21.1: *Invited Paper:* The Hopeful Future of High Dynamic Range Imaging

Greg Ward
BrightSide Technologies by Dolby

Abstract

This paper offers an overview of the challenges and opportunities presented by high dynamic range (HDR) imaging. We examine the length of the imaging pipeline, from creation and storage through image editing and viewing, and discuss how each stage is affected by a move to HDR.

Introduction

Since the first release of Photoshop™ in 1990, imaging has been well-grounded in an 8-bit/sample (i.e., 24-bit/pixel RGB) world. This is not without reason. An 8-bit integer provides enough levels on a standard display that banding is almost negligible using a *de facto* γ=2.2 encoding. Although considered hefty at its introduction, the 24-bit RGB standard is a reasonably efficient representation, which permits in-core editing of screen-sized displays, using less than 1 MByte for a VGA-resolution image. Years later, 24-bit color was also favored during the dissemination of color management systems, since it conveniently serves as an index in a 3-D lookup table.

As we adopt a wider conduit for imaging, many of the decisions that came before need to be re-examined, and some need to be reconsidered. Naturally, there are practical limits to the changes we can make. Much of the imaging pipeline is ingrained to the point where major changes would cause a technological upheaval whose short-term costs would undermine or even outweigh its long-term benefits. The purpose of this paper is to examine some of these technological trade-offs, compare an "ideal" HDR imaging pipeline to a "likely" one, and consider the need for backward-compatibility. We hope this exposition will spur additional ideas and solutions to the problems presented.

Image Creation

Most modes of digital image creation, including paint software, still cameras, video cameras, and low-end rendering and animation systems, work in a 24-bit "output-referred" color space, such as *sRGB* [Stokes et al. 96]. This is a convenient choice for viewing and editing on common video display devices, such as CRTs and LCDs, which have a limited dynamic range and color gamut. So long as the color depth resolution of the created imagery meets or exceeds the output device, the latter is unlikely to show up deficiencies in the former. However, as we graduate to higher bit depths and perceptual range in our display systems, digital cinema and high-end home theater will expose inadequate source materials for what they are.

The special effects industry was the first to recognize that greater bit depths were needed before computer graphics (CG) would blend seamlessly with live-action film footage. The greater color resolution, gamut, and dynamic range of film reveal the shortcomings of 24-bit output-referred encodings, which include the notion that a maximum value somehow corresponds to "white." The real world presents highlights that are 1,000 to 10,000 times brighter than the 18% gray level commonly used as a reference in cinematography, and these same highlights must be represented in special effects work to incorporate lens flare and similar cues that something is "brighter than white." In the absence of HDR, special effects lack the depth and realism of live action.

An important driver for HDR in this context is the image-based lighting (IBL) technique introduced in [Debevec 98]. Using IBL, one may capture an HDR image of the environment reflected in a mirrored sphere, and use this to illuminate CG elements so they will blend convincingly with the captured film footage. Figure 1 outlines the basic application of this method.

Image-based lighting is now a principal practice employed in special effects for film, and its reliance on HDR imagery in certain parts of the pipeline have led to a general migration towards floating point representation throughout. The gaming industry is also pushing HDR applications, especially since the majority of graphics cards now support 16-bit/channel floating point natively in their texturing and rendering engines.

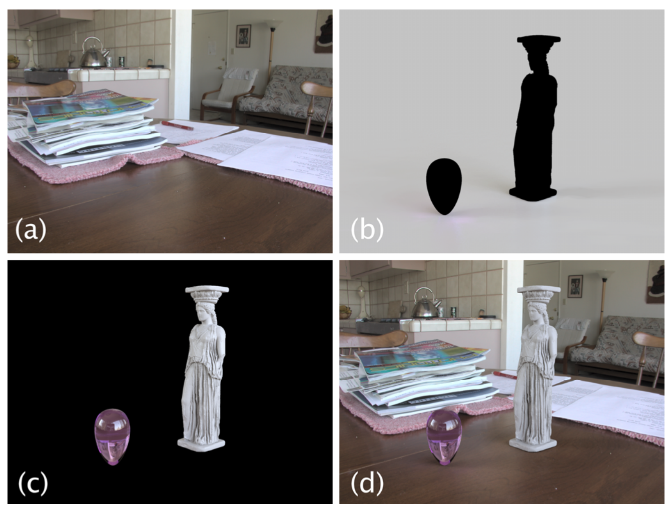

Figure 1. Image-based lighting [Debevec 98]: The background plate (a) is multiplied against a shadow and scattering image (b). The CG elements (c) are illuminated by an HDR *light probe* then composited into the final frame (d).

Capture Methods

While rendering and special effects applications have already made the transition to high dynamic range, capture technology appears to be lagging behind. The principal methods today for HDR moving and still image capture are (respectively) film scanning and multiple, bracketed exposures.

For movies, film offers a sufficient latitude that careful scanning can yield subtle shadows simultaneous with bright highlights, unlike today's digital video cameras. Unfortunately, scanning film reels with the necessary bit depth is an expensive and timeconsuming operation, thus it is presently limited to big-budget productions. Furthermore, film scanning presents many challenges due to the carefully tuned and highly non-linear response of film to scene colors and intensities. Merging scanned film imagery with the linear world of computer graphics falls somewhere between science and art, and only a few houses have mastered it. If an HD-resolution digital video camera with true HDR output were to arrive on the market tomorrow, it might well change the entire industry.

For still capture, the method of multiple exposures popularized in [Debevec & Malik 97] ten years ago has yet to be replaced. Because a sequence of exposures necessarily span a substantial time frame, camera/tripod shake and scene movement are perennial problems. In the case of camera movement, fairly robust image alignment is available using either the Median Threshold Bitmap technique [Ward 03], Scale Invariant Image Feature Transforms [Lowe 04], or image-flow methods [Kang et al. 03]. Image-flow methods are also useful for rectifying scene movement. Other “ghost removal” techniques include variance-based segmentation [Reinhard et al. 05], and robust estimators [Khan et al. 06]. An HDR still camera would eliminate the need for such workarounds, and open up a world of new possibilities beyond 24-bit RGB.

Image Transmission

The first viable HDR image format was introduced as a little-known “exponent” extension to the Utah Raster Toolkit [Petersen 04]. Independently, the author developed a nearly identical 32-bit RGBE format as part of the *Radiance* rendering system [Ward 91] [Ward 94]. Several years later, 24-bit and 32-bit LogLuv extensions were added to the TIFF library, covering the full visible color gamut and dynamic range [Ward Larson 98]. More recently, the 48-bit/pixel EXR format was made public by Industrial Light and Magic in the form of the excellent OpenEXR™ C++ library [Kainz et al. 2002].

The current trend in HDR image transmission is towards better, customized compression algorithms. For example, ILM’s EXR image data starts out with 50% more bits/pixel, but ends up taking the same or less space than *Radiance* RGBE thanks to its lossless wavelet compression. (*Radiance* uses a basic run-length encoding scheme.)

In the realm of “lossy” compression algorithms, there is a trend towards backwards-compatible formats, such as JPEG-HDR [Ward & Simmons 04] and MPEG-HDR [Mantiuk et al. 06]. These formats have the dual advantage of taking up a small fraction of the space of the lossless HDR streams, while being displayable using conventional hardware and software. They are thus ideally suited to the internet, video, and perhaps a new generation of digital cameras.

Backwards-compatible HDR formats are a win-win for manufacturers and consumers alike. Easing the transition from 24-bit RGB to full dynamic-range and gamut capture, editing, and viewing has the potential to greatly accelerate market penetration. Early adopters would gain immediate access to an HDR world, while others would not be inconvenienced, and could even benefit from improved tone-mapping in their legacy content. A well-chosen strategy may even simplify the eventual retirement of low dynamic-range data years hence [Ward 06].

Image Editing

As we noted above in our discussion of special effects, HDR offers numerous advantages for combining disparate image sources such as film and CG. In particular, the freedom from a hard ceiling on values and the inclusion of the entire visible gamut avoids a host of mapping problems. Similarly, HDR offers numerous opportunities for image editing, though it presents some challenges as well.

The first challenge for image editors is to overcome the notion of a maximum value corresponding to “white.” This concept does not apply in the real world, until and unless you print something on paper. Being restricted to a maximum white during editing, even in cases where you do not intend to print, is unnecessarily constraining. One approach to this problem is to provide a slider to adjust the range of display values, or to edit locally exposed subregions. Until desktop systems are equipped with HDR displays, this may be a necessary compromise.

A parallel challenge is the appropriate on-screen representation of out-of-gamut colors. To some degree, we face this already with 24-bit RGB in the disparities between input, display, and hardcopy devices. The usual approach is to map colors to the current display, optionally highlighting pixels that will be out-of-gamut upon printing. The situation becomes even more interesting when we consider colors outside the visible gamut, which most HDR formats can represent. Overall, HDR offers the simplicity of working in a scene-referred linear color space, which streamlines color management enormously.

In Photoshop today, 32-bit/channel features are about where 16-bit/channel features were 10 years ago. Floating-point has been introduced, but most editing functions are currently disabled or restricted, and we expect the support to improve gradually with time. Cinematic image editing and compositing tools are pushing the envelope more forcefully, and there are other consumer-level image tools that are quite powerful in the HDR domain (e.g., Idruna’s Photogenics™ HDR, Artizen™ HDR, and Cinepaint).

Our fondest hope is that the current RAW format craze will give way to a more productive and long-term investment in HDR. In many ways, RAW is just a poor man’s HDR, where every camera has its own representation and each software application has to extract the image as best it can. Building a scene-referred HDR pipeline is like building an autobahn – suddenly there is a compelling reason to engineer better systems. In contrast, RAW is like 1000 gravel driveways, ultimately leading us nowhere.

Image Viewing

Ideally, everyone would have a high dynamic range display that would permit users to view and work effortlessly with HDR images [Seetzen et al. 2004]. While we’re at it, let’s give our displays four or five spectrally-spaced primaries so they cover the entire visible gamut. One and possibly both of these wishes will come true in the next 2-7 years, but in the meantime, we need some method to represent colors that are outside the range of our display. Regardless of display advances, printing will still require some gamut mapping, at least until someone invents self-luminous paper.

Tone-mapping is the general technique employed to reduce an image’s dynamic range and color gamut to fit on a monitor, projection system, or print. It has a long history, extending back to the invention of negative photography. Tone-mapping was first introduced to computer graphics in 1993 [Tumblin & Rushmeier 93], and HDR tone-mapping has been explored extensively over the last five years [Reinhard et al. 05]. The best HDR tone-mapping algorithms also happen to be the most expensive, adjusting luminance values locally in a perceptually optimized way [Durand & Dorsey 02] [Fattal et al. 02] [Reinhard et al. 02]. By comparison, the older, global tone-mapping operators can be applied in real-time, providing convenient image editing and HDR video display [Krawczyk et al. 05]. Eventually, local TMOs may achieve real-time rates, but the path is not yet clear.

Digital cinema will probably be the first arena where medium-to-high dynamic range imagery will be presented. After all, cinematic prints already encompass a wider gamut and dynamic range than other media, and moviegoers have come to expect a certain richness. No theater owner wants to make an expensive upgrade that amounts to a backwards step in quality, so HDR could even be considered a prerequisite for digital cinema.

Once accustomed to a more exciting theater experience, consumers might start looking for equipment that offers similar dynamic range in their homes. In fact, the home theater market is currently moving towards HDR faster than either digital cinema or computer display equipment. This may be driven by vendors' desires to win a competitive advantage before large, flat-screen televisions saturate the market. If such a transition happens too quickly, we may find ourselves in the awkward position of having to synthesize all our HDR imagery from LDR content – essentially the inverse of the tone-mapping problem [Meylan et al. 06] [Rempel et al. 07].

Conclusion

It is clear that high dynamic range imaging will someday dominate the market. The question is, when? Our best estimate is between 2 and 7 years, and many things can influence its advance. It is preferable that HDR be introduced well rather than quickly. As engineers and imaging scientists, we are not powerless to affect this process. Through careful planning and intelligent standards-making, we can grow our businesses while delivering well-timed improvements in professional and consumer equipment and software.

The logical path is to introduce HDR to the high-end digital cinema market first, simultaneously with independent niche markets such as medical imaging, then allow a some years to pass before introducing HDR to the consumer. At that point, we will have settled the standards and worked out the kinks, the studios will have plenty of HDR content, and the job of consumer education will already be done.

References

DEBEVEC, P. 1998. Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography. In *Proceedings of ACM SIGGRAPH 1998*.

DEBEVEC, P., and MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proceedings of ACM SIGGRAPH 1997*.

DURAND, F., and DORSEY, J. 2002. Fast Bilateral Filtering for the Display of High-Dynamic Range Images. *ACM Transactions on Graphics*, 21, 3.

FATTAL, R., LISCHINSKI, D., and WERMAN, M. 2002. Gradient Domain High Dynamic Range Compression. *ACM Transactions on Graphics*, 21, 3.

KAINS, F., BOGART, R., HESS, D., SCHNEIDER, P., ANDERSON, B., 2002. *OpenEXR*. www.openexr.org/.

KHAN, E.A., AKYUZ, A.O., REINHARD, E. 2006. Ghost Removal in High Dynamic Range Images, *IEEE International Conference on Image Processing* (accepted for publication)

KANG, S.B., UYTTENDAELE, M., WINDER, S., SZELISKI, R. 2003. "High Dynamic Range Video," *Proceedings of ACM SIGGRAPH*.

KRAWCZYK, G., MYSZKOWSKI, K., SEIDEL, H.P. 2005. Perceptual Effects in Real-time Tone Mapping, *Proc. of Spring Conf. on Computer Graphics*.

LOWE D. G. 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 2.

MANTIUK, R. EFREMOV, A., MYSZKOWSKI, K., SEIDEL, H-P. 2006. Backward Compatible High Dynamic Range MPEG Video Compression, *Proc. of SIGGRAPH '06 (Special issue of ACM Transactions on Graphics)*.

MEYLAN, L., DALY, S., SÜSSTRUNK, S. 2006. The Reproduction of Specular Highlights on High Dynamic Range Displays, *IS&T/SID 14th Color Imaging Conf.*

PETERSEN, J. 2004. float_to_exp(3) man page. www.cs.utah.edu/gdc/projects/urt/help/man3/float_to_exp.html

REINHARD, E., WARD, G., PATTANAIK, S., DEBEVEC, P. 2005. *High Dynamic Range Imaging: Acquisition, Display, and Image-based Lighting*, Morgan Kaufmann Publishers, San Francisco.

REINHARD, E., STARK, M., SHIRLEY, P., FERWERDA, J. 2002. "Photographic Tone Reproduction for Digital Images," *ACM Transactions on Graphics*, 21,3.

REMPEL, A., TRENTACOSTE, M., SEETZEN, H., YOUNG, D., HEIDRICH, W., WHITEHEAD, L., WARD, G. 2007. "Ldr2Hdr: On-the-fly Reverse Tone Mapping of Legacy Video and Photographs," *Proc. of SIGGRAPH '07 (Special issue of ACM Transactions on Graphics)*.

SEETZEN, H., HEIDRICH, W., STUEZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., VOROZCOVS, A. 2004. "High Dynamic Range Display Systems," *Proc. of SIGGRAPH '04 (Special issue of ACM Transactions on Graphics)*.

STOKES, M., ANDERSON, M., CHANDRASEKAR, S., and MOTTA, R. 1996. Standard Default Color Space for the Internet. www.w3.org/Graphics/Color/sRGB.

TUMBLIN, J. AND RUSHMEIER, H. 1993. Tone Reproduction for Realistic Images, *IEEE Computer Graphics and Applications*, 13(6).

WARD, G. 2006. A General Approach to Backwards-Compatible Delivery of High Dynamic Range Images and Video, *Proceedings of the Fourteenth Color Imaging Conference*.

WARD, G., AND SIMMONS, M. 2004. JPEG-HDR: A Backwards-Compatible, High Dynamic Range Extension to JPEG, *Proceedings of the Thirteenth Color Imaging Conference*.

WARD, G. 2003. Fast, robust image registration for compositing high dynamic range photographcs from hand-held exposures, *Journal of Graphics Tools*, 8(2):17-30.

WARD LARSON, G. 1998. Overcoming Gamut and Dynamic Range Limitations in Digital Images. *Proc. of IS&T 6th Color Imaging Conf.*

WARD, G. 1994. The RADIANCE Lighting Simulation and Rendering System. , In *Proc. of ACM SIGGRAPH*.

WARD, G. 1991. Real Pixels. In *Graphics Gems II*, edited by James Arvo, Academic Press.