Chapter 1. Color Correction Workflows

Click here to gain access get your media files.

In order to start grading, you need to get the project you've been hired to work on into the application you're going to be working with. As unexciting a topic as this may seem, understanding how it works is vital, so this is where we shall begin. This chapter isn't meant to provide an exhaustive overview of postproduction workflows, nor is it intended to cover every format you can expect to have to deal with as a colorist. Instead, it's meant to give you a better window into how the colorist fits into the postproduction process and what decisions get made that will affect your well-being.

This chapter has been written with the medium-sized boutique facility in mind. Colorist/finishers working in smaller outfits will undoubtedly have much more on their plate, while colorists at top-level facilities get to ignore a lot of this and focus on the art of the grade. I'm guessing you're probably somewhere in the middle, so if you have any influence at all over how the footage you'll grade will be shot, or how the workflow that the post supervisor organizes will go down, read on.

Are You Grading for Cinema, Broadcast, or the Web?

This is a trick question, because it doesn't matter. Whether you're working on a movie that's going to 300 theaters, a network or cable television episode that's going to broadcast, or a science-fiction web series that's going to YouTube or Vimeo, your program needs to be graded. If you're treating your project in a professional manner, then the workflows are similar if not *identical*. This goes double if your client says, "I think my program actually looks fine as it is, I don't know what really needs grading...."

It's a misconception that color correction or color grading (whichever you want to call it) is primarily about fixing problems and dealing with broadcast-safe issues. The notion that the colorist's job is curative-only assumes that if your video is well-shot and headed straight to video/late-night cable/wherever-on-the-Web, then color grading doesn't matter much and you can skip it. The truth is, while fixing problems is a meaningful part of what the color correction process is all about, that's not actually the main reason you want to incorporate color grading into your workflow.

You grade projects because *you want them to look as good as they can*. You grade them because you want to emphasize and preserve image detail that's important and to lend a sense of style and occasion when necessary to finesse the project's look. And this is something you need to do whether your program was shot on a RED Dragon, an ARRI Alexa, a 5D Mark III DSLR, or a GoPro taped to someone's crash helmet.

Where Colorists Fit In

Color correction, as I'm fond of saying, ideally starts in preproduction. It doesn't have to wait until the very end of the postproduction process (although it all too frequently does). You, as the colorist, can meaningfully contribute to the overall process of preproduction, production, and postproduction in a variety of ways.

- Preproduction: At this stage, you can recommend shooting formats based on the postproduction needs for a particular project, work with the cinematographer and director during camera tests to evaluate how differently graded looks hold up with different cameras and shooting conditions, and turn these looks into evaluation lookup tables (LUTs) that can be loaded into on-set production displays and used by the digital imaging technician (DIT) during the creation of dailies.
- Production: DITs incorporate many skills of the colorist into their job description as they help the

cinematographer to evaluate the images being shot, manage the monitored image in video village, create on-set grades, and generate digital dailies for use in postproduction. On union shoots, colorists who aren't in the camera union cannot participate directly in this process on-set, but communication between a DIT and the colorist is often essential, especially in situations where grading data, whether LUTs, CDL, or other presaved grading formats, is being exchanged along with the recorded media.

- Editing: Prior to editing, productions using raw camera formats or highly compressed long-GOP or H.264-compressed media must have this potentially unwieldy camera original media turned into a set of easily editable QuickTime or MXF media files. This process is usually referred to as creating "digital dailies" and may involve syncing audio and applying or creating grades for the media so the editor and director aren't stuck looking at low-contrast, ungraded media for the duration of postproduction. Sometimes this work is done on-set by a DIT; sometimes it's done by a colorist back at the suite. As an edit nears completion, colorists are sometimes called upon to do a quick "offline color" pass for projects that will be undergoing test screenings. Sometimes this color pass is the beginning of the final grade, sometimes not.
- Visual Effects (VFX): During the process of creating visual effects, colorists are often called upon to help grade and match greenscreen elements and to grade (and regrade, and regrade again) finished VFX being incorporated back into the edit. There's usually a lot of back and forth between color and VFX when a program is undergoing final finish.
- Grading: For the colorist, this is the main event; once the edit has been locked (one hopes) with all final media and VFX elements (good luck with that), the edited timeline using the digital dailies is *reconformed* to use the camera original media instead. This is done within the grading application being used, usually by an assistant at a large facility or by the colorist at a tiny one. At that point, the colorist takes control of the timeline to adjust the color, contrast, look, and feel of every clip in the project.
- Finishing: Sometimes, some or all of the finishing process—finalizing titles; adding header elements such as the slate, bars, and tone; making last-minute edits; laying in the audio mix; taking care of VFX-like tasks such as light digital paint, compositing, or simple blurring to remove unlicensed materials from the image—falls to the colorist, depending on the size of your facility and on the skill set of the other folks you work with. Top-level colorists usually focus *only* on the grade, while colorists at smaller boutiques are often invited to a broader level of participation in these sorts of activities.
- Mastering: Increasingly, the tools for final mastering of a finished project are being incorporated into the very grading software that colorists use, whether the program is being output to tape, rendered and copied to SSD as a digital master, or exported as a Digital Cinema Package (DCP) for cinema distribution. Again, dedicated colorists aren't usually involved with this process, but colorists at smaller facilities may very well be.

These are all the aspects of preproduction and postproduction you may be expected to weigh in on, depending on the budget and the workflow of the project you're participating in. Of course, on most projects the post supervisor will simply wait until they're almost finished with the edit and then give you a call, in which case you get what you get.

Ideally, they call you early enough in the editing process so that you can talk them through what will be necessary to conform the project they're editing in their NLE of choice to your grading application of choice. Oftentimes, small decisions made early on can make all the difference between a five-minute conform and days of tedious work.

Before the Shoot: Choosing a Recording Format

If anyone actually asks you, the colorist, what format would be best to shoot for the best possible experience in post, then you're a lucky, lucky person. Unfortunately, the answer is not as simple as "the highest-quality format you can afford," although that is certainly a rational first response. In truth, there are a variety of ways to obtain high-quality media that's suitable for grading, and the right one depends on the budget, the schedule, and the style of shooting a production will employ.

I'm making the deliberate decision not to talk about specific cameras in this section. For you as the colorist, the data format of the recorded media is much more important than the particular camera being used, although one would hope the cinematographer chooses good glass (lenses) to shoot with. Moreover, digital cameras are such a moving target that it's pointless to discuss them in print; new models that have surpassed or succeeded any that I'd recommend would already be available by the time you read this. However, the media formats these cameras record to don't change nearly as quickly, and there are certain characteristics that distinguish one format from another that you should be familiar with.

Film

If anyone ever asks you whether they should shoot 35mm film and their budget is sufficient to the task, you should probably just say yes. While this book is primarily occupied with the process of digital grading workflows and techniques, film is easily (if not cheaply) made into digital media through the use of a film scanner, which ingests each reel of negative or reversal film frame by frame and produces 2K or 4K DPX image sequences with reel information and timecode based on the frame numbers that were scanned. Scanned film DPX sequences have lots of latitude for later adjustment, can be converted into more easily edited offline or online-quality formats using the same techniques discussed in this chapter, and can be either graded and finished via online-quality transcodes or matched back to the finally edited sequence for grading and finishing at maximum quality.

Film scanning workflows will continue to be with us long after film is no longer used as a regular acquisition format because of the necessity of digitally converting material from film archives around the world, either to remaster older programs or to incorporate archival media into contemporary projects. If you're interested in digital intermediate film workflows, consider reading Jack James' *Digital Intermediates for Film and Video* (Elsevier, 2006) to get started.

Raw vs. "Mastering-Quality" Codecs

When shooting digitally, one of the most basic decisions any production needs to make is whether to record to a raw format such as RED raw, ARRI raw, or CinemaDNG, or to a more highly compressed "mastering quality" format such as QuickTime or MXF (discussed in the next section). Most digital cinema cameras have the ability to do both, and increasingly smaller cameras are being made to record raw video data as well, such as those made by Blackmagic Design, and others thanks to third-party camera modifications available via the open source Magic Lantern project.

Raw camera formats record linear light data directly from the sensor to a file. Some use compression to make these potentially huge files more manageable in post (RED raw is a notable example); others don't. Since digital cinema cameras are so named because they use a single large-format image sensor—typically the equivalent of a super 35mm or super 16mm film aperture—the main thing you need to understand about raw media is that it must be *debayered* or *demosaiced* in order to produce an image that you can work with and watch.

The advantage of media in a raw format is that, since raw format records everything the camera sensor sees, it

provides the colorist with the maximum amount of image data to work with in postproduction. Typically, any ISO or aperture settings you make with a camera shooting raw affect only how that image is monitored on the set and what default camera metadata is written to each raw file, rather than modifying the image data that's being recorded. Changes made to the image being monitored from the camera will certainly change how the cinematographer is lighting the set, but the fact remains that you can change this metadata after the fact as you begin grading raw clips. This is a fantastic amount of flexibility.

One other advantage is that raw media formats that are also compressed (meaning that the color space of the media is raw even if the data format is not) can be quite small, easing storage space and bandwidth requirements necessary for working on a project.

However, raw formats have disadvantages, depending on your point of view. They are usually difficult or impossible to edit directly, necessitating conversion to a different format in a second step and possibly a reconform later when the time comes for grading, adding time and complexity to the postproduction process. Furthermore, recording in raw formats produces large amounts of data that need storage and backup. Perhaps the biggest liability of raw formats is that they require a certain amount of postproduction know-how to incorporate smoothly into a given workflow, and the benefits may not seem immediately worthwhile. For all of these reasons, cameras usually provide the option to record to other formats.

Quicktime or MXF Media

Alternately, most digital cinema cameras have the ability to record either QuickTime ProRes (usually either ProRes 422 [HQ] or ProRes 4444) or MXF (typically using DNxHD) video media. Furthermore, even if your particular camera is constrained in the media formats it supports, many productions elect to output uncompressed video data from a camera's HD-SDI or HDMI connection to an external digital recorder of some kind, which in turn records QuickTime or MXF media.

Either way, the advantage of recording to one of these formats, instead of raw, is that it dramatically simplifies your postproduction pipeline since you can copy the QuickTime or MXF media straight from the camera into your editor's system, and they can start cutting immediately. Furthermore, when they're finished, they can easily hand the resulting project and its media off to grading and finishing without the need to reconform to media in another format.

These workflow advantages are significant, but there is an additional consideration. When recording to QuickTime or MXF, you usually have the option (depending on the camera) to shoot either log-encoded media or Rec.709 (BT.709) media. Log-encoding is explained in more detail later in this chapter, but to simplify for now, the main difference is in how the contrast of the image is recorded, which is important.

Log-encoded media compresses the contrast of the image captured by the sensor in order to preserve the greatest amount of latitude for adjustment within the available bit-depth of these formats (either 10-bit or 12-bit depending on the codec you use). This gives the colorist the greatest amount of image data to work with when it comes time to do the grade. Even though log-encoded media looks strange when uncorrected, this is a good thing, although it means that the editor is going to have to apply either a correction or a filter, or they will have to enable a monitoring setting if they want to see what the media actually looks like straight out of the camera (as opposed to working with a duplicate set of normalized dailies that have cloned timecode and reel information). This isn't really a big deal but needs to be factored in. Friends don't let friends edit log-encoded media without correction.

The other option—recording QuickTime or DNxHD media as normal-looking Rec.709 video—is guaranteed to make your colorist cranky. While Rec.709 is drop-dead simple to monitor, simple to work with (no corrections necessary), and simple to conceptualize in terms of workflow, the reduction in available

latitude for adjustment between log-encoded and normalized Rec.709 media is *significant*. If you're shooting magazine show segments that need to be rushed to air, then the advantages of recording Rec.709 video probably outweigh the disadvantages. However, if you're shooting a music video or narrative production that would benefit from detailed grading, then shooting straight to Rec.709 would be doing the project, and the colorist, a disservice. And if you're a shooter and you do this without asking anyone, it just might get you fired.

Note

While camera menus typically present this option as Rec.709 (as we'll refer to it in the next paragraph as we discuss menu options), this book adheres to the naming convention of BT.709.

H.264 Media

On the opposite end of the scale from digital cinema cameras are ENG-style cameras, DSLRs, and crash-cams that record highly compressed H.264 media. However, be aware that not all H.264 formats are created equally. The H.264 standard can use any one of a variety of standardized *profiles*, each of which uses a higher or lower level of compression and ratio of chroma subsampling. This affects the data rate of the files that are recorded, resulting either in higher-quality files that are larger or in lower-quality files that are smaller. Furthermore, each profile can be encoded at one of a number of *levels*, meaning that any given profile can be throttled up or down to fine-tune the ratio of quality to size used to encode your media.

Practically speaking, different cameras record video using different profiles at different levels. The combination used by your camera will affect the quality of the media that's recorded in addition to the quality of that camera's lens system, digital sensor, and image processor. Taking this into consideration, your selection of camera has a big impact on your final visual result, as well as how much latitude for grading the recorded media will have.

Colorists in general tend to have a really bad attitude about H.264 video, and for good reason. The deadly combination of high compression and limited chroma subsampling (explained in the next section) usually means that the latitude for adjustment of DSLR media is significantly more limited than that of cameras shooting raw, QuickTime, or MXF media. And woe to the client who asks a colorist to match GoPro footage to their camera raw ARRI Alexa media (and they will). There's only so much you can do.

However, the prudent colorist will bear in mind that a job is a job, and while these media formats may be abominable from an imaging purist's point of view, the cameras that shoot these formats tend to be small, light, inexpensive, documentary- and low budget-friendly, and in some cases they are useful for situations in which no other camera would possibly work, such as the aforementioned GoPro cameras attached to every crazy place you can imagine. There are a lot of projects that simply wouldn't exist were it not for these cameras, and that's worth keeping in mind.

To provide a gentle attitude adjustment, I'll share some personal history. I came up as a colorist during the transition from analog Beta SP to DigiBeta, in the time when the industry first embraced DV-25 digital video as a production format. I did the vast majority of my early grading work for clients shooting DV-25, with latitude that was positively abysmal when compared even to today's H.264 media formats. And grade it I did, as hard and as detailed as I could. Sure, you can't do miracles, and any significant adjustment you make to the image will likely introduce more noise than you'd wish, but there's a lot you can still do if you stop complaining about it and use the tools. Tough love, my friends.

Of course, if anyone asks you, tell them to shoot raw.

Note

As mentioned earlier, the Magic Lantern project is an open source software hack for modifying various DSLR cameras to shoot raw. While I neither advocate nor discourage its use, it's worth pointing out that raw media recording is expanding to ever smaller cameras.

Understanding Chroma Subsampling

To enable recording to ever-smaller storage devices, different video formats discard varying proportions of chroma information from the signal, and this also affects how much you can stretch contrast before introducing noise. As with the difference between log-encoded and Rec.709 media, you *want* media that's recorded with the most chroma you can get, but what you get depends on the camera's capabilities and the media format being recorded.

To put things into perspective, 4:4:4 chroma-sampled media stores 100 percent of the chroma information and thus has an impressive amount of latitude for exposure correction. This allows the colorist to lighten dark shots aggressively before artifacts such as excessive noise become a problem. Digital cinema cameras shooting raw typically shoot 4:4:4 data, as do cameras that shoot ProRes 4444 and DNxHD 444. Most low-cost cameras do not (unless they're tethered to an external recorder).

The next step down is the 4:2:2 chroma-sampled media that's typical of high-end HD camcorders. Media encoded at 4:2:2 chroma subsampling has a fair amount of latitude within which a colorist can adjust contrast by a decent amount before noise becomes a problem. ProRes 422 and most other DNxHD formats are both 4:2:2 chroma-subsampled formats, which are considered suitable for broadcast. It's also worth mentioning that while there is an H.264 profile that utilizes 4:2:2 chroma-subsampling, few devices use it.

The majority of consumer-level and DSLR cameras that record to H.264-based video formats encode 4:2:0 chroma-sampled media. This discards three-quarters of the chroma data in a manner considered to be perceptually indistinguishable from the original, in an effort to shrink file sizes to create media that's more manageable in low-cost workflows. While in many cases 4:2:0-subsampled source media is considered suitable for professional work, the discarded color information makes it difficult to make significant adjustments to contrast without introducing noise. This type of chroma subsampling can also make various types of visual effects work, such as greenscreen compositing, more challenging to accomplish.

However, for many types of programs, the advantages in cost and ease of use far outweigh the disadvantages, and it is to your DP's advantage to shoot such highly compressed formats with the best lighting and exposure possible in order to make the most of the format's limited bandwidth and to minimize the necessity of the colorist having to make difficult decisions later.

Preemptively Upconverting 4:2:0 Media is Often a Waste of Time

If your video was originally recorded with 4:2:0 chroma sampling, preemptively converting it to a 4:2:2 or 4:4:4-subsampled format before you conform the media for grading can make these clips easier to work with in terms of real-time decoding performance, but it won't, by itself, improve image quality. Keep in mind that a 4:2:0 to 4:4:4 conversion is automatically performed as the first step of the internal image processing pipeline by nearly all grading applications, which typically work in 32-bit floating-point 4:4:4 color space internally as a matter of course. So, doing this conversion prior to grading as a separate step is unnecessary unless you need to convert the media from a format that's not supported by your grading software to one that is.

On the other hand, you *do* want to render your final color-corrected output out of your grading application to a 4:2:2 or 4:4:4 subsampled format in order to retain the higher-quality image processing that your grading software produces. One tip, though: It's tempting to render projects using media acquired with a 4:2:2 or 4:2:0 chroma subsampling ratio out to a 4:4:4 chroma-sampled format for mastering. This is fine, but given what you started out with, this will generate huge files that will likely have no appreciable increase in quality over mastering to a less storage-intensive 4:2:2 chroma-subsampled format. However, if your source media was originally acquired in a 4:4:4 sampled format, then mastering to a 4:4:4 format is obviously ideal if you want to retain all the quality of the original.

Compression and Bit Depth

Different cameras let you choose among formats using differing amounts of compression. It goes without saying that less compression is better than more. For DSLR cameras recording H.264 video, common data rates range from 17–42 megabits per second (Mbps), depending on the frame size, frame rate, and level you choose (if these qualities are user selectable). This is an important consideration that differentiates more expensive professional camcorders, which typically offer more choices for recording less compressed video that can range from 145–440 Mbps (although the differing compression technology means this is not an apples-to-apples comparison). Additionally, most low-cost video acquisition formats are nearly always 8-bit, while 10- and 12-bit video capture is available for more expensive camcorders and digital cinema cameras.

H.264 compression is designed to be as visually imperceptible as possible; however, it's processor intensive to decode and can exacerbate the kinds of artifacts you get from chroma subsampling in situations where you need to make dramatic adjustments. In the editorial process, it's common to transcode H.264 media to another, more editing-friendly codec to reduce the processing overhead (freeing resources for real-time effects performance), and usually you end up transcoding to a 4:2:2 subsampled format when you do this, but the advantage of this workflow at this particular stage of postproduction is to improve real-time performance, not to improve image quality (as discussed previously).

As with chroma subsampling, promoting your video from a highly compressed 8-bit format to a less-compressed 10- or 12-bits-per-channel format does nothing to immediately improve the image. Keep in mind that image data that was discarded by compression and chroma subsampling while recording is lost forever, and most modern grading applications automatically promote any clip to 32-bit floating-point within the image processing pipeline regardless. However, rendering your final color-corrected output to a 10- or 12-bit format preserves any higher-quality image processing or compositing done by your grading application and is highly recommended.

Log vs. Normalized Media

Most digital cinema cameras offer the option of recording either *logarithmically encoded* ProRes or DNxHD media. Furthermore, you usually have the option of debayering raw media to *logarithmic-encoded* formats in your grading software. In both cases, creating log-encoded media lets you preserve the greatest latitude for adjustments in grading.

While each camera has a differing method of log-encoding that is customized to take maximum advantage of its particular sensor, many are based on the Cineon log curve that was originally developed by Kodak for scanning the 13 stops of latitude that film is said to record into the Cineon and DPX image sequence formats, in an effort to preserve as much detail as possible within the 10 bits per channel of data these formats use.

Log-encoded media should be considered as a sort of "digital negative," and while the initial appearance of log-encoded media is unpleasant, being deliberately low-contrast and desaturated, the recorded image preserves an abundance of image data that can be extracted for maximum flexibility in the grading process.

When debayering raw media, these log standards are usually available as a gamma setting of some kind. As of this writing, these standards include the following:

- Log C: Media recorded by ARRI Alexa cameras can be recorded or debayered using Log-C gamma and color processing, which is similar to the standard Cineon Log gamma curve.
- REDLog Film: Media recorded by RED cameras can be recorded or debayered using this logarithmic gamma setting that's designed to remap the original 12-bit R3D data similarly to the standard Cineon gamma curve, and this media is compatible with most log workflows, including those intended for film output.
- S-Log and S-Log2: Sony's proprietary S-Log gamma settings for its digital cinema line of cameras are very different from the Cineon curve, owing to their wide dynamic range. The original S-Log was introduced with the Sony F3 camera. S-Log2 was introduced with the Sony F65 and F55 cameras, owing to the even greater dynamic range those cameras offer. There are two methods that Sony recommends for normalizing this media using LUTs. A 1D LUT can be used to transform S-Log and S-Log2 clips into a standard Cineon (or Log-C) curve first, if that suits your workflow. Or, you can use a dedicated LUT to normalize S-Log and S-Log2 media directly. For more information on these formats, search the Web for Sony's document "S-Log: A new LUT for digital production mastering and interchange applications."
- BMD Film: Blackmagic Design's logarithmically encoded gamma setting is a modified version of the standard Log-C curve. This modifications is designed to emphasize the strengths of the sensors used by the Blackmagic Design cameras.

Despite the differences among each camera manufacturer's brand of log-encoding, the process of *normalizing* log-encoded data to match the appearance of the original scene is essentially a careful contrast adjustment, and there are several ways you can accomplish it, depending on the capabilities of your grading system. The process of normalizing and grading log-encoded media is covered in both Chapter 3 and Chapter 4.

Preserving Quality by "Shooting Flat"

A common strategy for preserving image quality when shooting with DSLRs that lack log-encoding is to instead use the in-camera menu settings to create an image recording profile with "flat" contrast in order to preserve highlights and shadows at the extreme ends of signal bandwidth. The idea is that by not "precorrecting" your image in-camera to boost its contrast, risking clipping of highlight or shadow detail that you might want to keep, you improve your ability to make more careful exposure decisions later, during the color correction process.

It's important to point out that "shooting flat" really means "recording image data flatly." In other words, it is neither necessary nor desirable to deliberately light your scene in a low-contrast manner. Instead, light the scene however you like—low key, high key, whatever—and instead use the settings of your camera to record a low-contrast signal that will preserve as much as the image data as is possible.

Similar to log-encoded media, clips recorded using flat data result in lackluster images when you first see them. However, that's just because they haven't been color corrected yet. Such "flat" imagery can result in superior shadow and highlight detail when you readjust contrast during the color correction process, but there are a couple of things you might want to keep in mind.

First, even though you're adjusting your camera settings to shoot low-contrast data in order to avoid clipping the bottom and top of the signal, you don't want to shoot *too* low contrast, or you won't be using enough of the range of the 8 bits per channel that you have available to you to preserve artifact-free midtone detail. Second, in choosing to shoot low-contrast data, you're forcing the need for color correction later. While obviously the purpose of this book is to encourage the grading of every project, if you're working on a project that will be constrained for time, this is a consideration.

If you're considering recommending a DSLR shoot flat, there are three widely publicized profiles for doing so.

- Prolost Flat (www.prolost.com/flat): Filmmaker and photographer Stu Maschwitz has long advocated these camera settings for shooting DSLR media that's easier to grade, and he gives a lengthy explanation at this page.
- Technicolor Cinestyle (www.technicolorcinestyle.com/download/): Grading powerhouse Technicolor released a downloadable camera profile for recording low-contrast media with wider latitude. The profile and user guide are downloadable at this page.
- Canon EOS Gamma Curves (www.lightillusion.com/canon_curves.html): Created by Steve Shaw's Light Illusion, these in-camera curve profiles can be used to maximize the image data stored in the limited bandwidth these cameras record to.
- Flaat Picture Controls for Nikon DSLRs (www.similaar.com/foto/flaat-picture-controls/index.html): Available from Similarly, this set of picture styles for both Canon and Nikon cameras also aims to give you a low-contrast, latitude-increasing starting point, but also claims to handle skin tones well.

Digital Dailies: the Start of Postproduction

In the classic film studio workflow, when production stopped for the day, the camera negative would be rushed to the lab for developing, and then workprints would be made and synced with production audio. The workprints would then be assembled into a set of *dailies* that the film crew would watch either that evening or the next morning to evaluate the day's performances, check that there's enough coverage for the scene, and make sure there are no technical problems. The dailies would then be handed off to the editing team, which would begin cutting.

Unless you're shooting a camera with tethered audio that's recording ProRes or DNxHD media directly and handing that media directly off to the editor to begin work, chances are that some kind of digital dailies workflow will be necessary, since raw and log-encoded media still has to be adjusted and processed, and separately recorded audio still has to be synced. On a professional set, the day's work should still be evaluated.

As mentioned, digital dailies are created either during production by the DIT or later in more controlled conditions by a colorist at a post facility. The process of generating dailies for a production generally has three components, though there may be other creative workflow innovations pursued by specific facilities.

Note

The consolidation of film labs into fewer and fewer markets makes the notion of same-day film dailies unlikely anymore, although scanned digital deliverables can speed up the return trip.

Syncing Dailies

If the production tethered audio to the camera during the shoot, then you already have high-quality production audio incorporated into the camera media, and there's no other syncing work to do.

On the other hand, productions recording dual-system sound, where the audio is recorded separately from the visuals, are going to require that the audio and video be synced when you create the dailies. If you're lucky and the production recordist and assistant camera operator (AC) were on the ball, this process can be automated in a couple of ways. If not, then you (or your assistant) get to rock it like it's 1985 and sync each pair of video and audio clips manually by lining up the closing of the clapstick of a clapperboard slate in the video (now possibly a tablet running a clapperboard app) to the audible clap on the audio track that's visible as a spike in the waveform. Clip by clip. Good times, but hey, at least you don't have to line up mag tracks on a Steenbeck.

Preferably, the production will be using timecode-synced dual-system sound, where time-of-day timecode is synced between the digital cinema camera and the audio recorder and periodically jam synced over the course of the day to keep the sync relationship frame-accurate. While the equipment and expertise are generally more expensive for other methods of syncing dual-system sound, with careful file management of the video and audio media, the matching timecode on each pair of video and audio files makes syncing the audio fast and nearly flawless. I've synced three days' worth of production media in seconds using this method, with no problems. Needless to say I'm a fan of productions that work this way. Applications including DaVinci Resolve, FilmLight Baselight, and Assimilate Scratch (and Scratch Lab) are able to facilitate timecode sync.

If timecode sync wasn't in the budget, you can still get good automated audio/video syncing using the technique of waveform synchronization. Using this method, an on-camera microphone records lackluster audio along with the video. The on-camera audio can then be used, when syncing to superior audio recorded elsewhere, to line up the matching waveforms of each pair of on-camera and dual-system audio files. Applications such as Red Giant PluralEyes are dedicated to this task and are capable of processing batches of files all at once for linking inside of other NLEs. However, editing applications such as Final Cut Pro X have also jumped into the fray, providing similar functionality that's built in.

Grading Dailies

Another task awaiting the DIT or colorist who is creating the digital dailies is the grading of the dailies, if necessary. When it comes to on-set grading and dailies work-flows, there are all manner of ways to proceed, depending on the type of production, the budget, and the schedule. I'll restrict myself to a very high-level overview here; bear in mind that this only scratches the surface of the workflows that are possible.

On-Set vs. In the Suite

Once, most dailies were done by film labs, and later by postproduction facilities, where junior colorists would often cut their teeth doing unsupervised evening-shift work, while senior colorists would work on the supervised daytime projects. Increasingly, digital workflows using more affordably priced and portable grading workstations are enabling this work to be done on-set by the DIT.

If a DIT is involved, on-set color is typically restricted to setting primary grades so that the cinematographer can see how the scene is shaping up via the live camera display in video village; this is especially critical in

workflows where log-encoded data is being recorded, which looks terrible unless monitored with an appropriate correction. As grading applications reach farther back in the pipeline toward production, there are increasingly sophisticated workflows becoming available that enable more intricate grading work to be done, if desired.

On the other hand, grading software developers are increasingly facilitating a bidirectional workflow, where the facility colorist and cinematographer is able to set looks in advance based on footage from test shoots, which can be made available to the DIT for reference and as a starting point for the on-set work.

Grading Data Interchange

The grades created by the DIT may be "baked" into the digital dailies being created for the benefit of the editor and director, but these grades may also be handed off to the colorist who's doing the final finish as a starting point for the final grade. Granted, they won't always be used, but at least they'll provide valuable insight into what the cinematographer was thinking on the set. There are a variety of ways these grades can be preserved and handed off for later use.

- Camera metadata: Digital cinema cameras usually store the ISO, exposure, and other metadata inside each recorded media file. Grading applications that are compatible with a particular raw camera format will be able to both read and manipulate this image adjustment metadata, which affects how that media is debayered for use within the image processing pipeline.
- Lookup tables: Used extensively in commercial shoots with limited locations, LUTs are saved image-processing operations that can be created to set looks for how the scene should appear when viewed by the on-set display. LUTs are advantageous since they can be loaded directly onto a variety of production displays, and they can also be handed off to the colorist in post for use within a grading application, either for reference or to use as a starting point for ongoing work.
- Color Decision Lists (CDLs): CDLs are an industry-standard file format originally developed by the American Society of Cinematographers' technology committee. CDL files are formatted similarly to EDLs, with SOP (Slope/Offset/Power) and SAT (Saturation) values embedded as metadata in much the same way as comments are in a more typical EDL. CDLs are used in television and long-form programming to organize on-set color data. Using a CDL, primary grade adjustments can be organized for a collection of shots in different locations and retrieved by the colorist later for reference or as a starting point for ongoing work.

Using Grading Applications on the Set

Although there are dedicated applications for facilitating on-set color and digital dailies workflows, since this book focuses on dedicated grading applications, I'll focus on three studio grading applications covered elsewhere in this book that also facilitate on-set workflows as a subset of their functionality. Of course, if you have the infrastructure, you can use pretty much any grading application on the set, but some applications are more carefully designed for this than others.

- *Scratch Lab* is a version of Assimilate Scratch that is specifically set up for doing on-set work and that can be run on a variety of very portable Windows and OS X computers. Primary grading, import and export of LUTs, CDL support, grade matching, and other features for processing dailies is built in, and once you're done, either you can either export LUTs and CDLs for use in other applications or you can move your on-set project into the full version of Scratch and begin working that way.
- FilmLight has developed a self-contained "Baselight in a box" called the *Flip*, which essentially puts a full Baselight product in a form factor that's easily equipped on the DIT's cart. The Flip can ingest live output from the camera so that the DIT can set looks as the crew works, as well as record video, but it

also has all of the dailies synching, LUT and CDL exchange, and media processing features of Baselight. If you're using Baselight from end to end, the Flip can save Baselight Grade (BLG) files that contain the full Baselight grade, LUTs that were used, and keyframes that were set, along with two reference stills (one graded, one ungraded) and timecode metadata. These BLG files can be shared among all versions of Baselight including Baselight editions for Avid, Final Cut Pro, and Nuke, as well as the Baselight studio software used for the final finish, so that grades generated on the set can be carried through and refined during editorial and VFX and then used as a starting point during the grade.

- DaVinci Resolve (either the Full or Lite version), which can run on your choice of Linux, Windows, or OS X computers, has a feature called *Resolve Live* that lets you monitor and grade the live image coming from the camera from within Resolve, while simultaneously monitoring live HD-SDI output, creating and saving full-blown Resolve grades while the crew works, along with an image still and timecode metadata that enables easy syncing to the camera original media when ingested later. Since it's DaVinci Resolve, you also have all of the LUT and CDL compatibility workflows, dailies synching, and media processing capabilities you'd have in the studio.

Other applications are available, but these provide a good look at how the on-set and facility grading experiences are becoming increasingly connected.

Note

If you're going to do on-set color, it's crucial that you have a high-quality, color-critical display as described in Chapter 2 and that you do your best to shield it from ambient light reflecting on the front of it so that you can judge the contrast of the image reasonably well. Otherwise, the resulting grading information that is handed off will not truly reflect an accurate transformation of the image data as it will be viewed in studio conditions. Granted, you can do only so much in the hectic conditions of video village, but the closer you get to studio monitoring conditions on the set, the better and more useful the grades you'll be handing off.

One Light vs. Best Light

When generating digital dailies, there two approaches. If the on-set look or saved camera metadata applies equally to all media from a given reel, then a "one-light" grade, where a single adjustment is applied to an entire collection of similarly shot media, should be sufficient to create decent-looking media for editorial to start with. If it's not perfect, that's not a problem since the offline media will be replaced with online media during the reconform process, just prior to final grading.

On the other hand, if the director is especially picky or if the camera original media is being transcoded to an online-quality format that will be used for final finishing, then it might be preferable to apply more detailed grades to each clip to bring out the best the media has to offer. This is typically referred to as a "best light" grade and is more typically carried out by a colorist back at the grading suite.

Camera Raw Transcoding Workflows

When shooting raw media, another aspect of digital dailies creation is the decision concerning how to go about transcoding a set of useful offline or online media. Although nonlinear editing applications are beginning to get fast enough to allow direct editing of camera raw media, in my opinion doing so is inadvisable, at least as things stand at the time of this writing. The performance requirements for real-time debayering are significant, and the trend of editors using more affordable and portable equipment to work is at odds with having to go back to using a high-powered desktop system just to accommodate a camera format (although don't let me stop anyone who wants to do this).

Also, the storage requirements even of compressed raw formats, while smaller than uncompressed mastering formats, are still far larger than the low-bandwidth offline codecs that are available for editing. Consequently, using offline-quality media with smaller file sizes and lower bandwidth requirements can significantly accelerate the performance of one's editing software, as well as lowering your storage requirements by an order of magnitude.

This means that the creation of an alternate set of matching digital dailies is typical in raw workflows. If shooting raw and creating digital dailies is the preferred workflow, there are three ways of dealing with this.

- Debayer and transcode the raw media to an offline format that's lower quality, that's lower bandwidth, that has smaller file sizes, and that is overall easier to edit, but that will require reconforming to the camera original raw media when grading and finishing. While the reconform is an extra step, raw media gives the colorist maximum flexibility for making adjustments, assuming their grading software supports this workflow.
- Debayer to DPX image sequences (typically log-encoded) to create uncom-pressed, mastering-quality media that you can use for finishing if you're in a workflow where raw files are inconvenient but you need top-quality media, such as for VFX-heavy programs. DPX clips are huge and will require significant storage capacity for longer projects, and they aren't typically NLE-friendly, so you'll need to also generate a set of edit-friendly offline media with matching time-code and reel information to facilitate reconform to the DPX sequences during finishing. You lose the flexibility of raw media, but log-encoded DPX files will have all the image data you need if the media was competently shot with appropriate metadata settings from the camera.
- Debayer and transcode the raw media to a mastering-quality yet NLE-friendly format (QuickTime ProRes or MXF) at the resolution you'll be finishing the project at, ideally as log-encoded media if you want to preserve the maximum latitude for later adjustment. This generates media that won't need to be recon-formed to another format later but that will be compressed (depending on the codec) and will likely have higher storage requirements if you choose a high-quality codec like ProRes 422 (HQ), ProRes 4444, DNxHD 220Mbit/s, or DNxHD 444. However, since storage is pretty cheap these days, this may not be an issue depending on the type of project you're editing. The advantage of filling up your storage system with mastering-quality media is that you don't have to reconform to the camera original raw media when you finish, sacrificing the flexibility of raw (which you may or may not need) in order to save yourself the potential hassle of an extra step.

While there are vocal advocates for all of these workflows, I do not prefer one over the other. All have their merits depending on the nature of the project at hand, and I've done all three—raw to offline with a reconform, raw to DPX to offline with a reconform, and transcoding to and finishing with mastering quality media—with satisfying results relative to each project.

The key, if you're planning on debayering and transcoding to a mastering-quality format, is to do so using a high-quality codec, to check your metadata settings and perhaps make some simple adjustments (if necessary)

to make sure you're debayering files with appropriate ISO and exposure settings, and to preferably debayer log-encoded media (discussed next) that preserves the greatest latitude for adjustment later, during the grade.

Round-Trip Workflows

While there are some exceptions, grading applications are generally designed to import an EDL, XML, or AAF file that's been exported from an NLE after the edit has been completed and then to export a corresponding EDL, XML, or AAF file of the finished project to send back to the original NLE. This process is referred to as a *round-trip*, and this section illustrates a generic round-trip workflow to help you understand what is involved, overall, and where grading fits in.

An outline of the average round-trip workflow is as follows:

1. Lock the edit (you hope).
2. Prep your timeline for handoff.
3. Export the edit, and organize its accompanying media.
4. Grade the project.
5. Reconform last-minute VFX changes and stock footage buys.
6. Render the final graded media.
7. Export the graded timeline, and reimport it into your NLE or finishing application.

Every application does this differently. For example, Adobe SpeedGrade has the enviable ability to import a Premiere Pro file directly, rather than relying on XML or AAF interchange formats. However, if you're going to be organized (something I strongly suggest), variations of these steps apply no matter what combination of postproduction applications you plan on using.

Before You Begin: Locking the Edit

Much has been said about the benefits of applications that allow more flexible workflows in terms of passing timelines back and forth between the editorial and finishing departments. Ultimately, the desire is to eliminate the notion of "locking" changes made to an edit altogether, to enable any little alteration to be implemented at any time right up until the program is output.

While this sounds great on paper and this kind of flexibility is certainly welcome in terms of application interoperability, many key advantages to locking an edit remain, not the least of which is cost savings. At the risk of editorializing, locking the edit should not be looked at as a technological limitation but as a *scheduling milestone*.

Sooner or later, the director and producer will have to make up their minds, be done with the edit, and allow the project to go through finishing. The longer this is put off, extending reedits into the finishing stage, the more time the grading team will have to spend reconforming these last-minute changes, with their attendant alterations to grades within each updated scene, and the more hours of work will end up on the bill (or so the finishing facility hopes). If the project is a tentpole summer blockbuster with a $200 million budget, this may not be an issue, but if the project is a $500,000 indie feature or $100,000 documentary, those hours add up fast, and not just because of the grade.

Typically, audio is mixed at the same time that the visuals are being graded, and any changes made to the edit will affect the audio timeline just as much as the grading timeline. For the lower-budgeted project, locking the edit to avoid such last-minute changes reaps all kinds of efficiencies when it comes to the finishing process, and that will save you cash.

None of this is to say that you shouldn't expect updates to titles, visual effects shots, last-minute stock footage

purchases, or other media substitutions; these sorts of one-to-one clip replacements are inevitable, and modern grading applications make these sorts of changes relatively easy to handle. What becomes problematic are more sweeping changes made to the program that alter the total run time by rearranging clip position and duration for multiple scenes, which can become very complicated very quickly without careful organization.

Bottom line, if the edit can be locked, it should. If it can't, someone should ask why not. All things must end.

Preparing Your Edit for Grading

It's always a good idea to do some prep work to your edited sequence before handing it off for grading. Every grading application is compatible with a different collection of effects beyond edits and dissolves, but there are invariably going to be some effects and media types that won't be supported. While often these effects are either ignored or preserved and sent back to the original NLE after grading has finished, these sorts of things can cause unwanted complications, so a bit of timeline reorganization will always help keep things on the rails. In general, I always recommend making the following organizational changes to a duplicate of your edited sequence.

Move Clips to Track 1

It's best to move all noncomposited clips to track V1. Many editors use superimposed series of clips not to create layered effects but to edit a scene together. While this works well within a given NLE, color correcting a program with hundreds of clips spread across several video tracks can be a pain in the neck for the colorist to manage. It's much easier to keep track of what you're doing and to copy grades across multiple shots when they're all on the same video track.

On the other hand, superimposed clips that are part of a compositing or transparency operation should be left alone. Many grading applications have incorporated the ability to do basic compositing, so there is a good chance your application can import these effects, or at the very least re-create them.

Isolate Unsupported Effects

Checking ahead to see which timeline effects a grading application does and does not support is a great time-saver for the person doing the reconform at the finishing facility. One way to do this is to move all unsupported clips to a superimposed video track. For example, many grading applications don't support long-duration still image files, freeze frames created by the NLE, generators specific to a particular NLE, or exotic compositing operations. When you export a project from an NLE to a grading application using XML or AAF, either these types of clips will usually not appear or they'll appear as offline or disconnected clips. Sometimes these unsupported effects are preserved internally in order to let them be sent back to the original NLE on the return trip, but sometimes not.

If you have these types of clips in your project and you don't need to grade them, you can simply ignore them. For example, you typically don't bother color correcting clips such as titles, lower-thirds, or other computer graphics that have been created specifically for the program, on the premise that they've been created specifically with the limits of broadcast or film output in mind.

On the other hand, if there are freeze frames or composited effects that you do need to grade, a good workflow is to do the following:

1. Move the unsupported clip to a superimposed track.
2. Render it as a self-contained media file using a mastering-quality codec.
3. Reimport the media file movie you just rendered back into your project, and edit it back into the sequence on track V1.

At this point, you can delete the original, superimposed clip in track V2, but leaving it there makes it easier to locate and rework it should you decide you ever need to make a change to the original effect. Now that the effect has been turned into a plain old media file, it can be graded like any other clip.

Figure Out what to Do with Speed Effects

Grading applications all have varying support for speed effects, *especially* variable speed effects, so check ahead to see whether your grading application will work properly with imported speed effects coming from the NLE in question. If not, there are plenty of ways you can use plug-ins, built-in NLE functionality, or dedicated compositing effects applications to preprocess high-quality speed effects using optical flow processing in order to create a high-quality self-contained media file with which to replace the original speed-affected clip, all *before* you move the media into your grading application of choice (similar to the workflow presented in the previous section).

Sort Out an Edit's Effects Plug-Ins

Grading applications don't typically support the same effects plug-ins that are used by NLEs—at least not in the same format. If there are plug-in effects that you want permanently applied to a clip prior to grading, you'll need to render out and reimport "baked" versions of these clips as described previously. However, you should consider just what kind of effect it is. If it's a grading plug-in or an effect or look that your colorist might be able to do better, you should strip it out of the edit prior to handoff.

However, before you start stripping out an editor's grades and looks, it's often useful for the colorist to refer to such offline corrections and effects during the grade so they know what the client has gotten used to. Consider rendering a self-contained "reference" file of the entire sequence in QuickTime or DNxHD form to hand off along with the actual project and source media. Such a reference movie serves several purposes. It provides a reference of how the edits in the timeline are supposed to line up, in case something goes wrong during import. It gives a reference of which clips are supposed to go where, in case something goes wrong with the conform. And it provides a visual reference of what kinds of sizing changes, temp grades, and other effects have been applied to the project that you may need to re-create (or outdo).

Tip

Dave Hussey, senior colorist at Company 3, refers to the phenomenon of "temp love," where the clients have fallen so in love with temp music or a temporary grade during the edit that it's difficult for them to fairly evaluate the original work it is being replaced with. It's a genuine challenge you'll face time and time again, and it will test both your salesmanship and your patience.

Preparing Your Media, If Necessary

As with any finishing workflow, it's essential to use media that has been ingested and/or transcoded at the highest possible quality. In general, there are two ways to do this, depending on what kind of media was used for editing.

Programs Edited Using High-Quality Media

For example, if you're working on a program that was edited using mastering-quality media from the beginning such as ProRes 422 (HQ), ProRes 4444, DNxHD 220, or DNxHD 444 (or whatever codec the production decided was suitable for mastering), then probably all you have to do is to export an XML, AAF, or EDL of your edited sequence, media manage the accompanying media onto a portable hard drive, and hand those things over to the grading facility for them to be able to quickly reconform the project on their workstation.

Programs Edited Using Offline Media

If you're working on a project that, for convenience, was edited using an offline-quality format such as ProRes 422 (Proxy) or DNxHD 36, then you're going to want to conform the exported XML, AAF, or EDL of the edited sequence either to the high-quality camera original media that was originally recorded or to a set of online-quality transcoded media that was generated for that project, from which the offline clips were generated. Doing this kind of *reconform*, where you substitute the offline-quality clips with the camera original or online media, is a core feature of most grading applications.

Reconforming typically uses a combination of media filenames, unique identifier metadata (UUID), timecode, and reel name information to match each offline-quality clip to its corresponding online-quality clip in order to replace each low-quality clip in a timeline with its high-quality equivalent, ready for grading. For this reason, it's essential to manage the metadata of your media with care from production through postproduction.

If the offline-quality media was output by the same grading application that will be used for the final finish, then the easiest thing to do is usually to conform the edited project to the original database of media created to output the dailies. If this is the case, things should go quickly and easily. Otherwise, it's usually not that difficult to copy the camera original data to the facilities storage system in preparation for project reconform from scratch.

Automatic Edit Detection

Every once in a while, you'll end up needing to grade a program for which there's only a single rendered master media file that's available. If this is the case, you can sometimes use a grading application's ability to "notch," or cut up a single media file to match the edits specified by an EDL. If you don't even have an EDL, however, then some grading applications (including but not limited to DaVinci Resolve, FilmLight Baselight, and Adobe SpeedGrade) have the ability to automatically detect edit points in a movie file (based on changes to color and contrast) and build an edit list that you can then verify and use to cut the file into individual clips for easy grading. This is particularly useful for archival work-flows from film or tape, where only the final master remains.

Importing the Project for Grading

Nearly every grading application is capable of importing a variety of project interchange formats, including XML, AAF, and EDL files. None of these files is an actual format saved by a nonlinear editing application; they're file formats that are specifically exported by NLEs for the express purpose of converting a sequence of clips and effects in one application to a matching sequence of clips and effects in another. Again, importing one or more of these formats is a core feature of grading applications; you need only determine which format is best for your particular combination of editing and grading application.

Rendering the Final Graded Media

For most workflows, it is necessary to render a completely new set of graded media files from your grading application before sending the edit to the originating NLE or on to a finishing application. There are two ways you can do this.

Rendering Individual Clips

If the client is concerned about last-minute changes being made to the program, they may ask for individually rendered clips such that the edited timeline can be completely reconstructed, only with graded shots. This is a workflow that every credible grading application can support. In general, you'll end up with one new color-corrected media file for every shot in the edited sequence that's been graded, and it is these new media files that will be linked to when you send the finished XML, AAF, or EDL file from the grading application back to the NLE or finishing application where the application will be completed.

Rendering a Graded Master Media File

Alternately, if the project is well and truly finished, you may be asked to simply output the entire timeline as a single digital master file. Typically this will be a text-less master that can be loaded back into an NLE or finishing application for the final application of titles, end credits, and anything else that needs to be added prior to final output and delivery.

Send the Edit or Master Render File Back for Final Finish

After you've rendered the graded project, either as individual clips or as one great big one, you need to move the resulting media back into an NLE or finishing application in preparation for whatever else needs to be done (titling, format conversions, last-minute VFX, marrying the soundtrack, and so on). In particular, if you rendered individual media files to preserve the structure of the edit to facilitate truly eleventh-hour changes, then you'll send an XML, AAF, or EDL back to your NLE or finishing application—which is why it's called "round-trip."

Once the project is in the finishing application (assuming this is even necessary), the rest of the team can add whatever is necessary for final mastering and delivery.

Where Does Grading End and Finishing Begin?

Where once there was a hard dividing wall between grading applications and "finishing" applications, with the evolution of the main grading applications used by the industry, there is increasingly a gray area between grading and finishing applications that gets grayer every year. The definition of finishing depends on who you talk to, but in general it includes editorial changes, titling, clip resizing, multiple-format support and format conversion, effects such as digital paint to remove previously unnoticed flaws and trackable blurs to eliminate onscreen elements the producers don't have rights to, and the ability to output to tape, digital formats, and master other sorts of digital deliverables such as digital cinema masters (DCMs) for projects requiring digital cinema distribution.

Over time, more of these features have creeping into the toolkits of various grading applications, such that in some instances a full round-trip may be unnecessary assuming you can output directly from the grading application. Realistically, however, most workflows involve a return-trip to the editing application of origin.