White Paper: 4K Digital Capture and Postproduction Workflow

Lucian Ion and Neil Humphrey

4K digital capture and postproduction are both attractive and viable. With support from IT infrastructure leaders, 4K digital intermediate (DI) processes can bring new levels of power, possibility, and creative control. This paper outlines an advanced digital infrastructure to support 4K DI and exploit the power of DALSA's Origin camera.

Introduction: 4k Digital Capture

There is no question that digital capture is an attractive option for motion pictures. Digital capture can offer cinematographers more power, more possibilities, and more creative control in telling their stories. Not only can digital capture offer immediate feedback, the best digital capture can also offer unprecedented image fidelity.

DALSA's Origin® camera is the first and only digital camera to offer performance truly worthy of the big screen, but the camera is just the first link in a 4K digital chain. To exploit the Origin camera's performance, the motion picture industry needs viable options for every link. The chain is data-centric, based on digital data rather than film or videotape. A 4K workflow requires a sophisticated digital infrastructure, but the components for that infrastructure are already available. Leaders in IT equipment for the media industry already offer solutions to bring 4K digital workflows from design into reality.

4K Digital Production Workflow

Digital production follows the same basic flow as film production: Photography, processing, review (e.g. "dailies"), and material prep for post. Figure 1 outlines the flow. Perhaps the biggest difference with digital production is the possibility of immediate feedback. Set preparation becomes real-time interactive, making it possible to test multiple scene looks. Cinematographers can review every take immediately, allowing them to immediately make the decision to reshoot or strike the set. Background frames can even be layered onto the set "monitor" output for real-time preview during shooting. The "dailies" concept of reviewing material after some basic editing and arranging is still valuable, but a digital process removes the need for film processing, telecine, and their associated logistics, and makes dailies fully interactive. A true digital process on the set can provide a multitude of new tools for the cinematographer, from basic compositing of background templates to interactive synchronization of complex camera motion and lens setting to computer generated graphics for today's SFX shots

Figure 1. Digital Production Workflow

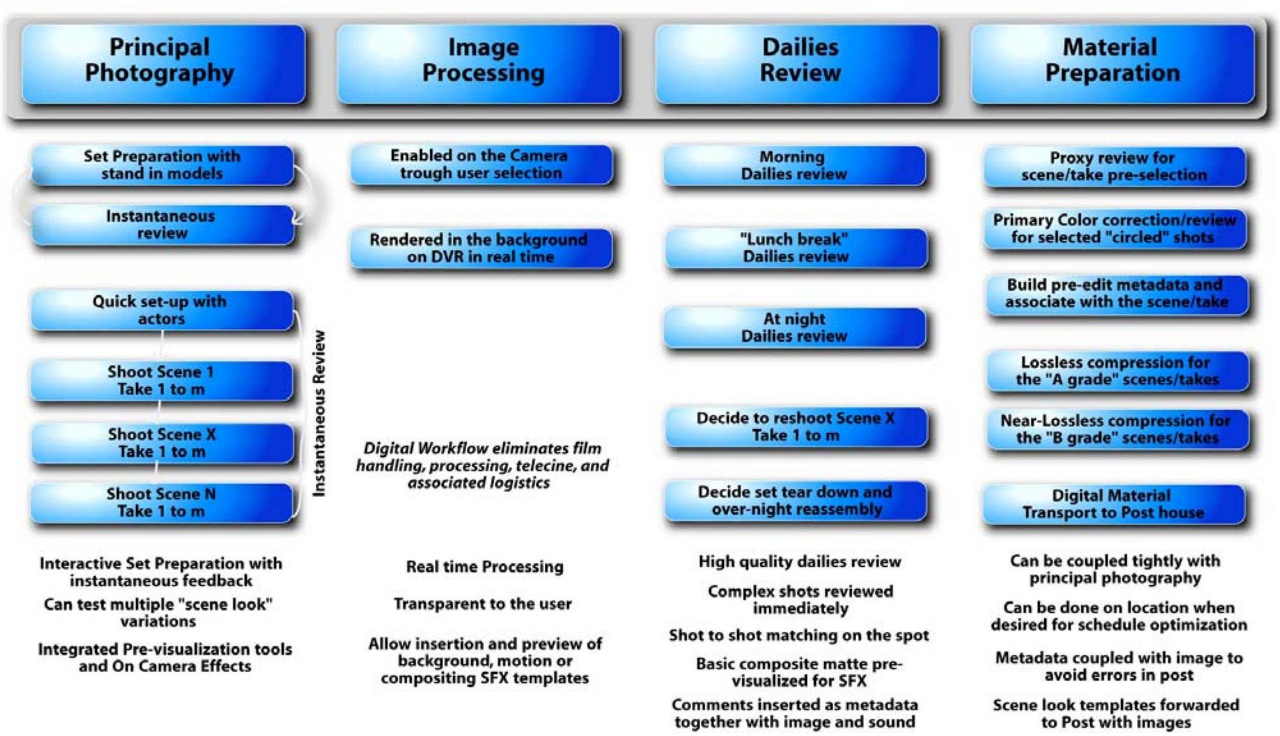

Bayer Image Data: the digital version of film negative

To allow the use of standard cine lenses, the Origin camera captures images with a single sensor chip and a Bayer pattern color filter (GRGB) like those used by the best professional digital still cameras. RGB values for each pixel are interpolated from the raw sensor data.

To take best advantage of the sensor's huge exposure latitude (at least 12 stops), Origin outputs 16-bit linear values in its primary data mode. Origin also offers 12-bit power law-encoded (gamma) data, but 16-bit linear offers the highest fidelity. Combined with the 4k x 2k resolution, each raw frame of 16-bit data is at least 16MB. At 24fps, that represents more than 400MBytes per second (1.44TBytes/hr). For in-camera interpolated 16-bit RGB, the figure rises past 48MB per file, or 1.2GBytes/sec.

One convenient way to reduce the bandwidth and recorder storage requirements is to record Bayer data from the camera instead of full RGB. Set monitors can be driven from Origin's 1K DVI feed or its HD-compatible output in RGB or YUV formats, but the main 4K data stream can be output to the recorder either as Bayer data or as interpolated RGB frames. The Bayer output can be considered as a form of lossless compression–it requires only one third of the bandwidth and storage capacity without compromising image quality. Bayer data acts as "digital negative" and makes an ideal original archive format. As we will see later in this paper, interpolation isn't necessary until each sequence goes through an editing, compositing, or color grading process in postproduction.

Metadata

One of the challenges of any movie production, film or digital, is managing the raw footage. Traditional filmmaking labels its film canister and uses paper camera reports to capture vital identifying information to pass on to the developing lab, post production editors and color graders. With a digital workflow it is possible to attach this information directly to each sequence and even embed it in-camera into every frame as metadata. Timecode information is the minimum, but a digital data workflow allows almost unlimited metadata, including lens and exposure details, scene and take, crew, location, and even the cinematographer's on-set commentary and instructions as audio clips. As part of the image file, this metadata is available all the way through the post process for use in managing digital assets and clarifying creative decisions. And compared to the volume of image data in a 4K frame, even extensive metadata is almost inconsequential in terms of additional bandwidth required.

Disk Capacity and I/O

At 400MBytes/sec., 60 minutes of footage would occupy 1.44 Terabytes. Origin's recorder must offer both tremendous capacity and tremendous I/O performance. Disk capacity has tended to double every 9 to 12 months, but the real limiter is disk I/O speed, something that has improved only slightly every year. Bandwidth requirements dictate a parallel disk recording configuration. Current Origin recorders use a RAID0 configuration while the new generation recorders will be configured as RAID3 or RAID5 to provide redundancy. Current HDD technology insures great reliability (>1,000,000hours MTBF), and advanced recorder designs with RAID3/5 configurations ensure that the valuable recorded information is much safer than on analog media.

In order to support the I/O bandwidth required for 4K data streams, Origin implements a 4xInfiniBand protocol over a single highly reliable fiber optic connection. InfiniBand is a well established networking technology that is well positioned to take full advantage of the emerging PCI-Express architecture that will re-define the IT word in the years to come. The most attractive features of InfiniBand are high QoS (quality of service), low latency and switched fabric architecture. Sustained data rates in excess of 850MBytes/sec have been widely demonstrated and ensure low-latency image data delivery.

In a studio setting, camera output could be streamed directly into a storage network, but when "on location" the camera needs a field recorder which must be offloaded when full (much like film magazines). Current recording units contain 1.2TBytes of storage, enough for approximately 50 minutes of shooting in 4K Bayer format. Table 1 compares storage with film footage in several scenarios.

Table 1. Film footage vs. 4K data storage

	Small project	Larger project	Intensive, multi-camera
Footage per day	6,000	10,000	25,000
Frames per day	96,000	160,000	400,000
Storage per day (16bpp raw Bayer)	1.59TB	2.65TB	6.63TB
Shooting days	17	50	60
Total Footage	100,000	500,000	1,500,000
Total Storage	27TB	133.5TB	398TB

Obviously, a high-performance storage network and a robust, scalable file system are vital. Fast and easy offloading is also important, and new generation recorders will allow real-time offload. Metadata and lower resolution proxies (ex. 1K dailies) can also be recorded on disk with the main image stream to simplify the postproduction process and provide better information for post decisions.

Postproduction Infrastructure

Postproduction includes editing, compositing, animation, special effects, color grading, mastering, and archiving (see Figure 2). To support all the functions of a full digital intermediate (DI) process, it goes without saying that a DI facility's network must be built on an advanced high-bandwidth backbone infrastructure. Technologies like InfiniBand, Fibre Channel and more recently 10GbEthernet are available to address these requirements. Each of these protocols offer stitched fabric for flexibility, high QoS, and guaranteed bandwidth on demand.

SAN: Near-Line Image Library

To handle the tremendous amounts of data produced by a DI process, advanced SAN structures are required. A near-line SAN based on low-cost SATA technology can optimize the process economically. Original camera material, animation, graphics, proxies and metadata can be stored in a near-archive mode in a compressed form and can be made available quickly on demand. The speed and cost benefits of using high quality proxies and metadata to sort, qualify and provide a rough cut of the available material are considerable.

"Work in progress," in which post applications are actively using image data, requires the highest possible performance, and true online SAN structures are necessary, built with the most advanced disk technologies, such as Fibre Channel, SAS, and SCSI.

Image Library Management

Beyond merely storing data, a reliable digital infrastructure must manage the data. The complex, often heterogeneous nature of media networks and the intermingled, even chaotic succession of editing, compositing, and color grading processes make it vital to carefully manage both the archived original digital assets and the evolving work-in-progress data. A secure image library management server should broker and log all requests and uses of assets from the near-line storage of the SAN to the various post applications and workstations. The library management server's functionality is very similar to the database management systems used widely in other high-volume IT networks for the past decade, such as transaction processing or claims management.

On the library management server, digital asset management applications provide check-in/check-out, versioning, search and preview functionality, as well as automatic backups. This server governs file permissions and mediates shared access, storing metadata about the assets in the image library. But because actual image data is not stored on it or transferred through it, it is not a data bottleneck.

Processing Render Farm

Render farms—collections of parallel processing nodes for handling the most demanding processing tasks in compression/decompression, color interpolation, conforming, color grading conform rendering, compositing, etc.—are standard elements in media networks. For a 4K digital infrastructure, the render farm can be developed in a scalable fashion as projects demand to contain tens to hundreds of processing

Figure 2. Digital Postproduction Infrastructure

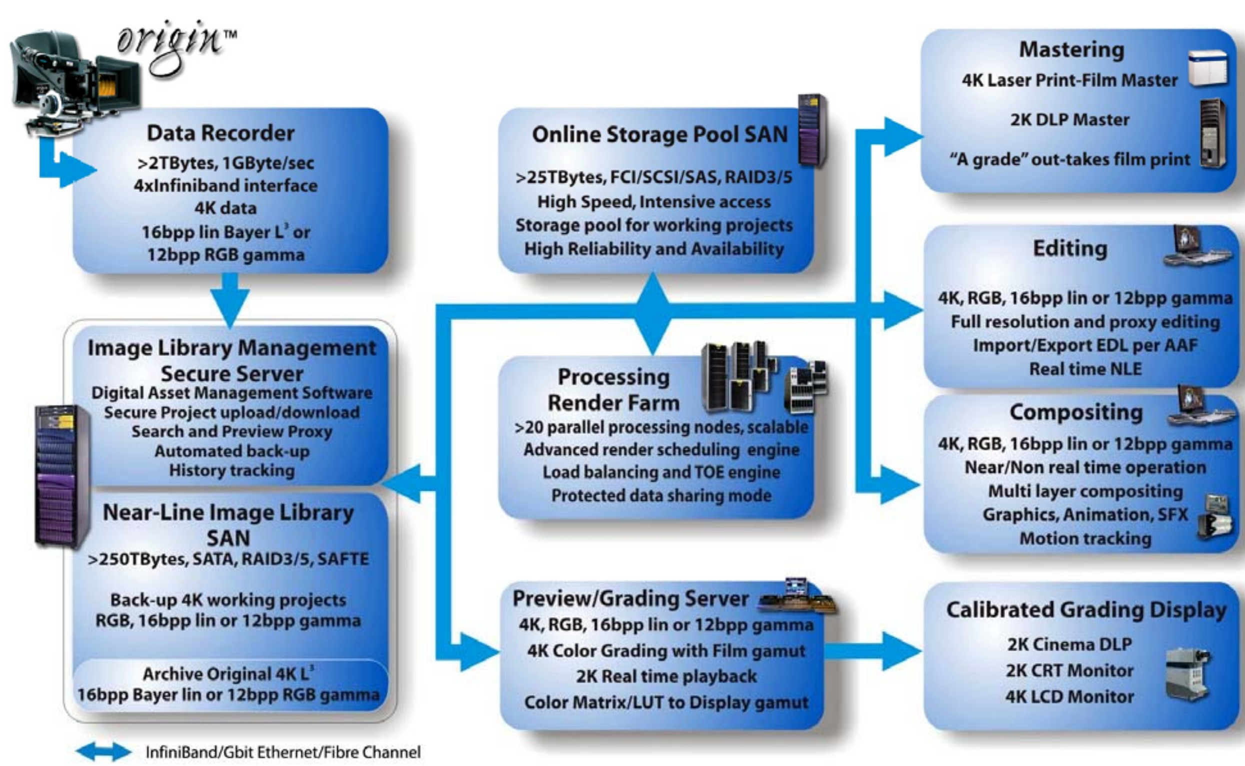

nodes. The farm can be built as an HPC cluster using mainstream 32-bit processor structures or it can use a more advanced configuration based on the newest 64bit processing engines. One powerful example for an advanced computing engine is SGI's® Altix™ system. Coupled to a heterogeneous infrastructure, it can provide high value performance right out of the box.

Editing/Compositing

Rendering 4K frames with 16-bit linear data requires considerable horsepower. A typical feature film production generates a very large amount of "original negative" to create a relatively small final cut. Shooting ratios of 3:1 for a modest production to a 30:1 for a high budget release are common. While there is little value in processing all the material at 4K, it is very important to preserve the quality of the original. Assistant editors can perform original material sorting based on circled take information, on set metadata, and proxy review (1K proxy full color can provide very high quality review). A detailed edit decision list (EDL) can be composed during the "Original material management" stage (see Figure 3); that EDL will provide an input to the full resolution editing/compositing suites as well as to the render farm for timely processing. Rendering and network management software will be running in the background, transparent to the user to manage the rendering queue and set project access rights for concurrent workstation access to material. As the production requirements for original material are developing during the postproduction process, the EDL can be updated in real time and the software can interactively retrieve the material from the offline SAN.

The first time an editing station retrieves Origin data from the library, that data will still be in Bayer format and may also be compressed with DALSA's L^3 algorithm. In this form the full 4K image files are not much bigger than uncompressed 2K RGB. On their way to the editing station they can be decompressed and interpolated on the network's render farm for manipulation as 16-bit linear data. (This could be done in software or with the assistance of hardware accelerators.) Most post applications already perform most of their operations in 16-bit linear space, although until now they have expected to load the data from 10-bit log files (e.g. Cineon/DPX).

Figure 3. Postproduction Workflow

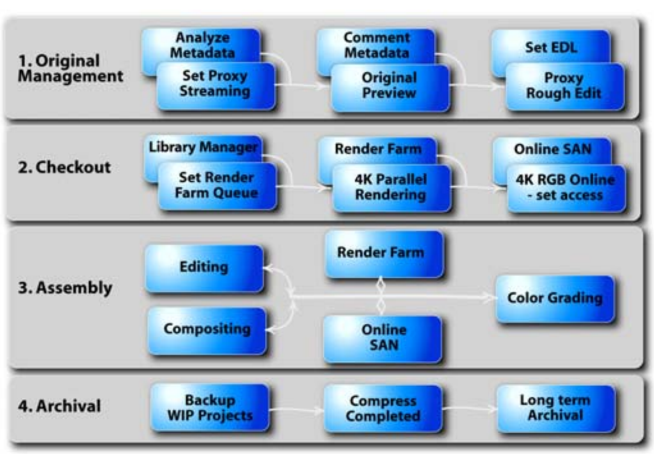

Bit Depth and File Formats

It is clear that the industry is outgrowing 10-bit log formats. The steps between 10-bit log values are quite visible to most observers viewing a high-quality digital projection, especially at the bright end of the range. A high quality digital imaging system should preserve an exposure latitude exceeding 12 stops. The bit distribution from black to extreme highlights must to match that goal with the needs and limitations of the selected display media to take best advantage of the characteristics of the human vision system. 12-bit gamma encoding provides steps smaller than most observers can detect in most situations, but any logarithmic or power-law encoding of linear data is necessarily both lossy and irreversible. To carry the maximum image fidelity all the way through the post process, data would ideally remain in a linear format as long as possible with the highest bit depth possible. This is why the 4K DI process requires so much more "heavy lifting": it is not just the number of pixels, but the bit depth of each pixel.

The Origin camera outputs 4K data in 16-bit linear format (Bayer data or 4:4:4 RGB), or 12-bit 4:4:4 RGB gamma encoded. The 12-bit gamma data can be mapped internally in the camera from the 16-bit linear data through application of a soft clip and gamma curve. The 16-bit linear data allows more flexibility for scenes with wide dynamics; it also allows the DP to precisely trim the desired look after the shoot. Metadata can be attached to each frame to preserve those preferences through post-production.

It should be noted that the DPX framework does not preclude 16-bit data, but 16-bit linear implementations have not been established yet. Whatever the file format selected, it must be flexible, allowing the insertion of a variety of metadata at different stages of the DI process. Another option is to preserve metadata information attached to the file sequence into a descriptive format like XML that provides unlimited flexibility and wide compatibility with IT software. It must be efficient and easily shared across operating and file systems among many post applications. A number of open file formats are well-suited to this process, including SGI's 16-bit file formats, 16-bit and 12-bit DPX formats (viable, though lacking industry support), as well as Open EXR, which provides even more flexibility in data packaging.

Online Storage Pool

The image library SAN is the ultimate repository for all the digital assets, but for the most intensive, highest-data-rate access, the network should include an online storage pool of perhaps 25TB (as opposed to the SAN's near-line 250TB) built for maximum availability, reliability, and performance. Active projects and "work in progress" data should live here.

Color Grading and Preview

As with editing/compositing, a color grading or preview application should be able to render original archived material that may still be in Bayer format, possibly also with L^3 compression. Again, render farm power and hardware accelerators can boost performance. Several advanced solutions already exist that allow high performance color grading at lower resolution (1K, 2K…) with precise conform rendering at 4K.

Each distribution channel (laser film writers, DLPs, monitors) has a different color space or gamut, so tuning colors for the intended artistic effect is not a single process but a series of processes. The grading process must map the colors in the original to the appropriate colors on the exhibit screen. The essential element is a calibrated viewing screen that emulates what the content will look like on the actual exhibit screen. For example, the final color grading on a digital release of a feature will be different than the digital data heading for a film writer to create the film release negative. A high quality digital display device (e.g. Cinema DLP) is a gate towards insuring consistent results and making the best use of the information available from the "negatives."

Conclusion

A 4K digital workflow is both possible and worthwhile. It is not futuristic—solutions exist now and are rapidly evolving and improving to meet the industry's demands for higher image fidelity, more power, more possibility, and more creative control. These ideals are independent of any particular technology—film or digital—but a pipeline from Origin's 4K capture through a new 4K workflow provides a path to the highest quality cinematic experience yet.

DALSA Corp.
605 McMurray Rd.
Waterloo Ontario, Canada
N2V 2E9
dc@dalsa.com
www.dalsa.com/dc

More Information

For more information on digital infrastructures and workflows, consider these sources:

www.sgi.com

- White Paper: Digital Infrastructure Solution for Production, J. Farney, SGI, 2003
- White Paper: E-cinema – The Next Digital Revolution, Louise R. Ledeen, SGI, 2003
- White Paper: A Data-Centric Approach to Cinema Mastering, Thomas. J. True, SGI, 2003
- White Paper: Shared Filesystems for Collaborative Production, Thomas. J. True, SGI, 2002

www.discreet.com

- White Paper: Digital Intermediates, Discreet, 2002
  - http://www.discreet.com/support/documentation/whitepapers.html

www.quantel.com

- http://www.quantel.com/domisphere/infopool.nsf/html/Papers