6.3: High Dynamic Range Imaging Pipeline

Helge Seetzen[1,2], Greg Ward[1], Lorne Whitehead[2]
[1] BrightSide Technologies Inc.
[2] University of British Columbia
Vancouver, Canada

Abstract: *A series of solutions for camera, graphics and display applications have been developed in order to create a full high dynamic range imaging pipeline. Every element of this pipeline features very high contrast, brightness and amplitude resolution as well as compatibility with existing infrastructure. A brief description of the capture, storage/process and display technology is provided as well as an outline of benefits of the combined system. Relevant performance metrics and comparison standards with other devices are also discussed.*

Keywords: display; camera; image processing; file format; high dynamic range

Introduction

The emergence of a variety of imaging technologies over the last few decades has shown that the paradigm of "one camera, one display" is doomed to disappear. Instead, each camera or display technology brings new performance envelops including brightness, contrast, color gamut and video quality. Nevertheless, a strong trend is noticeable for camera and display devices popular today – they are limited by their own past. Methods for capturing and subsequently displaying images have historically been based on output-referred solutions. Analog film cameras have provided us with image data that spans three to four orders of magnitude of dynamic range. Digital cameras today struggle to reach the dynamic range of analog film. Likewise, the first CRT displays provided a limited brightness range and today's line-up of display technologies remains bound to this low dynamic range.

As every photographer knows, such a limited dynamic range is insufficient to capture or portray the world around us. To overcome this limitation BrightSide Technologies has developed a series of imaging solutions designed to allow scene-referred imaging where the ultimate limit is the visual capabilities of the viewer rather than any elements of the imaging pipeline. The following briefly introduces the technologies making up the pipeline and provides an overview of their performance.

Imaging Pipeline

A useful high dynamic range (HDR) imaging pipeline requires three components: capture, storage/process and display [1]. Each component needs to support the very demanding levels of contrast, brightness and amplitude resolution accessible to our visual system. Performance aside, there are a number of necessary conditions for broad usage of such a pipeline:

- Compatibility with existing infrastructure (e.g. data, computing environments, TV signal, etc)
- Comparable cost to current solutions
- Independence of each component of the pipeline (i.e. the HDR camera should interface with a conventional display, etc)

In the absence of these conditions an HDR imaging pipeline would function as such but unlikely to be used in consumer applications. Especially backward compatibility is critical for this step as the emergence of High Definition TV (or the lack thereof for many years) has shown. The following describes each component and provides a brief overview of their compatibility limitations.

2.1 Camera

A wide range of high dynamic range camera solutions has been developed over the last few years. With the exception of higher amplitude resolution CMOS sensors most of these cameras rely on the capture of multiple image frames at different exposures [2]. A compositing algorithm can then be used to recombine these exposures into a single image with a higher dynamic range than any of the individual frames. Several specialty camera manufacturers have developed high frame rate cameras which can be used to capture such an exposure sequence in comparable time to a normal single exposure. Unfortunately, these devices are currently too expensive for the consumer market. In principle this technique can also be used by consumer cameras but the requirement for 3 or even 5 individual shots for a single HDR image limits its use to still scenes and often requires the use of a tripod.

A novel drive scheme for conventional CCD cameras eliminates the need for extended captures times and allows for an upgrade of consumer cameras to HDR cameras. Based on the functions of a conventional CCD camera this drive scheme reverses some of the image capture steps. During the time from opening and closing of the mechanical shutter a conventional CCD camera captures a single image on the CCD which is then shifted first under the shadow-mask and afterwards out of the CCD into memory. During the same time period the HDR camera captures a second image. This is achieved by shifting the CCD charges after some period of exposure without closing the mechanical shutter. Once the mechanical shutter is closed afterwards the CCD contains two images: One below the shadow-mask and one still in the active element of the CCD. Shifting both images out of the CCD and into memory provides two exposures separated by several f-stops which can then be combined as shown in Figure 1.

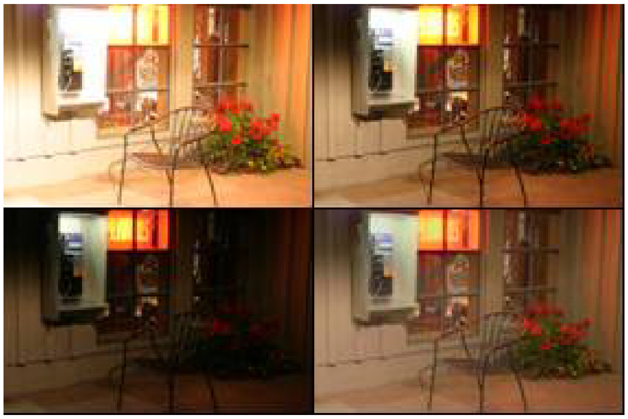

Figure 1: Multiple exposure combination (top left: long exposure captures dark areas, top right: mid exposure, bottom left: short exposure captures bright areas, bottom right: final combined image showing both dark and bright areas)

The benefits of this approach lie in its compatibility with existing cameras without restricting the cameras use. The combined exposure time is in line with conventional "point and shoot" style cameras. Furthermore, since there is no change to hardware of the conventional camera there is likely no cost increment either. Consequently this HDR camera method achieves the objectives of compatibility, no cost increase and independence outlined above.

Process and Storage

Acquiring HDR images results in substantial data file sizes. Many file formats exist for the storage of HDR data (OpenEXR [3], Tiff LogLuv [4], Radiance RGBE [5], etc) but similar to the higher end HDR cameras described above, their usefulness in basic consumer applications is limited. In most cases the compression factor is only modest and the resulting files are simply too large to be use by the average consumer. Conventional HDR file formats are also generally unreadable by most image viewers, which further limits their distribution. Generally only tone-mapped HDR images appear in the consumer space but these have of course lost any additional information.

The challenge of compressing HDR images lies mostly in the higher amplitude resolution. Noise and small image variations which would disappear in conventional 8-bit images are visible in HDR images and limit the efficiency of common compression techniques such as run-length encoding. A novel HDR file format has been developed to address the need for higher compression as well as compatibility with existing software solutions. A full description of this file format can be found in [6]. The format, called JPEG-HDR or MPEG-HDR for video version, offers very high compression factors and generally produces encoded HDR images which are comparable in size to low dynamic range images. Moreover, the format is fully backwards compatible and images will appear nicely tone-mapped on conventional viewers capable of opening JPEG files. As such it meets the requirements of compatibility, comparable cost (size) and independence.

Display

The final element in a full HDR imaging pipeline is a display device. This can be a front view display or projection system. In both cases conventional display technology is incapable of delivering the brightness, contrast and amplitude resolution available to our eyes in the real world. Novel HDR display and projection systems have been developed to overcome this limitation. Both devices are based on similar concepts of dual modulation.

Conventional displays generally produce light at a central source (e.g. projector lamp, LCD backlight) and then modulate this light through various means (e.g. DLP chip, LCD glass). Alternatively, some display technologies such as plasma displays create light at different intensities at the pixel level. Our display technology combined two such modulation steps to achieve a multiplication of the amplitude resolution and contrast of conventional displays [7,8].

In the production version of this technology a low resolution array of light emitting diodes (LED) replaces the normally uniform backlight of an LCD. The LED array is modulated to produce a low resolution copy of the desired image and the LCD superimposes the high resolution color information together with a correction pattern to compensate for the lower resolution of the LED array. Together the two layers produce a very high contrast image due to the multiplicative nature of serial modulation. Assuming LED modulation over a range of 400:1 and LCD contrast of 500:1 then the combined system delivers a dynamic range of 200,000:1. Similarly, the amplitude resolution steps are multiplied to go from 8-bit per channel to 16-bit per channel (with some limitations on achievable steps in the top end of the range where our eyes cannot perceive small step sizes in any case). Once such a high dynamic range is achieved, the brightness of the LED backlight can be increased without losing a very low black level. BrightSide's DR37 display uses this design to achieve a brightness of over 3000cd/m2 (10x higher than most conventional displays) and a black level of zero (where LEDs emit no light at all).

The projection alternative to the above design adds a second low resolution monochrome image modulator (e.g. a second projection LCD) to fulfill the same role as the LED array in the HDR display. Very low resolution monochrome modulators can achieve high transmission (especially if the polarization of the two modulators are aligned in the case of an LCD projector) so that the efficiency of the HDR projector is comparable to conventional systems. The same contrast and amplitude resolution benefits as in the HDR display design apply.

Figure 2: BrightSide DR37 high dynamic range display (note that print cannot show the full dynamic range of the display)

In both cases software algorithms have been developed to allow the display of legacy data (e.g. TV signal). Such algorithms extend the dynamic range of conventional data, smooth any visible steps and fill gaps in the data which have been lost due to saturation of the original camera limitation. As a result legacy content will look significantly better on an HDR display than on a conventional display while HDR content will truly shine [9].

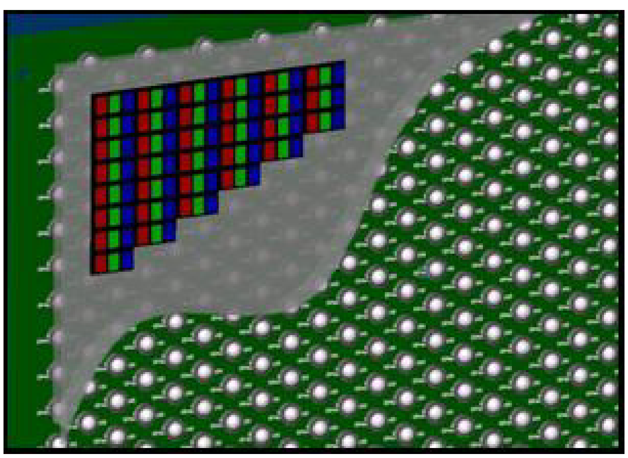

Figure 3: Concept drawing of high dynamic range display layout including LED array and LCD layer.

Benefits

The combined HDR imaging pipeline offers significant benefits in both general and speciality applications. In the latter case applications such as medical imaging, satellite survey and film post-production already employ HDR content and lack the ability to efficiently display (or often store and process) such data. Since the HDR imaging pipeline can effectively emulate the lower performance of most other displays or cameras, it can also be used in comparative work such as quickly confirming whether edited DVD content will look good on various choices of home entertainment devices (LCD, CRT, Plasma, etc).

Yet, the ultimate goal of a complete HDR imaging pipeline is the home entertainment market. From the early days of TV we have always tried to open a window from our living room to the world. An HDR imaging pipeline can make a major step into that direction and provide viewers with highly realistic representations of the world around us.

Performance Metrics

The usual performance metrics used in the display industry were developed with conventional (output-referred) imaging pipelines in mind. As such they are often inappropriate for emerging technologies. The most common confusion in this area arises then the fundamental metrics of brightness, contrast and amplitude resolution are de-coupled (as they are in most measurement standards). Since the metrics are clearly linked in terms of our perception of images, this often leads to very large difference performance specifications on devices that are visually very similar. For example, several emerging display technologies (e.g. SED, OLED) have claimed 100,000:1 or even 1,000,000:1 contrast ratios and yet visually look little different than a good 1000:1 display. The reason for this is that generally these displays have comparable or even lower brightness to conventional display so that the added range is entirely in the dark levels. A 300cd/m2 SED display for example might very well have a 0.003cd/m2 black level (and thus as 100,000:1 contrast) but in your average living room it will look if anything less dramatic than a good LCD TV with much lower contrast (~1000:1) but higher brightness (~600cd/m2 for newer LCD TV). In general, ambient lighting conditions will be bright enough that these extreme black levels are less important than higher brightness. Of course a very good black level is important for good image quality at comparable brightness but at some point black is black and remains (visually) black no matter how much lower the system can go. The following describes the metrics of contrast and amplitude resolution as commonly used in the industry and points out challenges.

Contrast

Contrast is usually quoted as the ratio of the brightest to the darkest achievable level of a display. Clearly, such a metric becomes meaningless if the black level is at zero luminance. At first glance it might be more appropriate to use the second achievable level instead of the first to avoid the division by zero but even then the HDR display yields astronomical contrast numbers. With a high end LCD such as the one used in BrightSide's DR37 a contrast ratio of 1000:1 is achievable on the LCD alone. The LED array is linearly controlled by 8-bit drivers so the second lowest black level of the DR37 is still less than 0.01cd/m2 (and thus the contrast according to this revised metric would be in excess of 300,000:1). Such numbers of course do not reflect a meaningful comparison. While an HDR display has indisputably higher contrast than any conventional display, the visual effect of the difference is unlikely to be 300x better than the conventional LCD alone.

Amplitude Resolution

Amplitude resolution (often called bit depth) refers to the number of distinct steps of luminance that a display can achieve. For most of displays and digital cameras today the amplitude resolution is 8-bit (256 steps) with some higher end solutions allowing for 10-bit (1024 steps). With the constant push for higher brightness in the industry the 8-bit standard becomes a real limitation. Distributing 256 or even 1024 steps over a high contrast/brightness range such as the one found in higher end LCD TV leads to banding and other perceivable artefacts. Attempting an 8-bit distribution over the even larger range of an HDR display is even less desirable. As a result most display designs use spatial and temporal dithering as well as a host of other techniques to give the impression of higher amplitude resolution. Add to this that most display technologies have inherently different response curves and the conventional language of amplitude resolution quickly becomes meaningless. Two displays might have the same 8-bit resolution but could look vastly different if different response curves and dithering techniques are used.

New Metrics

The extreme performance specifications of the HDR display under the conventional metrics create an opportunity to advocate meaningful new metrics for display quality and break away from basing new standards merely on the newest display design introduced by the manufacturer proposing the standard. Such a scene-referred standard should take into account the nature of the world around us as well as our ability to perceive such an environment. Such an approach is quite common in specialised imaging environment such as medical imaging or computer graphics where it is acknowledged that ultimately what matters is how the original content looks and how much of it the viewer can see. For such a metric the notions of contrast, brightness and amplitude resolution need to be merged into a general concept of Number of Distinguishable Grays (NDG) in order to establish how much distinctly visible information is available to the viewer. An outline of such a proposal can be found as a white paper on www.brightsidetech.com [10].

Conclusion

A new high dynamic range imaging pipeline has been developed to allow capture, process, storage and display of high fidelity images. The pipeline provides several orders of magnitude improvements in contrast, brightness and amplitude resolution over any conventional display technology while supporting the same infrastructure. All elements of the imaging pipeline are independently suitable for interaction with conventional devices and fully backward compatible.

With the availability of such an imaging pipeline it is necessary to rethink conventional image quality metrics which have been developed for the limit capabilities of today's cameras and displays. The basic problems of current standards and a brief outline of a possible solution have been provided.

References

1. Reinhard, E., G. Ward, S. Pattanaik, P. Debevec, High Dynamic Range Imaging, Morgan Kaufmann 2005.

2. Paul E. Debevec, Jitendra Malik, "Recovering High Dynamic Range Radiance Maps from Photographs," Proceedings of SIGGRAPH 97, Computer Graphics Proceedings, Annual Conference Series, pp. 369-378 (August 1997, Los Angeles, California). Addison Wesley. Edited by Turner Whitted. ISBN 0-89791-896-7.

3. F. Kainz, R. Bogart, D. Hess, "The OpenEXR Image File Format," SIGGRAPH Technical Sketches, 2003. (See also www.openexr.com.)

4. G.W. Larson, "LogLuv encoding for full-gamut, high-dynamic range images," Journal of Graphics Tools, 3(1):15-31, 1998.

5. G. Ward, "Real Pixels," In Graphics Gems II, edited by James Arvo, Academic Press, 80-83, 1991.

6. G. Ward, M. Simmons, "Subband Encoding of High Dynamic Range Imagery," First Symposium on Applied Perception in Graphics and Visualization (APGV), 2004.

7. H. Seetzen, L. A. Whitehead, G. Ward, "A High Dynamic Range Display System Using Low and High Resolution Modulators", Proceedings of the 2003 Society for Information Display Annual Symposium, May 2003

8. H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. A. Whitehead, M. Trentacoste, A. Gosh, A. Vorozcovs, "High Dynamic Range Display Systems", ACM Transactions on Graphics, special issue on Proc. of ACM SIGGRAPH 2004

9. P. Ledda, A. Chalmers, T. Troscianko, H. Seetzen, "Evaluation of Tone Mapping Operators using a High Dynamic Range Display", ACM Transactions on Graphics, special issue on Proc. of ACM SIGGRAPH 2005

10. G. Ward, "Defining Dynamic Range", White Paper, www.brightsidetech.com , May 2005