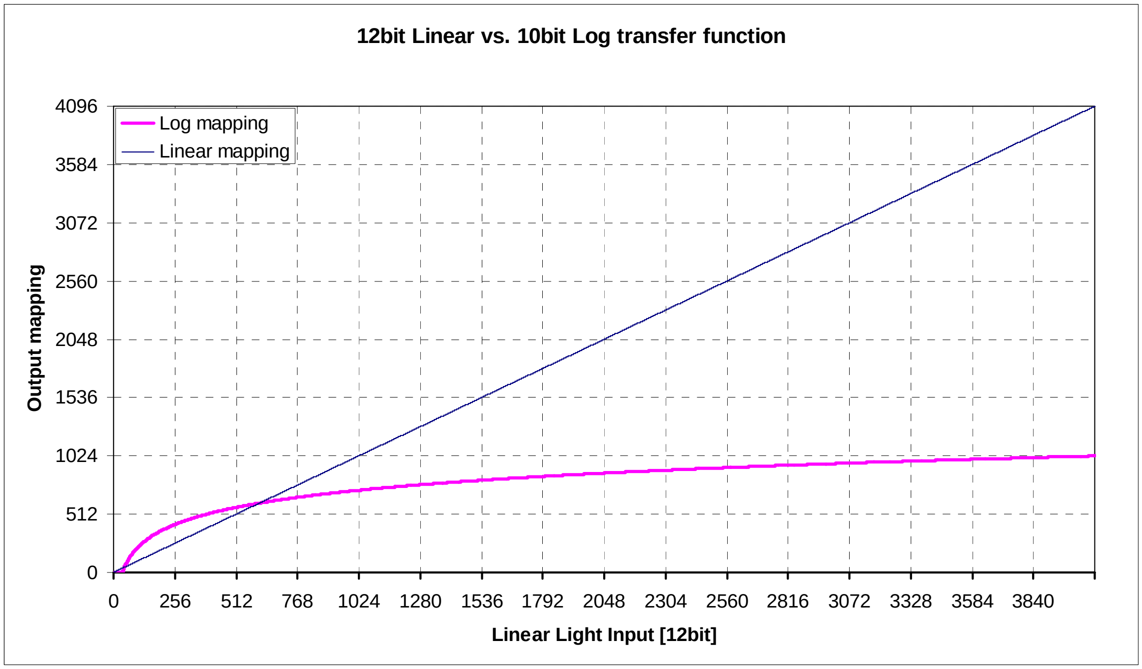

Figure 13. 12bit Linear vs. 10bit Log output mapping

steps the files need to be repeatedly mapped back into the 10bit log file envelope that is currently the standard. Log processing itself is not the problem but rather every time a file is saved within this bit envelope the precision is reduced due to rounding errors introduced in the float to integer number conversion. Figure 14 shows the relative error values in digital numbers (DN) of a 12 bit image mapped to 10 bit log space.

While the initial log mapping achieves visually lossless coding, each successive conversion step between the log file envelope and computers that perform linear mathematical operations will multiply the error and significantly reduce the quality of the output.

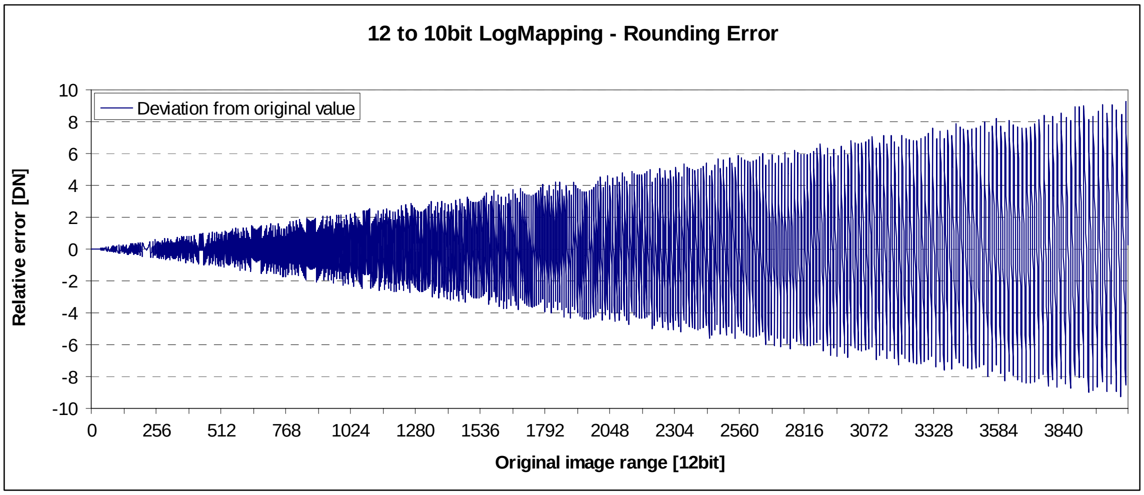

Figure 14. Rounding Error in 10 bit Log Mapping

High Quality Data Centric Workflow

A modern state-of-the-art workflow that is capable of taking advantage of today's advanced software and computing devices will look much like an enterprise class IT environment. It will be based on a data centric file based architecture that successfully integrates data management (i.e. metadata) with advanced image processing tools.

The primary objective is to implement a workflow design that allows integration of materials from multiple sources that are processed at different facilities by specialized teams working collaboratively and in parallel. Digital asset management tools that correlate metadata containers produced early in the production stages and then allow for expansion of the data set during downstream processing steps will not only ensure data integrity but allow collaborative data access and increase overall efficiency.

From a practical perspective the use of lower resolution proxy files is recommended throughout the off-line process to reduce the volume of data being processed. Metadata based conform using advanced EDL and CDL from the various off-line processes enables a high resolution master to be seamlessly produced in keeping with the original artistic intent.

Figure 15 is a block diagram overview of a 4K data centric workflow environment for producing high quality motion pictures based on existing hardware and software tools.

Using an HDR linear workflow philosophy deep into the post production chain ensures the highest quality end result. Regardless of the origination sources (film scans, electronic capture, CGI) the farther into the process the entire original image content is preserved the better element matching and overall quality that can be achieved. By utilizing soft LUT emulation tools for display mapping the original material is not altered until it is committed to the target presentation media. This minimizes the effort required to match multi-source elements by reducing the number of time consuming steps required to compensate for earlier processing differences and leaves more time for creative activities such as color grading.

Global Data Management Considerations

What is the intended primary distribution media…and how do I archive my original material? This section deals with a number of data management decisions that need to be considered in the context of distribution intent and budget constraints.

Depending on the specifics of each production the total amount of footage shot can vary widely from 20 hours (approx 120K ft) on a low budget feature to 250 hours (approx 1.5M ft) on a major multi-camera production loaded with VFX. Between the dailies review, editorial and VFX work only 3 to 10 hours of footage ever

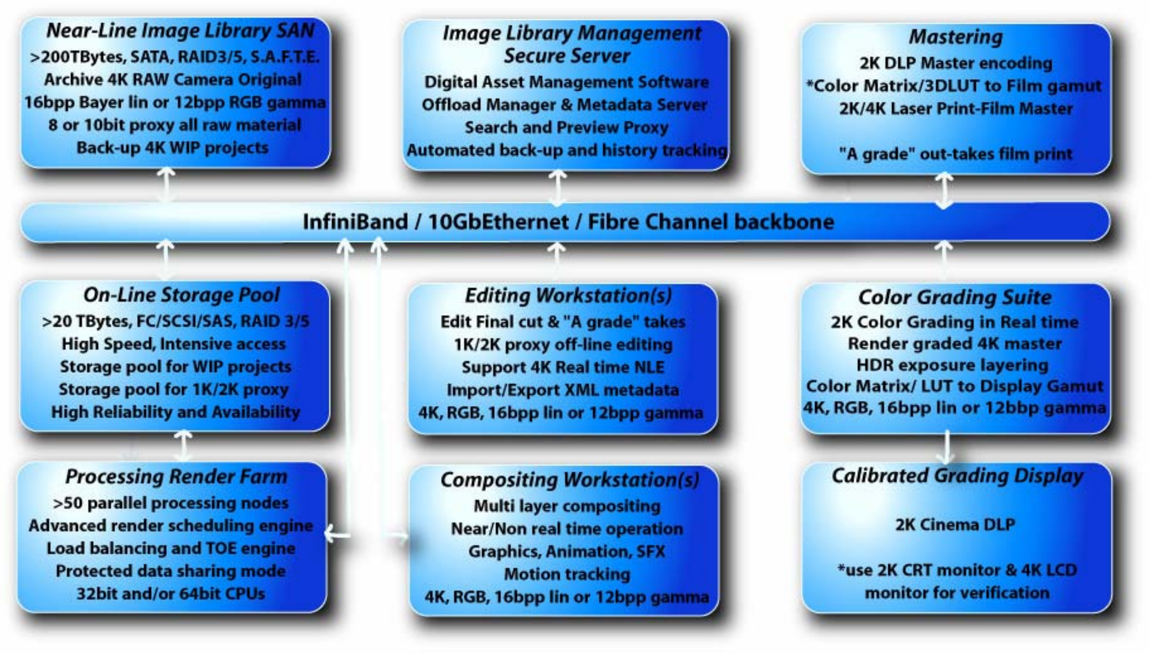

Figure 15. 4K Data Centric Workflow Environment

make it into the latter stages of post production for assembly and compositing. Once the key segments of the final cut are established only the final cut makes it into final color grading so it is important to make decisions about the appropriate resolution, bit depth and compression ratios for each of the elements as they work through the pipeline.

Table 1 illustrates the effect such decisions have on the amount of data storage required based on a large 1,000,000 foot production. The assumptions behind the figures presented are;

- “A Take” ratio is 5:1 (i.e. 10 hours for a 2 hour feature)
- “B Takes” equal all footage minus A Takes (approx 173 hours)
- “L^3” refers to mathematically lossless 2.5:1 data compression
- “L^3bounded” refers to “visually” lossless 8:1 compression
- “MPEG” refers to 10:1 “lossy” compression

The columns are arranged from left to right in order of lowest to highest “quality” wherein resolution and bit depth increase and compression ratios decrease as you move across the table. For example, the “Low” column enumerates the storage requirements in Tera-Bytes for each of the constituent elements based on the formats specified in each and shows that based on these decisions, the total data storage is almost a factor of 10 less than the same 1,000,000 “feet” treated as shown in the “High” column.

From a cost perspective it is important to consider the speed requirements for accessing various image elements and recognize that not all data needs to be stored in an on-line fashion. Current pricing HDD storage suitable for off-line usage has now dropped to the point were it is competitive with 35mm film. 1 hour of 4K RAW data requires approximately 1.5TB of disk space which today will cost roughly $3,000 complete with controllers and RAID cards. This compares to 1 hour of 35mm film negative at approximately 6,000ft that will cost the same $3,000 so from an origination perspective HDD storage is now competitive with film on an outright purchase basis. Depending on the business models that emerge HDD may in fact be cheaper if they are made available on a “rental” basis that models the temporal realities of the production cycle relative to the amortization of the hardware costs.

Other alternatives to HDD for storing material offline during post include the various linear data tape systems such as LTO.

Preservation of Original Source Material

Regardless of the mastering decision a decision also needs to be made about archiving the source 4K RAW files. The 4K RAW data makes an ideal data archive format as it enables future repurposing of the material. It is important to note that in addition to the image data, the compression algorithm used (ideally open source code with mathematical formulas describing the algorithm) as well as metadata required for color reconstruction be packaged together with the source material to ensure that changing technology will have no impact on long term material preservation.

For practical reasons, it is recommended that the recorded source material be graded for archival and preservation purposes along the lines of the table above. RAW files can then be compressed for archival depending on quality and budget considerations and distributed to multiple locations for disaster risk mitigation (consider what just happened to the various local film archives in Louisiana). It is also recommended that a laser output from the 4K master to 35mm negative be struck to complete the archival and preservation strategy.

Total Storage Requirements [Tbytes]	**Low**	**Medium**	**High**
Release print *	**2.6**	**6.44**	**8.59**
	4K, 12bit log RGB, L3	*4K, 12bit log RGB*	*4K, 16bit lin RGB*
A Takes ***(per above)***	**9.43**	**15.27**	**50.91**
	4K, 16bit lin RAW, L3	*4K, 12bit log RGB, L3*	*4K, 16bit lin RGB, L3*
B Takes ***(per above)***	**15.53**	**49.69**	**137.88**
	2k,10bit log L3 bound	*2k,10bit log L3*	*4K,16bit RAW, L3 bound*
Proxy ***(entire OCN)***	**0.86**	**1.54**	**2.52**
	8bit, 600x300MPEG	*8bit, 800x400MPEG*	*8bit, 1024x512MPEG*
Total Feature Film	**28.40**	**72.94**	**199.90**

Table 1. Data storage requirements

*This table is intended to illustrate the relative impact of image quality decisions on temporal storage requirements during post production phases and higher compression ratios are clearly possible for B Takes and all proxies. Also, it does not address the lower bits rates and storage requirements allowed under the DCI specified image container for distribution masters.

Summary

As motion picture production makes the transition to an all-digital process there are a number of important factors to consider. High dynamic range, high resolution cameras such as the DALSA Origin offer significant benefits relative to traditional 35mm or HD systems both in terms of quality and economics but the choice of camera cannot be made in isolation. The camera supplier has an obligation to understand the requirements of the entire pipeline and work collaboratively with production companies and their post production service providers to help ensure the best possible result within the budget constraints of the project. As technology inevitably marches forward new tools and efficiencies will be introduced to the pipeline but the following four fundamental considerations will remain intact:

- With no irreversible baked-in image processing, metadata rich High Dynamic Range (HDR) RAW image files are the best way to preserve flexibility and convey creative intent
- Maintaining data in a linear format deep into the pipeline enables the highest quality end result
- 4K RAW files are required for archival purposes and to future proof assets
- Data management decisions are required throughout the process

More Information

DALSA Digital Cinema.
6160 Variel Ave.
Woodland Hills, CA 91367
Tel: 818 884 7000
dc@dalsa.com
www.dalsa.com/dc