

A Window Manager for High Dynamic Range Display Systems

Andriy Pavlovych

Andrejs Vorozcovs

Wolfgang Stuerzlinger

*York University, Toronto, Canada
{andriyp, av, wolfgang}@cs.yorku.ca*

Abstract

The dynamic luminance range of many real-world environments exceeds the capabilities of current display technology by several orders of magnitude. Recently, new display systems have demonstrated, which are capable of displaying images with a dynamic luminance range much more similar to that encountered in the real world.

The paper summarizes how the human eye perceives high dynamic luminance ranges, sources of high dynamic range data, how the new display systems work, as well as their limitations. The paper discusses the need for a high dynamic range window manager and presents an initial implementation. Finally, the results of a preliminary evaluation are presented.

1 Introduction

1.1 High Dynamic Range content in computer graphics

In the past few years, this issue of limited dynamic range of both imaging devices and displays has been extensively studied in the computer graphics community. In addition to being able to produce such imagery via methods such as physically based rendering [7], algorithms have been developed for capturing both still images [1, 5, 6, 8] and videos [3] of real environments with extended dynamic range.

As the dynamic range of luminance of such real and synthetic images often exceeds the capacity of current displays by orders of magnitude, new approaches to enable their presentation were also developed. One of the ways to display high dynamic range images is to transform the original range of intensities into a significantly smaller range of intensities a common desktop monitor can reproduce. Such process is called *tone mapping* and a number of tone mapping operators

have been developed to date. While these tone mapping operators (e.g. [2, 4, 10, 11] 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 inadequate for creating a visual sensation of watching a real sunset or driving a car into oncoming traffic at night. To ease this problem, a new class of displays has recently been demonstrated [9], which allow for a contrast ratio of more than 50000:1, and have peak intensities in the range of 2700 cd/m² to 8500 cd/m², while lowering the black level to 0.05 cd/m². For comparison, traditional displays usually reproduce a contrast of about 300:1 with a luminance range of approximately 1-300 cd/m².

1.2 High Dynamic Range Display Technology

The principle underlying the devices in [9] is the use of a specialized high-intensity backlight for a transmission LCD panel. In one of the versions, a Digital Light Projector (DLP) was used for that purpose, in another – a grid of high-intensity white light emitting diodes, each of which can be controlled individually. Now, if the maximum contrast of the backlight image is $c_1:1$, and the transmission ratio of the front LCD panel is $c_2:1$, then the theoretical contrast ratio of the system is $(c_1 \cdot c_2):1$. The maximum luminance of such system will increase linearly with the maximum luminous power of the backlight. The reason that the resolution of the backlight image can be lower than the front panel is based on findings from the field of psychophysics, which show that very high contrast, although important on a global scale, cannot be perceived by humans at high spatial frequencies.

Displaying images on such a screen then requires the following technical steps:

- Obtaining a linearly encoded high dynamic range image (*radiance map*).
- Generating the background image.
- Generating the foreground image.

1.3 The Human Visual System

The human visual system is a remarkable apparatus, which allows us to perceive objects under a wide range of ambient illumination, from starlight to daylight, with a resolving power of up to 1'. However, it has several important limitations, which we need to be aware of in the context of displaying high-dynamic range content.

Adaptation Luminance

The human visual system is useful over a wide range of luminance values, and at any given time we can perceive no more than 5 orders of magnitude of dynamic range [4]. With the effect of time-adaptation, this range can be shifted up and down to cover 10 orders of magnitude.

Despite the wide visual field of view of the human eye, it is not possible to observe the whole scene simultaneously. Rather, we sequentially fixate our attention on local areas of the field of view, where the eye rapidly adapts to the average [9] brightness in the neighbourhood of 1–1.5° of visual angle centred at the fixation point. The adaptation luminance determines what part of the overall intensity range the eyes can be sensitive to at that given moment.

Dynamic Range and Local Contrast Sensitivity

Furthermore, there is a limit to how much contrast can be perceived in a very small neighbourhood of the visual field. That is, when the contrast between adjacent spots on the retina exceeds a particular threshold, we will no longer be able to perceive the relative magnitude of that contrast (roughly speaking, the spot on one side will appear white and the one on the other – black). If you separate the spots in space, you will again be able to see their variations in brightness. The threshold at which this occurs, the maximum perceived contrast, is reported to be around 150:1 [9].

Disability Glare

Another major cause of human inability to distinguish detail in areas of high contrast is the phenomenon of *disability glare* [12]. It is caused by light scattering inside the liquid medium of the eye, in the atmosphere, and sometimes the surface of the display. The effect of disability glare is to form a constant *veiling luminance* across a large part of the image area that obscures any detail that has a lower luminance value.

1.4 User Interface Issues

In traditional user interfaces, coupled with traditional displays, user interface elements can afford to have a constant brightness without the danger of becoming poorly distinguishable due to visual interference from the content that is being displayed. This is because the ability of traditional displays to reproduce contrast is not far from the 150:1 threshold. With HDR-capable displays, the dynamic range of the content displayed on the screen can have a substantial effect on the visibility of the user interface elements, and vice versa.

In this paper we demonstrate how the brightness of the non-HDR elements on the high dynamic range display can be compensated to reduce the negative effects of visual glare. We present several requirements that such an adjustment must fulfill and a preliminary implementation of the method on the projector-based version of the HDR display.

2 High Dynamic Range Window Manager

A simple way to combine a high dynamic range image with a low-dynamic range user interface would be to assign a constant average brightness to the LDR content, perhaps matching that of a standard office display (~150 cd/m²). However, for reasons discussed in the previous section, significant visibility problems can arise in cases where windows or interface elements (e.g. icons on a desktop or text in a word processor) are located close to the edge of a window that contains high dynamic range content.

It is non-trivial to decide what the brightness of the user interface should be. If the intensity of the user interface elements is significantly lower than the intensity of the adjacent portion of the HDR window, then these elements will be invisible due to effects of glare in human visual system. On the other hand, if their brightness is too high, they will themselves generate parasitic glare on the other areas of the screen, including the HDR content. Hence, we should limit the brightness of the user interface elements if we wish to make use of the lower end of the display's luminance capability.

In summary, our goal is to maximize the visibility of user interface elements without adversely affecting the presentation of the HDR content.

2.1 Technique

Until fully HDR-aware user interfaces come into existence, we present an implementation of a “HDR window manager”, a background application that is retrofitted to an existing windowing system, and permits an operator to display and manipulate HDR content, as well as use standard, non HDR-aware applications without modification, on the HDR display.

The window manager includes an algorithm that adjusts the relative brightness of different parts of the screen. In practical terms, we would like the adjustment algorithm to have the following properties:

1. It should leave the HDR content unchanged.
2. It should limit the brightness of the user interface elements to avoid light scatter into the HDR image.
3. It should attempt to keep the *local* contrast between the HDR image and the user interface elements to below the local contrast perception threshold (~150:1).

The first requirement stems from the fact that the existing software driving the HDR display is already optimized to deliver the most accurate rendition of HDR content. Even though it would be possible to alter the presentation of the HDR content according to particular viewing conditions, that task falls outside scope of this paper.

The requirement to limit the brightness of user interface elements is explained by the fact that the “bottom end” of the HDR display capability extends to as low as 0.05 cd/m^2 . User interface elements at standard brightness levels (150 cd/m^2) would cause significant glare, which would make the HDR image effectively invisible.

Finally, the third requirement stipulates that the contrast on the boundaries between the HDR windows and the rest of the screen needs to be decreased in order to keep the user interface elements visible.

2.2 Implementation

As mentioned in the introduction, the projector-based HDR display contains two imaging planes, which are optically combined in a multiplicative fashion to obtain a single high-contrast image. The HDR rendering algorithm decomposes the input HDR image into two synthetic images, corresponding to the back and front plane. As a rough approximation, the front plane usually contains the high-frequency image information and the back plane contains low-frequency intensity variations. For the purposes of displaying a LDR user interface, we need to render the user

interface on the front plane while controlling the average intensity using the back plane.

We satisfy the above requirements by manipulating the image on the rear of the two planes of the HDR display. The user interface elements are present on the front surface of the display. The processing of the background image consists of interpolating the intensities outward from the window boundaries, for 5 millimetres of screen distance, until the magnitude reaches the average intensity level already present in the background image. The distance chosen, 5 millimetres, is on the order of the size of the adaptation region, and it allows for a gradual change in background intensity.

2.3 Examples

We demonstrate the ideas presented by considering a problematic high contrast edge that arises when a bright HDR image is located near a page of text in the user interface. Figure 2 shows a simulated picture of how the human visual system would perceive this kind of high contrast edge with and without our correction. Applying the correction improves visibility of the part of the text immediately next to the image by reducing local contrast.



Figure 1. Simulated screenshots of the HDR display

- Top left: No adjustment (flat)
- Top right: Perceived image (too dark)
- Bottom left: Corrected image (brightened)
- Bottom right: Perceived corrected image (flat)

Figure 1 shows images of the projector-based HDR display with and without running the algorithm, taken with a conventional digital camera. Note that these pictures cannot accurately reproduce the effects of light scatter in human vision.




Figure 2. Appearance of text next to HDR content (enlarged)

In the top image, the non-HDR content was not altered. The bottom image illustrates how text legibility was improved around the high-intensity parts of the image as a result of applying the algorithm.

3 Summary and Future Work

In this paper, we have drawn attention to unique challenges in displaying content on high dynamic range displays. We also presented an approach that addresses this challenge via automatically controlling the intensity of the non-foreground elements on the high dynamic range display in order to compensate for the effects of glare. The approach was described as a part of a window manager system, which adjusts the brightness of the surrounding content to assure that both the HDR content as well as the normal content are visible.

In our current implementation we ignored the content of the HDR window itself, assuming, mainly for simplicity, that it was “perfect” and was to be rendered as it was. For a more general implementation, we would have to consider the elements inside that window as well. For example, if one has instrument

palettes on top of the HDR content, the brightness of these palettes should be adjusted so that they are visible and do not adversely affect the main image (i.e. not too dark and not too bright respectively).

4 References

1. Debevec, P., Malik, J., “Recovering high dynamic range radiance maps from photographs”, *Proc. of ACM SIGGRAPH '97*, pp. 369-378.
2. Durand, F., Dorsey, J., “Fast bilateral filtering for the display of high-dynamic-range images”, *ACM Trans. Graph. (special issue SIGGRAPH 2002)* 21, 3, pp. 257-266.
3. Kang, S. B., Uyttendaele, M., Winder, S., Szeliski, R., “High dynamic range video”, *ACM Trans. Graph. (special issue SIGGRAPH 2003)* 22, 3 (2003), pp. 319-325.
4. Larson, G. W., Rushmeier, H., Piatko, C., “A visibility matching tone reproduction operator for high dynamic range scenes”, *IEEE Trans. on Visualization and Computer Graphics*, 3, 4 (1997), pp. 291-306.
5. Mann, S., Picard, R., “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 (1995).
6. Mitsunaga, T., Nayar, S. K., “Radiometric self calibration”, *Proc. of IEEE CVPR (1999)*, pp. 472-479.
7. Pharr, M., Humphreys G., *Physically Based Rendering : From Theory to Implementation*, Morgan Kaufmann, 2004.
8. Robertson, M., Borman, S., Stevenson, R., “Dynamic range improvements through multiple exposures”, *Proc. of International Conference on Image Processing (ICIP)'99*, 1999, pp. 159-163.
9. H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, A. Vorozcovs, “High Dynamic Range Display Systems”, *SIGGRAPH 2004*, 23, 3, ACM TOG, pp. 760-768.
10. Schlick, C., “Quantization techniques for visualization of high dynamic range pictures”, *Proc. of Eurographics Workshop on Rendering '94*, 1994, pp. 7-20.
11. Tumblin, J., Turk, G., “LCIS: A boundary hierarchy for detail-preserving contrast reduction”, *Proc. of ACM SIGGRAPH '99*, 1999, pp. 83-90.
12. Vos, J., “Disability glare - a state of the art report”, *CIE Journal* 3, 2, 1984, pp. 39-53.