

High Dynamic Range Display Systems

Helge Seetzen^{1,2}, Wolfgang Heidrich², Wolfgang Stuerzlinger³, Greg Ward¹, Lorne Whitehead²,
Matthew Trentacoste², Abhijeet Ghosh², Andrejs Vorozcovs³

¹) Sunnybrook Technologies, ²) The University of British Columbia, ³) York University




Figure 1: From left to right: our projector-based display showing an HDR image; our LED-based HDR display showing a discrete and a smooth intensity ramp (the top half of the discrete ramp and the bottom half of the smooth ramp have each been covered by a 1% transparent filter to illustrate high luminance content on the left side of the image, which cannot be captured by the camera); a color-coded original HDR image; HDR photograph taken off the screen of our projector-based system; HDR photograph taken off a conventional monitor displaying the tone-mapped image.

Abstract

The dynamic range of many real-world environments exceeds the capabilities of current display technology by several orders of magnitude. In this paper we discuss the design of two different display systems that are capable of displaying images with a dynamic range much more similar to that encountered in the real world. The first display system is based on a combination of an LCD panel and a DLP projector, and can be built from off-the-shelf components. While this design is feasible in a lab setting, the second display system, which relies on a custom-built LED panel instead of the projector, is more suitable for usual office workspaces and commercial applications. We describe the design of both systems as well as the software issues that arise. We also discuss the advantages and disadvantages of the two designs and potential applications for both systems.

Keywords: Hardware – Novel Display Technologies; Rendering – Perceptually Based Rendering; Image and Video Processing – Image Processing; Methods and Application – Signal Processing; Hardware – Framebuffer Algorithms.

CR Categories: B.4.2 [INPUT/OUTPUT AND DATA COMMUNICATIONS]: Input/Output Devices—Image display; I.3.3 [COMPUTER GRAPHICS]: Picture/Image Generation—Display algorithms; I.3.4 [COMPUTER GRAPHICS]: Graphics Utilities—Device drivers; I.4.0 [IMAGE PROCESSING AND COMPUTER VISION]: General—Image displays.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2004 ACM 0730-0301/04/0800-0760 \$5.00

1 Introduction

In the past few years, the limited dynamic range of both imaging devices and displays has received a lot of attention in the computer graphics community. Algorithms have been developed for capturing both photographs [Mann and Picard 1994; Debevec and Malik 1997; Robertson et al. 1999; Mitsunaga and Nayar 1999] and videos [Kang et al. 2003] with extended dynamic range.

Simultaneously, tone mapping operators have also been developed for compressing the dynamic range so that the images can be displayed on the familiar 8 bit/channel displays with typical contrast ratios of about 300 : 1, including any conventional Cathode Ray Tube (CRT), Liquid Crystal (LCD), and projector-based display. While these tone mapping operators (e.g. [Schlick 1994; Larson et al. 1997; Tumblin and Turk 1999; Durand and Dorsey 2002] among others) allow for displaying high-dynamic-range (HDR) images in a recognizable and even aesthetically pleasing way, nobody would confuse a photograph rendered in this fashion with, say, watching the same scene through a window. The dynamic range of conventional displays is simply insufficient to create the optical sensation of watching a real sunset or driving a car into oncoming traffic at night. Note that this is not just an issue of top intensity: simply increasing the brightness of a conventional display would wash out the dark tones and turn them into a medium gray. What is needed is a significant expansion of the contrast or dynamic range of the display.

In this paper we describe two alternative designs for HDR display systems. We have built prototypes of both, and discuss both the optical design, and related software issues such as display calibration and the rendering of HDR images on both displays.

Both display systems are based on the fundamental idea of using an LCD panel as an optical filter of programmable transparency to modulate a high intensity but low resolution image from a second display. For example, assume we have any display with a contrast range of $c_1 : 1$ between the darkest and the brightest intensity pro-

ducible by that display. If we now put an LCD panel with a contrast ratio of $c_2 : 1$ in front of the first one, then the (theoretical) contrast of the combined system is $(c_1 \cdot c_2) : 1$. In practice, the first display needs to be able to produce a very high intensity image, because color LCD panels only have a transparency of about 3-8%, even when switched to 'white', so that most energy is actually absorbed. Another reason for using a display with a very high base intensity is that a lot of the HDR images we would like to show have, in fact, very bright regions in them.

Based on this principle, we have derived two alternative designs for HDR displays. In the first design (Section 4), a video projector based on Digital Light Projector (DLP) technology serves as the base display. In this version, we directly focus the projector onto the back of the LCD panel. Since the illuminated area is much smaller than during regular use of a projector, the light density is dramatically improved, yielding the high top intensities that we are aiming for. While this design works well in a laboratory setting, it has several drawbacks that restrict its use for a wider class of applications. In particular, these are a large form factor, significant power consumption and heat development, as well as calibration issues.

To overcome these issues, we have devised a second design (Section 5), in which the projector is replaced with a low-resolution array of ultra-bright LEDs. The intensity of every LED can be programmed individually, yielding a low resolution version of the desired image. High frequency features are introduced by attaching a high-resolution LCD panel to the front of this LED array, and adjusting its transparency accordingly. This design makes use of results from psychophysics, which show that very high contrast, although important on a global scale, cannot be perceived by humans at high spatial frequencies (see Section 3).

The two displays we developed have dynamic ranges well beyond $50,000 : 1$, and a maximum intensity of $2700\text{cd}/\text{m}^2$ and $8500\text{cd}/\text{m}^2$, respectively. This compares to a typical dynamic range of about $300 : 1$, and a maximum intensity of about $300\text{cd}/\text{m}^2$ for a typical desktop display.

The design of both systems and their appropriate rendering algorithms, as well as the advantages and disadvantages are detailed in Sections 4 and 5. In Section 6 we discuss possible applications for our display technology before we conclude with some remarks and future research directions in Section 7.

2 Related Work

The class of image processing techniques for coping with the discrepancy between real world luminances and those that fit within the limited gamut of a conventional output device is collectively called *tone-mapping*. Tumblin and Rushmeier [1993] introduced this concept to computer graphics, though their early work did not address dynamic range limitations per se. The first tone-mapping operator to tackle dynamic range reduction was Chiu et al. [1993], who used a spatially varying exposure ramp over the image. However, this approach led to disturbing "reverse gradients" typically seen as halos around light sources. Later work by Larson et al. [1997] returned to a global operator for dynamic range reduction based on histogram adjustment to avoid these artifacts, with local variations to simulate disability glare due to high contrast boundaries in a scene. Pattanaik et al. [1998] developed what some researchers consider the ultimate still image operator based on the human visual system, incorporating color adaptation, local contrast, and dynamic range. However, even this operator exhibited some reverse-gradient effects near high contrast boundaries due to its local spatial filters, leading other researchers to take a different approach.

The basic challenge for a spatially varying tone-mapping operator is that it needs to reduce the global contrast of an image without

affecting the local contrast to which the human visual system (HVS) is sensitive. To accomplish this, an operator must segment the HDR image, either explicitly or implicitly, into regions the HVS does not correlate during dynamic range reduction. The first researchers to successfully accomplish this in an automatic tone-mapping were Tumblin and Turk [1999] with their LCIS operator. However, LCIS sometimes produces odd-looking images, which bear little relation or resemblance to the original scene brightnesses. More recent operators by Ashikhmin [2002], Fattal et al. [2002], Reinhard et al. [2002], Durand & Dorsey [2002], and Choudhury & Tumblin [2003], are much more successful in separating contrast differences that matter to vision from those that do not.

Regrettably, none of these techniques has been validated or supported by psychophysical research, so the resulting images remain as the only evidence that these methods have any real validity for reproducing our experience of an HDR scene on a low dynamic range device. Most of the methods contain free parameters that must be set by the user based on preference. Usually, little can be said about the visual impact that these parameters have on the image in terms of visibility, contrast, brightness, or human visual response in general.

HDR displays offer two benefits in this area. The first, immediate benefit is to researchers, who until now have had no means for the controlled display of dynamic HDR imagery in their studies. Our HDR displays are already helping researchers to test out their hypotheses regarding the effects of tone-mapping on HDR scenes. The longer term benefit will be felt when HDR displays penetrate the professional and eventually the consumer markets, reducing or eliminating the need for tone-mapping. Although we expect tone-mapping to continue to be a requirement for many types of output, such as hard copy, the availability of HDR displays will likely reduce the need for dynamic range compression for many critical applications in the years to come.

3 Remarks on Human Perception

The human visual system has tremendous capabilities but also some limitations. Some of these limitations are an integral part of the theory underlying the HDR devices presented in this paper and the following sections will describe these in detail. In general, our eye has evolved to deal with the vast dynamic range available to us in our daily environment, ranging from starlight to sunlight over at least an eight order of magnitude luminance range. To cope with this range, the eye uses a complex adaptation system. For the purpose of this paper, we use a simple model of adaptation with two time scales: those mechanisms working at time scale of the order of minutes and those with a shorter time scale. The former are of little interest from the point of view of HDR display development. The latter are very interesting as they are the primary reason why current displays cannot provide realistic representations of real world HDR scenes. The eye can capture approximately 5 orders of magnitude of dynamic range effectively simultaneously. No conventional display technology comes close to this. Yet, there are limitations to this capability as described below.

3.1 Local Contrast Perception

While we can see a vast dynamic range across a scene, we are unable to see more than a small portion of it in small regions (corresponding to small angles). Different researchers report different values for the threshold past which we cannot make out high contrast boundaries, but most agree that the maximum perceivable contrast is somewhere around $150 : 1$ [Vos 1984]. Scene contrast boundaries above this threshold appear blurry and indistinct, and the eye is unable to judge the relative magnitudes of the adjacent regions.




Figure 2: The point spread function of the human eye according to Moon&Spencer [1945].

This inherent limitation can be explained locally by the scattering properties of the eye. From Moon & Spencer’s original work on glare [1945], we know that any high contrast boundary will scatter at least 4% of its energy on the retina to the darker side of the boundary, obscuring the visibility of the edge and details within a few degrees of it (Figure 2). If the contrast of an edge is 25 : 1, then details on the darker side will be competing with an equal amount of light scattered from the brighter side, reducing visible contrast by a factor of 2 in the darker region. When the edge contrast reaches a value of 150 : 1, the visible contrast on the dark side is reduced by a factor of 12, rendering details indistinct or invisible.

However, we cannot claim high contrast content has no effect – clearly it does. An observer will notice when one region is much brighter than another, both by the challenge it creates in viewing the boundary, and by the accommodation that goes on when shifting from side to side. When the threshold is very large, observers may even experience discomfort as they attempt to see detail near a bright source, as any driver knows from their nighttime travels. A photographic print of oncoming headlights is merely an allusion to the real experience – it cannot duplicate the visceral experience of glare, or reproduce the effect it has on a human observer. It is exactly this kind of experience that an HDR display can uniquely reproduce.

The HDR display technology described in this paper only exploits the inability of humans to see detail in the immediate vicinity of a high-contrast boundary; it makes no assumptions about our overall response to varying brightnesses. Relative (and even absolute) brightnesses are maintained, and edges will be reproduced exactly when they are below the maximum contrast of the front display – about 400 : 1 in our current prototype. Only when this range is exceeded is some fidelity lost near high contrast boundaries, but this effect is well below the detectable threshold, and has not been visible in any of our experiments [Seetzen et al. 2003].

3.2 Just Noticeable Difference Steps vs. Contrast and Bit Depth

A key question to ask in designing any display system is: how many distinct input/output levels are necessary to cover the desired range without banding or similar quantization artifacts? For conventional displays, this question is often answered by considering a single viewer adaptation level and the number of bits required to represent suitable steps on a particular gamma curve. This may be adequate if the dynamic range being considered is small, but fails when a

display is capable of levels much brighter and much darker than ambient.

To answer this question in the context of a HDR displays, we turn to psychophysical research in the area of Just Noticeable Differences. A JND is the smallest detectable luminance difference at a given luminance level. Visual psychologists have established a complete theory for calculating JNDs over the luminance range relevant for our HDR displays [Barten 1992; Barten 1993], which allows us to establish the useful luminance level of the HDR display. Adding a JND to a particular luminance level effectively defines the next useful step on the luminance scale of the display since it is clearly redundant to provide addressable luminance levels between those two levels if the eye cannot perceive any difference.




Figure 3: The number of just noticeable difference (JND) steps for different maximum intensities according to the model by Barten [1992,1993].

Based on Barten’s original work, an analytical formula for JNDs was derived for a DICOM standard [2001] (see Figure 3). This model predicts 962 JND steps in the luminance range of the projection based HDR display described in Section 4 and 1139 JND in the luminance range of the LED based HDR display (Section 5). Our goal therefore is to reproduce at least this many steps on each display from the darkest to brightest output level, and in both cases our step sizes are maintained well below a JND throughout the luminance range. In the following sections we will use the notion of JNDs instead of contrast and bit depth to characterize the performance of the HDR systems, but continue to use these more familiar terms for the individual components (LCD and projector) that make up the HDR display.

Recently, Muka and Reiker [2002] have argued that, for conventional displays with a typical dynamic range of 300 : 1 or so, an 8-bit representation of images is sufficient for medical diagnosis. They argue that the difference between an 8-bit digital display and a 10-bit or higher bit depth is minimal, and perhaps not noticeable at all. However, as the range of displayable luminances increases, so does the number of JND steps required to cover that range, which is reflected in the numbers presented above.

4 System 1: Projector-based Display

As outlined in the introduction, the first HDR display system we built modulates the image from a projector with a transmissive LCD. This system is detailed in the following.

4.1 Hardware Setup

In a conventional LCD, two polarizers and a liquid crystal are used to modulate the light coming from a uniform backlight, typically a fluorescent tube assembly. The light is polarized by the first polarizer and transmitted through the liquid crystal where the polarization of the light is rotated in accordance with the control voltages applied to each pixel of liquid crystal. Finally, the light exits the LCD by transmission through the second polarizer. The luminance level of the light emitted at each pixel is controlled by the polarization state of the liquid crystal. It is important to point out that LCDs cannot completely prevent light transmission - even at the darkest state of a pixel, light is emitted and as such the dynamic range of an LCD is defined by the ratio between the light emitted at the brightest state and the light emitted in the darkest state. For a high end LCD, this ratio is usually around 300 : 1, with monochromatic specialty LCDs (e.g. those for medical imaging) going up to 700 : 1. The luminance range of the display can easily be adjusted by controlling the brightness of the backlight, but the dynamic range ratio will remain the limiting factor. In order to maintain a reasonable 'black' level of about $1\text{cd}/\text{m}^2$, the LCD is thus limited to a maximum brightness of about $300\text{cd}/\text{m}^2$.

The basic modification introduced by the HDR technology involves inserting a second light modulator and increasing the brightness of the backlight. These two modulators in series provide an extremely dark state with a very low light emission, which then makes it possible to increase the brightness of the backlight dramatically without losing the 'black' state. Optically, this series of modulators results in multiplication of the individual dynamic ranges.

For the projector-based HDR display presented in this paper, the backlight and the first modulator are combined into a single DLP using a Digital Mirror Device with a dynamic range of 800 : 1. The three central components of the HDR display are then the projector, the LCD and the optics that couple the two. Using these components, each image on the HDR display is the result of modulated light coming from the projector which is directed onto the rear of the transmissive LCD by the optics system, modulated a second time by the LCD, and properly diffused for viewing.

The projector used in the HDR display is an Optoma DLP EzPro737 digital mirror projector. To reduce unnecessary light loss, we have removed the color wheel from the projector, resulting in a monochrome display system with a threefold increase in brightness due to the absence of the color filters. New control electronics have been integrated into the commercially available projector to re-synchronize it in absence of this color wheel.

The LCD panel is a 15" XGA color LCD made by Sharp (Sharp LQ150X1DG0). It is driven by an EarthVision AD2 LCD controller, which allows a direct VGA connection. The LCD panel has been separated from the conventional backlight and all of the optical layers behind the display have been removed to create a transmissive image modulator.

The optics used in the HDR display include the conventional projection lens of the EzPro projector, and a Fresnel lens directly behind the LCD display to collimate the projected light into a narrow viewing angle for maximum brightness of the HDR display and to avoid color distortion due to diverging light passing through the color filters of the LCD. Finally, a standard LCD diffuser has been used to redistribute the collimated light into a reasonable viewing angle.

All three components have been installed in a single housing with appropriate alignment mechanisms to create a close matching of the DLP and LCD pixels. The alignment can be fine-tuned through the controls of the DLP projector. However, a perfect match is impractical as alignment at the sub-pixel level is hard to achieve and almost impossible to maintain. To avoid moiré patterns and alignment artifacts associated with even a minor misalignment, we have deliberately blurred the projector image and compensate

for that blur in the LCD image as described in the following section.





Figure 4: Top: schematic of the HDR display including the projector, LCD and optics. Bottom: actual photograph of the display. Both projector and LCD are driven by a single dual-VGA graphics card.

Using this configuration, the light output of each pixel of the HDR display is effectively the result of two modulations, first by the DLP and then by the LCD pixel, along the same optical path. The upper boundary of the dynamic range results from full transmission of both pixels (i.e. the 255th level on both modulators), and the lowest boundary from the lowest possible transmission of both modulators (i.e. the 0th level on both modulators). Since the DLP has a dynamic range of 800 : 1 and the LCD a dynamic range of 300 : 1, the theoretical dynamic range of the HDR display is 240,000 : 1. Imperfections in the optical path introduce noise that reduces the dynamic range to a measured 54,000 : 1. The luminance values matching these boundaries are a result of the brightness of the projector and the transmission of the LCD. In this case, the Optoma EzPro737 is rated at 1200 Lumens, or approximately 2400 Lumens once the RGB color filters are removed (since each filter for red, green and blue eliminates approximately 2/3 of the incoming light). The Sharp LCD panel has a measured transmission of approximately 7.6% in the white state (this is quite high for an LCD since even the theoretical maximum for a color LCD without any losses is only 16% due to the light reduction of 50% at the polarizer and another 66% due to the RGB color filter). Assuming that the light emitted by the HDR display is diffused across a solid angle ω , the maximum luminance is then given by:

$$L_{\max} = \frac{\Phi_{\max}}{A\omega}$$

where A is the area of the LCD and Φ_{\max} is the maximum outgoing flux. In the HDR display prototype, the flux is approximately 182 Lumens (2400 Lumens \times 7.6%). The area A is the area of the

15" LCD (697cm^2) and the solid angle of diffusion ω is approximately 0.66sr (40° diffusion horizontally, 15° vertically). The maximum luminance for this particular configuration is then approximately $3,956\text{cd}/\text{m}^2$. We actually measured a luminance of $2,700\text{cd}/\text{m}^2$ Lumens. The theoretical minimum luminance is less than $0.01\text{cd}/\text{m}^2$, while our measurements resulted in a value of $0.05\text{cd}/\text{m}^2$. Clearly, a shift of this range toward even higher luminance values would be possible with a brighter projector or a more transmissive LCD. Unlike a standard low dynamic range display, even an order of magnitude increase of the maximum luminance would not significantly reduce the quality of the 'black' state since $1\text{cd}/\text{m}^2$ is still a very satisfying 'black', especially if other parts of the image contain very high luminance values.

Within that luminance range, a very large number of different combinations of output settings for the DLP and LCD can be achieved. If both systems were linear 8-bit devices then the total number of combinations would be 256^2 , over 17,000 of which are distinct. Due to the non-linear gamma of each system, the actual range of distinct addressable steps is different, but still significantly larger than what is needed to display the 962 JND steps necessary to provide all visible and distinguishable luminance steps in the measured luminance range of the system (including all losses) of $0.05\text{cd}/\text{m}^2$ to $2,700\text{cd}/\text{m}^2$.

4.2 Driving the Projector Display

To correctly render HDR images on this display, we need to analyze the image formation process of the system. Let us assume for the moment that both the projector and the LCD panel are perfectly linear, and that both have the same dynamic range. For now, let us also neglect the blurring of the projector image. Under these assumptions, we can achieve the target intensity by normalizing the intensity range of the HDR image to $0 \dots 1$, and using the square root of this normalized intensity to drive both the projector and the LCD panel. This even split between pixel values on the projector and the LCD panel is preferable to a scenario where one value is very large and the other is very small, since quantization artifacts are relatively larger for small values.

In reality, neither the projector nor the LCD have a linear response, and we also need to compensate for the blurring of the projector image. We do this in the following way: we choose a simple estimate of what the projector intensity should be, and then simulate the effect of response function and blurring. Finally, the pixel values of the LCD panel are chosen such that they compensate for these effects.




Figure 5: Rendering algorithm for the projector-based display.

The complete rendering algorithm then works as follows (also see Figure 5): we take the square root of the original HDR image with intensity I (1). The resulting image (2) represents the target intensity \sqrt{I} for the projector. We map these intensities into projector pixel values by applying the inverse of the projector's response function r_1 (3). The projector now produces an image of intensity $r_1(r_1^{-1}(\sqrt{I})) = \sqrt{I}$, except that the image is actually blurred according to p_1 , the projector's point spread function (PSF). To simulate this blurring, we convolve the projector intensities with the PSF (4) and divide the result out from the original HDR image to get the

target LCD transparency (5). For the final pixel values of the LCD, we apply the inverse of the panel's response function r_2 (6).

An exposure sequence of the PSF p_1 is depicted in Figure 6. Note the vertical lines visible in images with larger exposure times. These are the RGB subpixels of the LCD panel. In order to speed up the computation, we do not use the measured PSF directly, but fit a tensor-product Gaussian to p_1 . We usually set the focus of the projector such that the fitted Gaussian has a standard deviation of about 2-3.5 pixels, so that we can use a 2D separable filter of width 13 for the convolution.




Figure 6: Point spread function of the projector.

These image processing steps are possible both completely in software and using recent programmable graphics hardware (graphics processing unit, GPU). The convolution is the most difficult part of the GPU implementation, but with the separable approximation it can be readily implemented as a pixel shader on both the latest ATI and NVIDIA chips. With this approach we achieve frame rates of 30 frames per second or more on current hardware.




Figure 7: Response function of both the LCD panel and the DLP projector in the projector-based display.

Figure 8 shows the results of this image factorization on a portion of Debevec's Stanford Memorial Church HDR photograph. On the left side, you can see a grayscale image that corresponds to the square root of the original intensity values. Convoluting that image with the PSF of the projector yields the center image. This is the predicted image that will be produced by the projector. Finally, the right image is the color LCD panel image that corrects for the blurriness of the projector. It is interesting to note that the LCD panel image is essentially an edge-enhanced image with low frequency components attenuated or removed. This is particularly noticeable for the widths of the window frames. Interestingly, the algorithm that we apply here is very similar in principle to a local tone mapping operator. This means that our LCD panel image is almost a tone mapped version of the original HDR image, although our method is clearly not designed for that purpose.

4.3 Discussion of the Projector Display

The projector based HDR display hardware provides a tool to present high quality HDR images but has several drawbacks. In ad-




Figure 8: Factoring an HDR photograph for our projector-based display. Left: square root of the intensity. Center: blurred image which is predicted to be the image generated by the de-focused projector. Right: edge enhanced LCD panel image that corrects for the blurriness of the projector image.

dition to the obvious form factor problem due to the optical length required by the projector, the power consumption, cost, thermal management and video bandwidth requirements are high compared to a conventional display.

High power consumption and the resulting thermal management requirements are a consequence of the image creation mechanism inside the projector. Unlike a cathode ray tube (CRT) display, where light is created only in the regions of the image that are supposed to be bright, an LCD or DLP projector creates a uniform light distribution that is then modulated by the LCD or DLP mirror chip. The power consumption of an LCD or DLP projector is thus independent of the image and always very high as there has to be enough light produced by the lamp such that a full screen ‘white’ can be shown. Combined with the low modulation efficiency of the LCD or DLP this causes the high power consumption. In the HDR display the situation is worse than in a conventional, single-modulator display. The lamp of the projector has to emit enough light to allow a full screen image at the highest possible brightness of the HDR display. To achieve $10,000\text{cd}/\text{m}^2$ on a 15” screen we would need an outgoing flux of approximately 500 Lumens (see Section 4.1). Even with a very high transmission LCD this requires at least 5000 Lumens to be emitted from the projector. In the prototype presented in Section 4.1 the color wheel/filter of the projector has already been removed to reduce the losses in the projector but even so the modulation efficiency of the projector is slightly less than 50%. The lamp thus has to produce in the order of 10,000 Lumens. Yet, in almost all HDR images the area that is actually at such a high brightness of $10,000\text{cd}/\text{m}^2$ is very small. In fact, a random selection of 100 HDR images indicated that average HDR images have less than 10% of the image content in the high luminance range (above $3000\text{cd}/\text{m}^2$) and that the average luminance over all images was less than $800\text{cd}/\text{m}^2$ for indoor scenes and $2,100\text{cd}/\text{m}^2$ for outdoor scenes. The projector HDR display consequently creates a factor of between 12.5 and 4.75 too much light at any given time.

The projector is also the cause of the high bandwidth requirements. Even though the image projected by the projector onto the back of the LCD is blurred, the projector itself is still a high resolution display which requires high resolution input data. As a result, the projector-based HDR display needs a high resolution video stream going to the LCD and a similar size video stream to the projector. This creates a requirement for a dual output graphic card and imposes limits to the frame rate of the display due to the computational requirements.

Finally, the cost of a high brightness projector is very high which makes this version of the HDR display unsatisfactory for commercial purposes. Cost also presents a barrier to larger screen sizes as the brightness requirements increase linearly with area, and the cost curve for projectors is very steep with brightness (while the step from 2,000 Lumens needed for a 15” display to a 4,000 Lumens projector for a 20” display is only twice as high, the next step

from 20” to 40” TV size would require a 15,000 Lumens projector at a price that is over 20 times higher than that of a 4,000 Lumens projector).

Yet, for research applications, the display is a valuable tool. The high cost is in part due to the use of fully finished consumer products instead of individual components, but this also makes it possible to assemble the system without significant development of custom electronics. Since the drawbacks mentioned above in no way diminish the actual image quality, the projector-based HDR display provides researchers with a relatively simple to build solution with very high image quality.

5 System 2: LED-based Display

As seen in the discussion of the projector-based HDR display in Section 4.3 there are significant obstacles to overcome.

To realize the dream of television or computer displays presenting images that look indistinguishable from the real world, it is not sufficient to merely show images with the appropriate luminance range and resolution; it is also necessary to make a commercially viable system that achieves these higher quality images within today’s hardware and software infrastructure and market price points. The version of the HDR display described in this section retains the high image quality of the projector display and overcomes the commercialization barriers: power, thermal, cost, form factor and infrastructure demands on the graphics card.

5.1 Hardware Setup

We have already seen in Section 4.1 that software correction can compensate for a low resolution of the rear image of the HDR display. It is important to realize that this correction works perfectly, as long as the local image contrast does not exceed the dynamic range of the front modulator. From the psychology theory presented in Section 3 we can establish the largest size of a rear image pixel. For this version of the HDR display we have used light emitting diodes (LED) at the largest possible size allowed by the veiling luminance effect which has been validated previously through experimental tests [Seetzen et al. 2003].

A prototype has been build using 12mm Lumiled Luxeon 1 Watt white LEDs (LXHL-PW01) on a hexagonal close-packing matrix where each LED is individually controlled over its entire dynamic range with 1024 addressable steps. 760 LEDs have been mounted behind an 18.1” L.G. Philips LCD with a 500 : 1 dynamic range and 1280×1024 resolution. On full screen white/black the maximum luminance is measured as $8,500\text{cd}/\text{m}^2$ and the minimum luminance is zero, since all LEDs are off. The minimum luminance is less than $0.03\text{cd}/\text{m}^2$ on a checkerboard pattern larger than 20mm. As mentioned in Section 3.2, the Barten model [Barten 1992; Barten 1993] predicts 1139 JNDs for this luminance range. The system is capable of displaying images at video rates.

5.2 Driving the LED Display

The principal rendering algorithm for the LED-based system is quite similar to that for the projector-based display, as shown in Figure 9. The primary difference between the two display systems from a rendering perspective is that the PSF of an LED has a much wider support than the one for a pixel of the projector. Also of importance is the fact that the LEDs are arranged on a hexagonal grid rather than a rectangular grid. These differences have two consequences. Firstly, because of the wider support of the PSF, it is advisable to come up with a better way to choose the LED values. Since the supports of the PSFs for neighboring LEDs overlap, determining the optimal LED value is essentially a de-convolution

problem, as explained below. Secondly, because of both the hexagonal geometry and wider support of the PSF, the convolution (4) has to be implemented differently.




Figure 9: Rendering algorithm for the LED-based display.

We address the first issue by adding an additional stage (2a) to derive the target intensities I_L for every individual LED. To this end we first down-sample the image to the resolution of the LED array, and then solve for the values, taking the overlap of the PSFs into account. This is essentially a de-convolution problem, the full solution of which would require solving a sparse linear equation system with as many unknowns as there are LEDs. This is not an option for interactive applications, and furthermore de-convolution algorithms are known to be numerically unstable. We can approximate a solution with a single Gauss-Seidel iteration over neighboring LED pixels. This amounts to a local weighted average of neighboring LED target values, where some of the weights are negative. In our design, these LED values (2a) can be forwarded directly to the controller electronics, which correct for nonlinearities in the LED response in hardware.

As before, we rely on the LCD panel to compensate for any differences between the LED values and the target image. To this end, we need to forward-simulate the low-frequency image (4) generated by the LED panel in order to derive the LCD pixel values. We use two different approaches for software and hardware implementations. With GPUs, we use a splatting approach, and simply draw screen aligned quadrilaterals with textures of the PSF into the framebuffer. Alpha blending is used to accumulate the results. In software, we approximate the measured PSF with a Gaussian, and implement the reconstruction by convolving a low resolution image with the approximate PSF, and then up-sampling to the full resolution.




Figure 10: Response function of the LCD panel in the LED-based display. The response of the individual LEDs (not shown here) is designed to be linear.

Figure 11 shows the factorization of an HDR image into LED component and LCD panel component. Due to the wider support

of the PSF of the LEDs compared to the PSF of a projector pixel in the first setup, the LED image is even more low-pass filtered than before. As a result, the compensation performed in the LCD panel is more significant, and is visible in the right image.



Figure 11: Factorization of an HDR photograph for the LED-based display. Left: LED contribution as a result of convolving LED values with the LED point spread function. Right: edge enhanced LCD panel image that corrects for the blurriness of the LED image.

5.3 Discussion of the LED Display

The use of LEDs overcomes the power consumption and thermal problem as a result of intelligent light production, and it eliminates other commercialization barriers outlined in Section 4.3. Form factor is no longer an issue as the LED matrix can be the same thickness as a conventional LCD backlight. The video bandwidth requirements are dramatically reduced due to the lower resolution of the LED matrix. As a result, only a few hundred 8-bit values are needed over a standard image, and these are added to the front of the DVI video stream going to the LCD and stripped off by the controller hardware inside the display. Cost remains an issue but significantly less so than with the projector system.

One slight disadvantage is that the rendering algorithm for the LED-based display is computationally more demanding than the algorithm for the projector-based display due to the larger support of the PSF. On a GeForce FX we currently achieve about 10 fps. for a fullscreen image factorization into LCD and LED components. The rendering time is mostly limited by having to work around limitations in the support of floating point framebuffers and textures on current GPUs. The next generation of GPUs will support a more complete set of operations, including blending and bi-linear interpolation, which should improve the performance of our algorithm by a factor of 3-4.

6 Applications

We have developed four simple applications to test our display technology, and to demonstrate its potential in a number of application domains. The first one is a simple HDR image viewer that works with both displays (Figure 12, top left). It allows the user to load an HDR image and show it while interactively adjusting the exposure settings (i.e. the absolute scale of intensity). This application was also used to generate the images in Figure 1. The three images to the right of that figure show a color-coded comparison of an original radiance map, an HDR photograph taken off our projector-based display, and finally a photograph taken off a standard LCD screen. The intensities of the photograph from the HDR display are similar but not identical to the values in the original radiance map. The differences are mostly due to imperfections in both the calibration of the display to absolute intensities, and in the image acquisition process. Clearly both intensity and dynamic range of our display are vastly superior to the standard monitor.




Figure 12: Screen photographs of the different applications we implemented. The exposure times of the two images in each pair differ by 4 stops. Top left: HDR image viewer. Top right: interactive rendering of measured BRDFs. Bottom left: volume rendering. Bottom right: medical image viewer.

The second application we developed is related to interactive photorealistic rendering. We modified a DirectX application for displaying BRDFs measured with linear light source reflectometry [Gardner et al. 2003], and replaced its tone mapping step with the rendering algorithm for our display (Figure 12, top right). Other interactive applications that can render into floating point buffers can easily be modified in a similar fashion. As pointed out in Section 5.2, the rendering performance on current GPUs is limited due to missing features such as alpha blending with floating point framebuffers and lack of bi-linear interpolation for floating point textures. These issues should be resolved in upcoming generations of GPUs.

Based on a similar principle, we have built a simple volume ray-caster that runs on a GPU (Figure 12, bottom left). It allows for operations such as rotation, slicing, and adjustments to the transfer function. Both the actual volume rendering algorithm and the processing for our displays is implemented in a single Cg shader [Mark et al. 2003].

Finally, a simple medical imaging viewer and browser has been implemented to allow radiologist to view medical images in high dynamic range (Figure 12, bottom right).

7 Conclusions and Future Work

In this paper we have presented two designs for HDR displays, one based on a projector setup and one based on an LED array. The two displays we developed have dynamic ranges of well beyond $50,000 : 1$, and a maximum intensity of $2700\text{cd}/\text{m}^2$ and $8500\text{cd}/\text{m}^2$, respectively. We have described both hardware and software aspects of these display systems, and have described a number of applications for this kind of technology.

There are plenty of opportunities for future work. On the hardware side, the image quality of the LED-based HDR display can be further enhanced by replacing each white LED with a triplet of red, green and blue LEDs or a larger number of colored LEDs if more primaries are desired (e.g. red, blue, blue-green, yellow-green). Color LEDs have very sharply defined spectral emission pattern as opposed to the broad band emission of fluorescent tubes

used in conventional LCD backlights. This creates pure primaries and allows for a significant increase in the display color gamut. A conventional LCD achieves approximately 66% of the NTSC color gamut while a tri-color LED based system can achieve 98% with additional red and blue outside the NTSC gamut [Ohtsuki et al. 2002].

The actively controlled matrix of LEDs can also be used to very easily implement backlighting schemes proposed by the industry such as flashing the backlight in segment in sync with the LCD refresh to reduce motion blur artifacts [Fisekovic et al. 2001]. Likewise, the LED array simplifies problems of global luminance non-uniformity found in conventional displays resulting from a non-uniform light output of the fluorescent tube backlight, lifetime problems due to backlight failure, and many other limitations of a conventional LCD backlight.

At this point the calculation of the rear and front image layer are executed by the graphics card but standardized calculations of this kind can easily be done by dedicated hardware inside the display as soon as standards for floating point video signals (such as extensions to the DVI standard) become available. This will reduce the load on the graphic card, allow software to directly send floating point data to the display, and have the display handle the necessary computations for blur correction.

In addition to these directions for further developing the hardware, there are a number of potential applications of this display technology in computer graphics, visualization, human-computer interaction, and perception research. Commercial applications include proofing for the film and special effects industries, since the dynamic range of film significantly exceeds what can currently be shown on standard displays, and eventually even home theater. For volume rendering in particular, the higher dynamic range might open up possibilities for completely new rendering algorithms. We also plan to explore new interaction metaphors, such as use of HDR imaging in user interfaces. The extra brightness should be useful for directing the user's attention to important events. Finally, an interesting research direction is to expand the work on human perception research in order to challenge and expand on assumptions that have been made in the context of traditional display technology. Other researchers have already used our display to run perceptual

comparisons of tone mapping operators.

8 Acknowledgments

We would like to thank Chris Tchou for his work on the initial interactive rendering demo for the linear light source reflectometry data [Gardner et al. 2003]. Paul Debevec made the Daguerreotype reflectance data from the same project available to us (see Figure 12 and the video). The HDR photographs used throughout this paper were also provided by Paul Debevec, while the CT data set was provided by Klaus Engel. Thanks to the anonymous reviewers for their valuable comments.

This work was in part supported by NSERC, IRIS/Precarn, and the BC Advanced Systems Institute.

References

- ASHIKHMIN, M. 2002. A tone mapping algorithm for high contrast images. In *Proc. of Eurographics Workshop on Rendering 2002*, 145–156.
- BARTEN, P. 1992. Physical model for the contrast sensitivity of the human eye. In *Proc. SPIE*, vol. 1666, 57–72.
- BARTEN, P. 1993. Spatio-temporal model for the contrast sensitivity of the human eye and its temporal aspects. In *Proc. SPIE*, vol. 1913-01.
- CHIU, K., HERF, M., SHIRLEY, P., SWAMY, S., WANG, C., AND ZIMMERMAN, K. 1993. Spatially nonuniform scaling functions for high contrast images. In *Proc. Graphics Interface '93*, 245–253.
- CHOUDHURY, P., AND TUMBLIN, J. 2003. The trilateral filter for high contrast images and meshes. In *Proc. of the Eurographics Symposium on Rendering*, 186–196.
- DEBEVEC, P., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proc. of ACM SIGGRAPH '97*, 369–378.
- DICOM. 2001. *Digital Imaging and Communications in Medicine (DICOM)*. <http://www.medical.nema.org/dicom/2001.html>, ch. Part 14: Grayscale Standard Display Function.
- DURAND, F., AND DORSEY, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Trans. Graph. (special issue SIGGRAPH 2002)* 21, 3, 257–266.
- FATTAL, R., LISCHINSKI, D., AND M. WERMAN. 2002. Gradient domain high dynamic range compression. *ACM Trans. Graph. (special issue SIGGRAPH 2002)* 21, 3, 249–256.
- FISEKOVIC, N., NAUTA, T., CORNELISSEN, H., AND BRUININK, J. 2001. Improved motion-picture quality of AM-LCDs using scanning backlight. In *IDW 2001 Proceedings*, 1637–1640.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *ACM Trans. Graph. (special issue SIGGRAPH 2003)* 22, 3, 749–758.
- KANG, S. B., UYTENDAELE, M., WINDER, S., AND SZELISKI, R. 2003. High dynamic range video. *ACM Trans. Graph. (special issue SIGGRAPH 2003)* 22, 3, 319–325.
- LARSON, G. W., RUSHMEIER, H., AND PIATKO, C. 1997. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Trans. on Visualization and Computer Graphics* 3, 4, 291–306.
- MANN, S., AND PICARD, R. 1994. Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. Tech. Rep. 323, M.I.T. Media Lab Perceptual Computing Section. Also appears, IS&T's 48th annual conference, Cambridge, MA, May 1995.
- MARK, W. R., GLANVILLE, R. S., AKELEY, K., AND KILGARD, M. J. 2003. Cg: A system for programming graphics hardware in a C-like language. *ACM Trans. Graph. (special issue SIGGRAPH 2003)* 22, 3, 896–907.
- MITSUNAGA, T., AND NAYAR, S. K. 1999. Radiometric self calibration. In *Proc. of IEEE CVPR*, 472–479.
- MOON, P., AND SPENCER, D. 1945. The visual effect of non-uniform surrounds. *Journal of the Optical Society of America* 35, 3, 233–248.
- MUKA, E., AND REIKER, G. 2002. Reconsidering bit depth for radiological images – is eight enough? In *Proc. SPIE*, vol. 4686, 177–188.
- OHTSUKI, H., NAKANISHI, K., MORI, A., SAKAI, S., YACHI, S., AND TIMMERS, W. 2002. 18.1-inch XGA TFT-LCD with wide color reproduction using high power led-backlighting. In *Proc. Society for Information Display International Symposium*, 1154–1157.
- PATTANAİK, S. N., FERWERDA, J. A., FAIRCHILD, M. D., AND GREENBERG, D. 1998. A multiscale model of adaptation and spatial vision for realistic image display. In *Proc. of ACM SIGGRAPH '98*, 287–298.
- REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. *ACM Trans. Graph. (special issue SIGGRAPH 2002)* 21, 3, 267–276.
- ROBERTSON, M., BORMAN, S., AND STEVENSON, R. 1999. Dynamic range improvements through multiple exposures. In *Proc. of International Conference on Image Processing (ICIP) '99*, 159–163.
- SCHLICK, C. 1994. Quantization techniques for visualization of high dynamic range pictures. In *Proc. of Eurographics Workshop on Rendering '94*, 7–20.
- SEETZEN, H., WHITEHEAD, L., AND WARD, G. 2003. A high dynamic range display using low and high resolution modulators. In *Society for Information Display International Symposium Digest of Technical Papers*, 1450–1453.
- TUMBLIN, J., AND RUSHMEIER, H. 1993. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications* 13, 6, 42–48.
- TUMBLIN, J., AND TURK, G. 1999. LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Proc. of ACM SIGGRAPH '99*, 83–90.
- VOS, J. 1984. Disability glare - a state of the art report. *CIE Journal* 3, 2, 39–53.