Available online at www.sciencedirect.com
ScienceDirect
J. Vis. Commun. Image R. 18 (2007) 439–451
JOURNAL OF VISUAL Communication & IMAGE Representation
www.elsevier.com/locate/jvci

Photometric image processing for high dynamic range displays

Matthew Trentacoste [a,*], Wolfgang Heidrich [a], Lorne Whitehead [a],
Helge Seetzen [a,b], Greg Ward [b]

[a] *The University of British Columbia, 2329 West Mall, Vancouver BC, V6T 1Z4 Canada*
[b] *Dolby Canada, 1310 Kootenay Street, Vancouver BC, V5K 4R1 Canada*

Received 16 November 2006; accepted 12 June 2007
Available online 17 July 2007

Abstract

Many real-world scenes contain brightness levels exceeding the capabilities of conventional display technology by several orders of magnitude. Through the combination of several existing technologies, new high dynamic range displays have been constructed recently. These displays are capable of reproducing a range of intensities much closer to that of real environments. We present several methods of reproducing photometrically accurate images on this new class of devices, and evaluate these methods in a perceptual framework.

Keywords: High dynamic range; Displays; Image processing; Photometry

---

1. Introduction

The high dynamic range (HDR) imaging pipeline has been the subject of considerable interest from the computer graphics and imaging communities in recent years. The intensities and dynamic ranges found in many scenes and applications vastly exceed those of conventional imaging techniques, and the established practices and methods of addressing those images are insufficient.

Researchers have developed additions and modifications to existing methods of acquiring, processing, and displaying images to accommodate contrasts that exceed the limitations of conventional, low dynamic range (LDR) techniques and devices. Methods exist for acquiring HDR images and video from multiple LDR images [4,13]. New cameras are capable of capturing larger dynamic ranges in a single exposure [1]. File formats have been designed to accommodate the additional data storage requirements [7,8,18]. Most relevant to this paper, high dynamic range display systems have been developed to accurately reproduce a much wider range of luminance values. The work done by Ward [17] and Seetzen et al. [14,15] has provided devices that vastly exceed the dynamic range of conventional displays. These devices are capable of higher intensity whites and lower intensity blacks, while maintaining adequately low quantization across the entire luminance range.

HDR displays are constructed by optically combining a standard LCD panel with a second, typically much lower resolution, spatial light modulator, such as an array of individually controlled LEDs [14]. The latter replaces the constant intensity backlight of normal LCD assemblies. Due to this design, pixel intensities in HDR displays cannot be controlled independently of each other. Dependencies are introduced since every LED overlaps hundreds of LCD pixels, and thus contributes to the brightness of all of them. It is therefore necessary to employ image processing algorithms to factor an HDR image into LDR pixel values for the LCD panel, as well as LDR intensities for the low resolution LED array.

In this paper, we discuss algorithms to perform this separation and to accurately reproduce photometric images. Achieving this goal entails designing efficient algorithms

---

[*] Corresponding author.
E-mail addresses: mmt@cs.ubc.ca (M. Trentacoste), heidrich@cs.ubc.ca (W. Heidrich), whitehead@physics.ubc.ca (L. Whitehead), helge.seetzen@dolby.com (H. Seetzen), gward@lmi.net (G. Ward).

1047-3203/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jvcir.2007.06.006

to produce the best images possible; characterizing the monitor; and calibrating it to reproduce the most faithful approximation of appearance, compared to the input image. We evaluate our methods by comparing the output image to the input using perceptual models of the human visual system.

The remainder of this paper is structured as follows: Section 2 covers the topics related to the work presented. Section 3 describes the task of rendering images and details the difficulties faced in doing so. Section 4 details the measurements required to correct for the actual hardware and calibrate the output, and how those measurements are incorporated into the image processing methods. Section 5 presents the results of the work, and evaluates them using a perceptually-based metric.

2. Related work

2.1. Veiling glare and local contrast perception

Any analysis of the display of images includes an inherent discussion about the viewer: the perceptual makeup of the human observer. While the human visual system is an amazing biological sensor, it does have shortcomings that can be exploited for the purpose of creating display devices. One such shortcoming is that, while humans can see a vast dynamic range across a scene, they are unable to see more than a small portion of it within a small angle subtended by the eye. This inherent limitation, called *veiling glare* can be explained by the scattering properties of the cornea, lens, and vitreous fluid, and by inter-reflection from the retina, all of which reduce the visibility of low contrast features in the neighborhood of bright light sources.

Veiling glare depends on a large number of parameters including spatial frequency, wavelength, pupil size as a function of adaptation luminance [10], and subject age. While different values are reported for the threshold past which we cannot discern high contrast boundaries, most agree that the maximum perceivable *local* contrast is in the neighborhood of 150:1. Scene contrast boundaries above this threshold appear blurry and indistinct, and the eye is unable to judge the relative magnitudes of the adjacent regions. From Moon and Spencer's original work on glare [11], 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. 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. This limitation of the human visual system is central to the operating principle of HDR display technology, as we will discuss in the following section.

2.2. HDR display technology

In a conventional LCD display, 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 transmitted at each pixel is controlled by the polarization state of the liquid crystal.

It is important to note that, even at the darkest state of a LCD pixel, some remaining light is transmitted. The dynamic range of an LCD is defined by the ratio between the light transmitted at the brightest state and the light transmitted in the darkest state. For a typical color LCD display, this ratio is usually around 300:1. Monochromatic specialty LCDs have a contrast ratio of 700:1, with numbers exceeding 2000:1 reported in some cases. The luminance level of the display can be easily adjusted by controlling the brightness of the backlight, but the contrast ratio will remain the limiting factor. In order to maintain a reasonable 'black' level of about 1 cd/m^2, the LCD is thus limited to a maximum brightness of about 300 cd/m^2. Approaches such as the dynamic contrast advertised in recent LCD televisions can overcome this problem to a degree and increase the apparent contrast across multiple frames. However, such methods can only adjust the intensity of the entire backlight for each frame displayed depending on its average luminance, and provide no benefit for static images or scenes without fast-moving action.

The fundamental principle of HDR displays is to use an LCD panel as an optical filter of programmable transparency to modulate a high intensity but low resolution image formed by a second spatial light modulator. This setup effectively multiplies the contrast of the LCD panel with that of the second light modulator such that global contrast ratios in excess of 100,000:1 can be achieved [14]. In the case of an HDR display, each element of the rear modulator is individually controllable, and together these elements represent a version of the 2D input image. Currently, this second modulator consists of an array of LEDs placed behind the LCD panel, as depicted in the upper left panel of Fig. 1. The array of LEDs is placed on a hexagonal grid for optimal packing, and the upper right panel of Fig. 1 demonstrates the LEDs of different intensities in the hexagonal arrangement that forms the backlight.

In order to ensure uniform illumination upon the LCD, the LED grid is placed behind a diffuser to blur the discrete points into a smoothly varying field. This lower-frequency illumination reduces artifacts caused by misalignment of the LCD and LED grid, and parallax from viewing the display from indirect angles, which would be very difficult to compensate for, and would perceptually be much more noticeable than low frequency errors. The width of the point spread function (PSF) is quite large compared to the spacing of the LEDs, as seen in the lower left panel of Fig. 1 which shows the point spreads of two adjacent