High Dynamic Range Display and Projection Systems

by

Helge Seetzen

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate Studies

(Interdisciplinary Studies)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

April 2009

ABSTRACT

The human visual system can perceive five orders of magnitude of simultaneous dynamic range of luminance values. Recent advances in image capture and processing have made it possible to create video content at this level of dynamic range but conventional displays are unable show this rich content. Modern display and projection systems cannot deliver more than three or four orders of dynamic range and are usually limited to lower luminance values than those found in real world environments.

This dissertation describes high dynamic range display and projection systems which resolve this bottleneck in the video pipeline and deliver luminance ranges to the limit of human perception. The design of these systems is based on the concept of dual modulation which combines several lower dynamic range image modulation components to achieve higher dynamic range. An overview of relevant perceptual mechanisms, viewer preference with respect to higher dynamic range images, and perceptual validation studies are discussed in addition to several implementation examples of dual modulation systems.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABBREVIATIONS

ASL – Average Surrounding Luminance: The average luminance level around a digital display. Usually this is the result of ambient lighting reflected of room walls.

AR – Amplitude Resolution: Number of distinctly addressable steps within the amplitude range of a display. Often provides as binary bit-depth (e.g. 8bit = 2^8 = 255 steps).

CR – Contrast Ratio: Ratio between the brightest and darkest achievable image section on a display.

CRT – Cathode Ray Tube (Display): Older display technology that uses a magnetically steered electronic ray to energize light emitting phosphors for image creation.

DLP – Digital Light Projection: Texas Instruments' brand name for a popular digital mirror device that uses an array of actuated micro-mirrors to modulate light.

HDR – High Dynamic Range: In this context HDR describes a wide range of luminance values stretching over at least 4 orders of magnitude.

HVS – Human Visual System: The entire psychophysical apparatus used for vision including the eye and core neural processes.

JND – Just Noticeable Difference: The smallest change in a continuous range that a human observer notices. For display devices this usually refers to luminance steps.

LCD – Liquid Crystal Display: Image modulator technology that uses liquid crystal material between two polarisers. The liquid crystal rotates the polarization of light according to an applied electric field and thereby modulates the transmission.

LCoS – Liquid Crystal on Silicon: Image modulator technology that uses a reflective liquid crystal arrangement directly on a drive chip.

LED – Light Emitting Diode: Solid state semi-conductor that emits light proportional to the current bias at its p-n junction.

PL – Peak Luminance: Highest achievable luminance of a display.

PSF – Point Spread Function: Spatial distribution of light from a single point such as a point light source. Also the response of a modulator to a single point of incoming light.

ACKNOWLEDGEMENTS

The history of high dynamic range (HDR) displays is long and very little of this work would have reached fruition without countless contributions by dedicated individuals and organisations. First and foremost is the original inventor of the dual modulation concept, Dr. Lorne Whitehead for his continuous support of the technology in the academic and commercial arenas. Further thanks need to go to Greg Ward for his extension of the general dual modulation concept to different resolutions and Dr. Wolfgang Heidrich for his work on core algorithms for HDR displays.

The technical development of HDR dual modulation system is also inseparably tied to the commercial entities which brought the basic concept from the University of British Columbia Structured Surface Physics Laboratory to the consumer product domain. Much of the original technical work and innovation described in this dissertation has progressed to the commercial stage through the engineering effort of Sunnybrook Technologies Inc. from February 2002 to April 2004, BrightSide Technologies Inc. from May 2004 to April 2007, and Dolby Canada Corp. and Dolby Laboratories Inc. from May 2007 onwards. Everybody who worked at or contributed to these organisations ultimately contributed to this dissertation and deserves my whole-hearted appreciation.

We have changed the world and I am proud to have worked with all of you over the years!

DEDICATION

To Dominique for her unwavering commitment and love.

To Michaël and Clara – may you find as much passion in your dreams as I have.

1 Introduction

Displays allow us to see things that we otherwise couldn't. This broad statement reflects the motivation behind the development of display or image visualization solutions over the centuries, ranging from static images such as photographs to modern display technologies. What we see, how we see, whether it is real or artificially created – those are questions that depend on the use of the display device. But regardless of the usage mode, we employ displays to show us images that we cannot easily see because they are either scenes of the past, physically inaccessible to us or simply do not exist in the real world. The only constant aspect in this broad definition is the fact that displays always send their images towards human observers.

The ultimate display system is therefore one whose only limit is the perceptual capability of the human visual system (HVS) itself. Today the performance of most displays is of course also limited by the type of content available to it, the format of the signal delivery, the type of display technology and so forth but those are all technical, not fundamental limitations. The history of improvement of display technology and related infrastructure has repetitively shown that performance specifications that are below the capability of the HVS are eventually overcome. For example, colour displays completely replaced monochrome systems, high definition is replacing lower resolution devices, display brightness is increasing continuously, and so forth. Improvements of display specifications only stop once they match or exceed what we can see.

1.1 State of the Art

Display technologies today offer a dazzling array of features but most image performance can be reduced to only a small number of fundamental specifications: size, luminance and colour range (or gamut), spatial resolution, and temporal resolution. Technical improvements in the latter two areas have already brought modern displays to, or above, the capability threshold of the HVS. Emerging 240Hz Liquid Crystal Display (LCD) panels and 360Hz digital mirror devices offer the promise to address the temporal resolution issue once and for all. Likewise, spatial resolution is today only held back by infrastructure limitations such as available broadcasting bandwidth. Higher end commercial systems have shown resolutions far beyond the ability of the eye to resolve at moderate distances[1]. The size of displays also continues to increase with up to 100" front view displays commercially available today. Projection systems remove even that size barrier and allow installation of near arbitrary image sizes.

The state of the art in terms of luminance and colour range is less impressive. The HVS has evolved in a world that features an enormous dynamic range of luminance values from that of bright sunlight to that of dark starlight. To cope with this environment the human visual perception system has evolved the ability to resolve up to five orders of magnitude of simultaneous dynamic range of luminance and, using slower adaptation processes, to span a much greater range. Modern display devices on the other hand fall short of this wide range by at least two orders of magnitude. Colour range reproduction is also limited in part because of the limited gamut available to most display technologies

[1] The IBM T221 for example delivers 9 million pixels on a 22.2" panel.

and also indirectly because luminance is a key component of the full three dimensional colour gamut.

For the first several decades of electronic display development this dynamic range limitation had no impact because the capability of most other electronic imaging devices was limited to approximately the same range. Video and still image camera sensors were limited to at most four orders of magnitude of dynamic range and thus the need for higher dynamic range displays was never pressing. As a result, most signal standards were arbitrarily limited to a similar range and in fact the entire imaging pipeline aligned to the input and output devices – a characteristic sometimes called “device referred”.

During that period only a few speciality applications required the ability to visualize higher dynamic range images. Medical imaging sensors and similar devices could capture a greater dynamic range than conventional video cameras and thus required some form of HDR visualisation on existing low dynamic range displays. Researchers also developed restoration techniques to generate HDR images from conventional captures (Mann & Picard, 1994) while others used multiple image exposures for still (Debevec & And Malik, 1997) (Robertson & Borman, 1999) (Mitsunaga & Nayar, 1999) or video (Kang, Uyttendaele, Winder, & Szeliski, 2003) HDR capture.

Displaying the resultant HDR images on conventional displays requires some form of adjustment. The class of image processing techniques for coping with the discrepancy between real world luminance 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 artefacts, 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 HVS, incorporating colour adaptation, local contrast, and dynamic range. However, even this operator exhibited some reverse-gradient effects near high contrast boundaries due to its local spatial filter, 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 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 brightness. 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.

More recent tone mapping algorithms have used perceptual models to relate the free parameters of the operators to the visual impact that these parameters have on the image in terms of visibility, contrast, brightness, or human visual response in general. While there has been some success, ultimately nobody would confuse a tone mapped image with the actual scene that it tries to portray.

Despite this limitation, tone mapping provided a pathway for higher dynamic range imaging from the specialty fields into the entertainment space. Camera manufacturers recognized the consumer desire for higher dynamic range capture and developed sensors with increased dynamic range. While today's consumer cameras cannot quite reach the dynamic range available to the HVS, they do outclass conventional display devices by at least one order of magnitude.

The increasing popularity of higher dynamic range capture also prompted the development of new compression and signal formats. Most new formats were aimed towards still images to capitalize on the availability of higher dynamic range digital cameras (Larson G. , 1998) (Fattal, Lischinski, & Werman, 2002) (Li, Sharan, & Adelson, 2005) (Ward & Simmons, August 2004)(Ward & Simmons, November 2005) while others provided solutions for as-yet to emerge HDR video cameras (Ward G. , November 2006) (Mantiuk, Efremov, Myszkowski, & Seidel, 2006). Many industries such as the film special effects world wholly converted to HDR imaging using new formats such as EXR (Kainz, Bogart, & Hess, 2003). As a result of these efforts there no longer exists any technical barrier to a full HDR video pipeline except for the display itself[2,3].

Thus, for the first time in digital imaging there is now a need for a true HDR display system to allow presentation of content that is indistinguishable from the real environment in terms of dynamic range. Such an imaging solution is often called "scene-referred" to contrast with the "device referred" current imaging pipeline. Some have argued that the accurate portrayal of a real scene is not the objective of many display applications including entertainment (e.g. cinema or TV). Rather, the objective is to entertain through visually enjoyable images. While a true statement, this misses the point. Image adjustment continues to be a core part of the creative process regardless of the capability of the display device. A HDR display simply provides more freedom for the same adjustment. Since it surpasses the visual capability of the viewer there will never again be a situation where a particular image adjustment cannot be done because the device cannot reach the desired specification. In many ways, the better term would be "human referred" or "human limited" in the sense that a true HDR display should only be bounded by the observer limits.

Outside of the field of digital displays, solutions for HDR image portrayal have already been found. For example, starting the 1950's, cinemas would advertise movies in backlit poster boxes and increase the dynamic range of the poster by printing a second,

---

[2] For more information on a complete high dynamic range pipeline see (Seetzen, Ward, & Whitehead, November 2005).
[3] A good overview of the different high dynamic range pipeline components can be found in (Reinhard, Ward, Pattanaik, & Debevec, 2005).

slightly blurred copy of the image onto the backside of the poster[4]. Similar techniques are used in a number of other parts of the motion picture industry to enhance the dynamic range of images. Matte artists will often unevenly illuminate matte painted backgrounds to increase the perceived dynamic range (Barron & Cotta Vaz, 2002). In a more related application, the principle of static dual modulation was used by Greg Ward for a stereo HDR viewer featuring two stacked transparencies (Ward G. , April 2002). While each of this solution provides an increase in dynamic range, they are all inevitably limited to static images and do not address the need for HDR displays.

The following chapters provide several avenues to achieve HDR displays for both front view and projection applications. All concepts are based on a common paradigm: Conventional display systems cannot currently achieve the required dynamic range and near term improvements are unlikely to resolve this 2 to 3 order of magnitude gap in dynamic range. Yet, certain optical and algorithmic arrangements enable the coupling of several conventional devices in such a way that the effective dynamic range of the display device becomes the product of the dynamic ranges of the individual components. Generally two systems are combined and the arrangement is termed dual-modulation. The multiplicative of dual modulation can bridge the massive dynamic range gap. For example, two devices, each with a 1,000:1 dynamic range, can be combined to yield 1,000,000:1 high dynamic range. Generally, such a device is described as a HDR display.

[4] More information at
http://www.learnaboutmovieposters.com/NewSite/INDEX/ARTICLES/doublesided.asp

1.2 Contributions

This dissertation covers a wide range of topics related to high dynamic range displays and in several cases commercial development of the concepts has occurred subsequent or in parallel to the original research presented in this document. Some of photographs of HDR devices in the document where built by commercial entities and are indentified as such. For the avoidance of doubt, the following outlines the direct contributions made by the author of this thesis to the content presented herein.

All human factor studies presented in this dissertation were designed, executed and evaluated by the author. Ms. Hiroe Li assisted with the execution (i.e. subject testing) of the luminance study described in section 4.1.1 and her corresponding co-authorship is recognised in (Seetzen, Li, Ye, Heidrich, Whitehead, & Ward, June 2006).

The list of incorporated patents and application at the end of this dissertation provides a good overview of the original inventions described in this document. Dr. Lorne Whitehead is the sole inventor of the original dual modulation concept. The author is a named inventor or co-inventor on all of the remaining 42 patents or applications relevant to the concepts disclosed in this thesis. Specifically, the author is a named co-inventor of the advanced dual modulation concept found in section 6.2, the distributed projection concept in section 7.2 and the active screen concept in section 7.3. Many of the related inventions concern implementation details for these technologies including image processing solutions, optical configurations and system architecture aspects. In most cases the concepts were co-invented by Dr. Lorne Whitehead, Mr. Greg Ward and in some cases by Dr. Wolfgang Heidrich. A variety of other co-inventors are named on a few additional patents as outlined at the end of this document.

The author personal designed and fabricated the first physical prototypes and related image processing algorithms for each of the five HDR display and projection concepts described in chapters 6 and 7. Subsequent commercial development of some of these technologies was carried out by BrightSide Technologies Inc. and later Dolby Laboratories Inc. or its subsidiary Dolby Canada Corp. In same cases the author was involved in the development of these second generation prototypes. Specifically, at the time of writing the following follow-on generation systems have been developed.

- Thomas Wan and Henry Ip of BrightSide Technologies Inc. collaborated with the author to fabricate a second generation design of the basic display concept described in section 6.1. Gerwin Damberg of Dolby Canada Corp. later fabricated a third generation design without assistance by the author (shown for illustration only at the bottom of Figure 15).

- BrightSide Technologies Inc. developed a series of second to fourth generation prototypes of the advanced display concept described in section 6.2. Further generations were developed by Dolby Laboratories Inc. as well as several major consumer electronics companies such as Samsung, Sharp, Lg. Philips and others.

- Gerwin Damberg and Michael Kang of Dolby Canada Corp. fabricated a second generation prototype of the basic projection concept described in section 7.1 without assistance by the author.

- No second generation prototypes were developed commercially for the distributed or active screen concepts described in sections 7.2 and 7.3.

As Chief Technology Officer of BrightSide and Director of HDR Technology at Dolby Laboratories, the author was involved in several of these development projects but none of them were related or are described in this dissertation beyond specifically indicated illustrative photographs. The algorithms, system designs and optics described in this dissertation are limited to the first generation designs developed exclusively by the author and often do not correspond in detail to later generation commercial systems.

Sections of this document contain relevant portions of publications by the author as indicated in the bibliography. The majority of these publications have co-authors including Dr. Lorne Whitehead and Dr. Wolfgang Heidrich. In all cases the included work is primarily the work product of the author of this dissertation and permission was granted by co-authors to use relevant sections.

2 Basic Imaging Concepts

The display industry uses a range of terms to describe the performance of imaging devices. Often, these terms are used without adequate qualification to achieve a better marketing position for a particular device. This has led to an erosion of the quality of terminology and rendered many metrics nearly meaningless. For example, the term "contrast" carries an intuitive meaning but is used in many different ways in the industry to the point that the same display can have a "contrast" of 200:1, 2,000:1 or 2,000,000:1. Many of these terms are necessary to describe the performance of imaging devices. The next sections provide working definitions of key terms in the context of this dissertation. Details on measurement techniques for these metrics can be found in ASTM (American Society for Testing and Materials (ASTM), 1987) or the VESA Standard (Video Electronics Standards Association (VESA), 2001).

2.1 Dynamic Range

Dynamic range is a measure of the range (usually the ratio) between the highest and lowest luminance achievable under any device setting. Usually this means the luminance range between a full white screen and a completely black screen. For displays that cannot show a full brightness white screen, another suitable pattern is used (e.g. Plasma displays cannot show a full white screen at peak luminance due to power limitations and consequently use a mostly black pattern with a small full white area that is used to measure peak luminance).

Note that dynamic range is often incorrectly presented as “contrast” as it is in general a larger number. For example, most displays have a brightness adjustment slider that can be used to reduce or increase screen luminance. The darkest level of this display would then be a completely black screen shown at the lowest slider setting. For the brightest level the adjustment slider would be moved to peak so that a full white screen becomes as bright as possible. If this adjustment were manual then nobody would be bold enough to argue that the contrast of the system is influenced by the setting of the slider. After all, nobody watches television with a finger constantly shifting the brightness slider corresponding to dark and bright scenes. This distinction becomes less obvious if the slider adjustment happens automatically based on the image histogram. Dynamic range is therefore the widest luminance range of the device including adjustments over time while contrast describes the simultaneous luminance ratio at a fixed parameter setting (i.e. at one particular point in time).

2.2 Contrast

Contrast measures the ratio between a simultaneously shown bright and dark section of an image. A common way to measure contrast is to use a pre-defined black and white pattern such as an ANSI checkerboard[5].

As outlined in the last section, dynamic range is always at least equal to contrast and generally higher due to internal light leakage from the bright sections of contrast patterns. The fact that technologies differ significantly in the severity of light leakage has led to confusion in the marketplace. Manufacturers of devices with a high dynamic range but very low contrast, such as "dynamic iris" projection systems, tend to mislabel dynamic range as "on/off" or "sequential" contrast. The term contrast is therefore used in at least two different variants in the consumer industry: "Sequential Contrast" (also referred to as "on/off Contrast" or "Inter-frame Contrast", usually comparable to dynamic range unless a display system has significant manual luminance control) and "Simultaneous Contrast" (also referred to as "Intra-frame Contrast" and usually specified as ANSI Contrast).

Ultimately, neither metric can stand alone. The best viewing experience comes from a display that has a high contrast within the range that generally matters. Few conventional TV sequences have average luminance levels below 10% or above 50% of the highest data value. There are exceptions to both but rarely for longer time periods on

[5] For details see
ANSI/NAPM IT7.228 (IEC Draft 61947-1) Section 4.3 Contrast ratio
ANSI/NAPM IT7.227 (IEC 61947-2) Section 5.3 Contrast ratio
VESA FPDM Ver 2.0 (2001) Sections 302-3 Darkroom Contrast Ratio of Full Screen, and
304-9 Checkerboard Luminance and Contrast (n×m)

the screen. A good display maintains a high contrast inside that range. In other words, a good display can show sparkling bright highlights, a few bright areas and still maintain a very dark black.

The black level of a display is also critically influenced by ambient illumination. Some displays are so reflective that even modest ambient illumination can completely mask most of the dark range. Most metrics for contrast assume a completely black viewing environment which is rarely the case in normal viewing but unfortunately the only reasonable reference assumption[6]. The same assumption is used in this dissertation for all measurements, but ambient illumination can have a more major impact on HDR displays than on conventional displays due to the presence of very low luminance values in dark regions.

[6] More information on the impact of ambient illumination and incorporation of ambient factors into contrast measurements can be found in (Becker, 2006), (Kelly, 2006). Measurement standards are contained in Commission Internationale de L'Eclairage (CIE) references (1977), (1979) and (1987). Specific information for different display technologies can be found in (Kelley & Jones, 1997) and (Krantz, Silverstein, & Yeh, 1992).

2.3 Colour Gamut

The term "colour gamut" refers to the complete range of colours that a display device can present to the viewer. This straightforward definition is complicated by the fact that the definition of "colour" itself is far from simple. Colour is an attribute of visual perception and thus influenced by measurable physical parameters as well as subjective psychophysical factors. This complex interaction is described in more detail in section 3.4.

To avoid the complexity of subjective colour perception, the display industry generally measures "colour gamut" as a coverage factor of a particular chromaticity space. Usually these spaces are two dimensional chromaticity diagrams that represent the full three dimensional gamut scaled by some perceptual factors and mapped into two dimensions. A common chromaticity diagram is based on the CIE 1931 Standard Colorimetric Observer (2 Degree Observer). The colour gamut of conventional three colour displays (i.e. red, green and blue sub-pixel) is defined by the location of the three primaries in the chromaticity diagram. These three points form a triangle on the diagram and the area bounded by the triangle can then be compared to the total area of the space or, more commonly, a reference chromaticity triangle (i.e. ITU-R Recommendation BT.709 for high definitional television[7]). The size of the colour gamut of a display device can then be provided as the ratio of these two areas.

---

[7] More information can be found on http://en.wikipedia.org/wiki/Rec._709

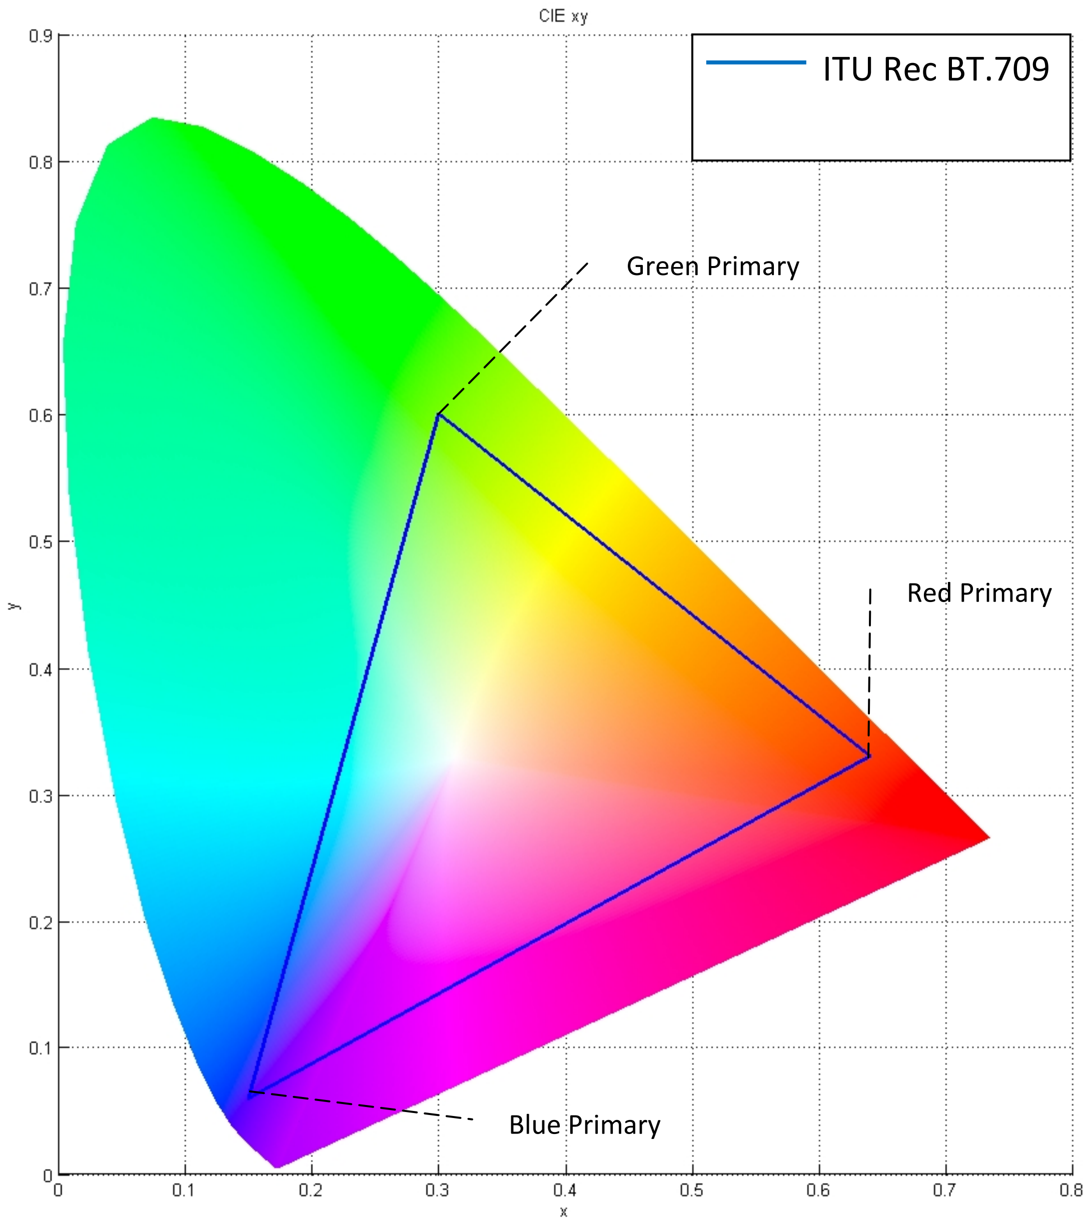

Figure 1: ITU-R Recommendation BT.709 reference primaries shown in the CIE 1931 xy colour space. The red, green and blue primaries form a triangle and the area of this reference triangle is commonly used as the denominator in expressing the colour gamut coverage percentage for display devices.

It is important to understand the severe limitation of this method of specifying colour gamut. Aside from overlooking the perceptual aspects outlined in section 3.4, this approach also inappropriately compares colour gamuts. Displays specifications often advertise a colour gamut of “120% Rec.709” or “80% NTSC” which are the chromaticity area ratios described above. Comparisons of this kind ignore the fact that the CIE 1931 colour space, in which both Rec.709 and the older NTSC standard are defined, is not perceptually uniform. A chromaticity shift near the central area of the space produces a much larger perceived change of colour than an equal sized shift near the edge of the gamut. This non-linearity varies in all areas of the space. It is therefore possible to achieve 120% coverage of Rec.709 in many different ways including many that would be perceived has having smaller colour gamut than a 100% coverage device.

Figure 2 shows a fictitious example of a display that would achieve a much higher than 100% of Rec.709 rating based on this metric but would clearly not be an appropriate display due to the lack of red. In this example all three primaries are shifted towards green in such a way that the lack of relevant coverage in the red region is “offset” in area by extra coverage in the less relevant green region.

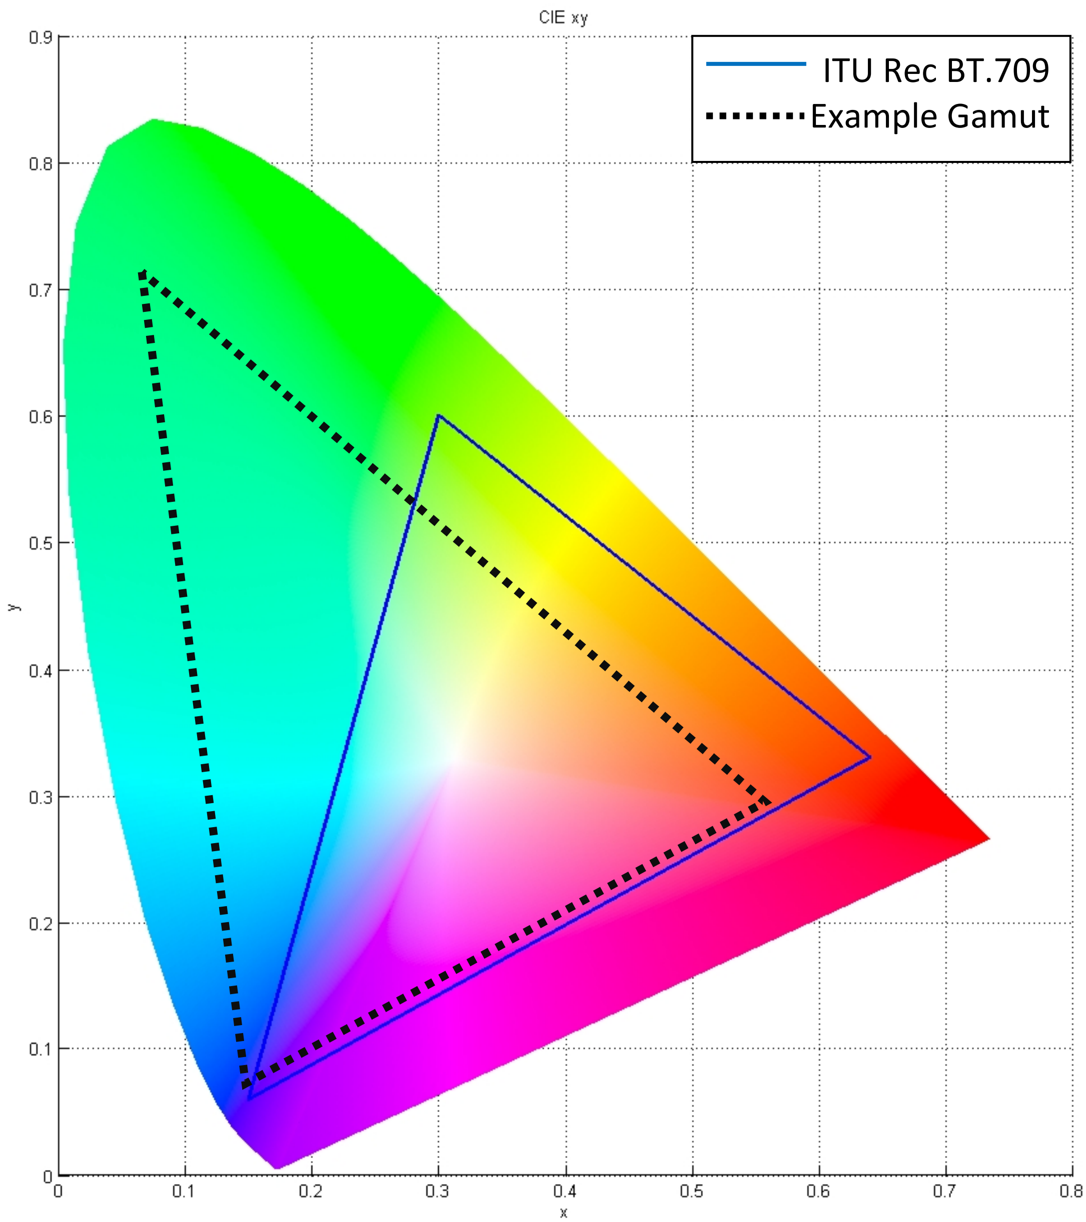

Figure 2: Illustration of limitations of area comparison for colour gamuts. The example gamut has a larger area than the Rec.709 reference gamut and would thus achieve a rating of over 100% Rec.709. But would clearly not be an acceptable display device due to the lack of red. Cases like this illustrate the limitations of the simplistic gamut area specification as a percentage of Rec.709 or other reference gamuts.

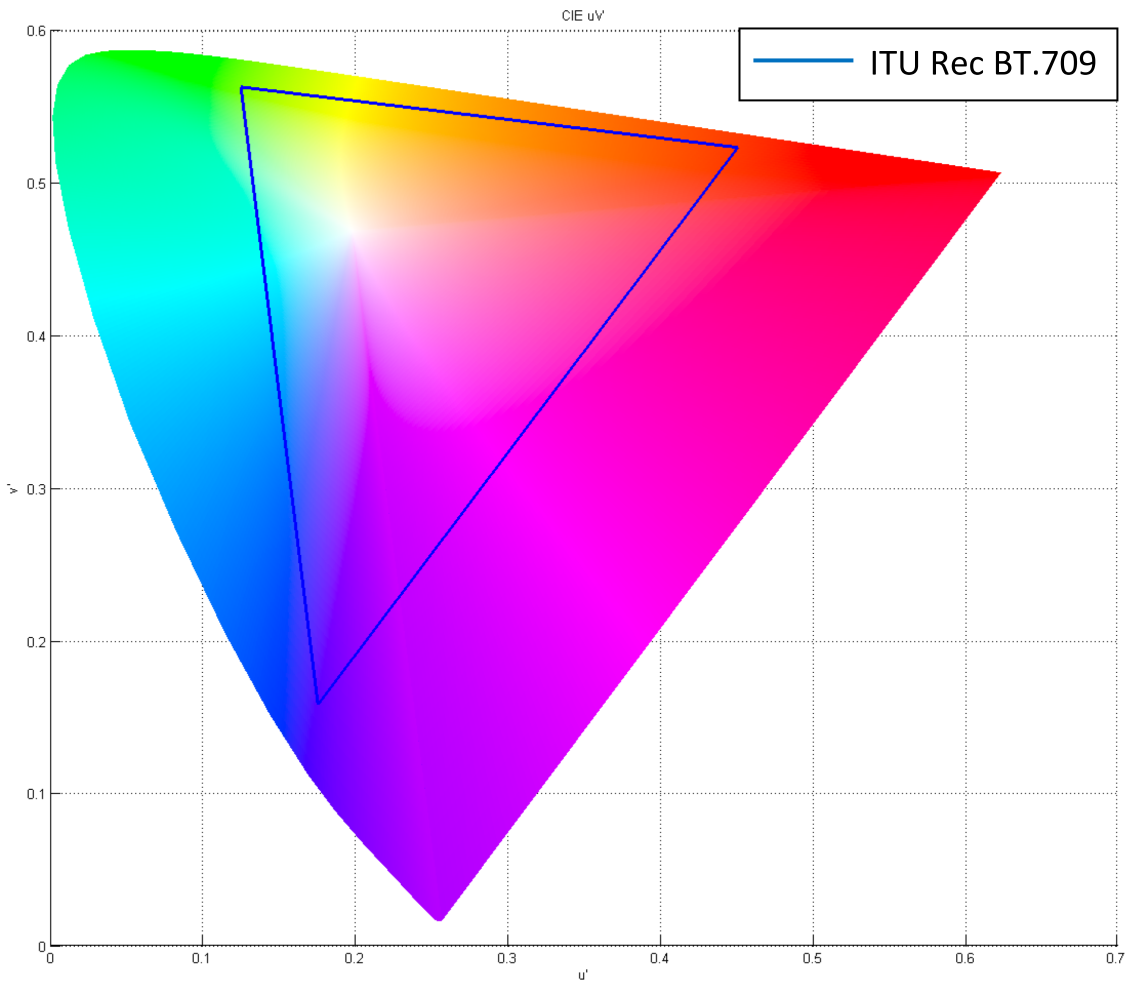

Figure 3: ITU-R Recommendation BT.709 reference primaries and white point shown in the perceptually more uniform u'v' colour space.

This problem can be reduced by using a more perceptually uniform chromaticity diagram such as CIE u'v'. This space has been shown to yield gamut comparisons that more closely match actual visual impressions (Hunt, 2005). The basic problem of lack of positional information for the device gamut remains though. Fundamentally, the area coverage metric for display colour gamut only has limited meaning if the chromaticity diagram and reference triangle are well document. Without this documentation the metric is entirely meaningless.

2.4 Amplitude Resolution

Amplitude resolution or amplitude quantization defines the number of levels between the maximum and minimum value of a range. In the context of display devices this term is usually used to describe the number of control values that produce a distinct luminance value at a pixel. Since most display control is digital, the amplitude resolution is often provided as a bit depth specification (binary digits). A conventional display system generally has a bit depth of 8 or 10 bit per colour sub-pixel which provides for 256 or 1024 controllable intensity levels. For RGB colour systems the effective amplitude resolution is often quoted as 24 bit (8 bit red + 8 bit green + 8 bit blue = 24 bit colour) or sometimes erroneously as 32 bit (24 bit colour + 8 bit grey = "32 bit"). Such nomenclature is confusing and designed mostly for marketing specifications where high numbers are viewed favourably. In the following, bit depth or amplitude resolution is used as a description of addressable steps per individual channel. Clearly, such channels can be combined to make different combinations.

It is also important to note that amplitude resolution has no direct connection to display contrast or dynamic range even though it is often mistakenly used to imply higher contrast. A low dynamic range of 1 to 2cd/m^2 can be sub-divided into a million steps while a very wide dynamic range could have only a single step (e.g. a light bulb turning on or off). For actual display devices the response curve of the modulator is also usually non-linear to mimic the perceptual response to luminance. A common response curve is a power function with an exponent of 2.2 (often called "gamma" γ) where L is the ultimate luminance level of the display and V is the incoming video signal.

$$L = V^{\gamma} \quad \textbf{(1)}$$

This choice of non-linear response is primarily driven by the much lower luminance range of the output device compared to real world scenes. The preferred gamma tends to approach 1 for higher luminance displays (linear response). In addition, spatial filtering such as dithering or anti-aliasing is often used to increase the effective amplitude resolution of display devices. Overall, a good conventional display should have an effective bit depth of 8 bit or more while much higher amplitude resolution is required for proper HDR imagery.

3 The Human Visual System

The human visual system (HVS) has tremendous capabilities. We can see during star-lit nights and in bright sunlight, discern a wide range of colours, and resolve the finest details. Moreover, we clearly enjoy exercising our visual apparatus to near its limits. We like colourful compositions; design expensive lighting environments, and in general pattern our world to include all the detail that our eyes can see. Much of this is the result of our evolutionary history. We have evolved in a wide range of environmental conditions and excellent visual abilities were a cornerstone of our evolutionary success.

Like all system with such a massive and complex capability range, the HVS also has some shortcomings. Visual artefacts, neurological masking and a range of other limitations are a normal part of our everyday visual experience. This chapter outlines the capabilities and limitations of the HVS as they related to HDR imaging.

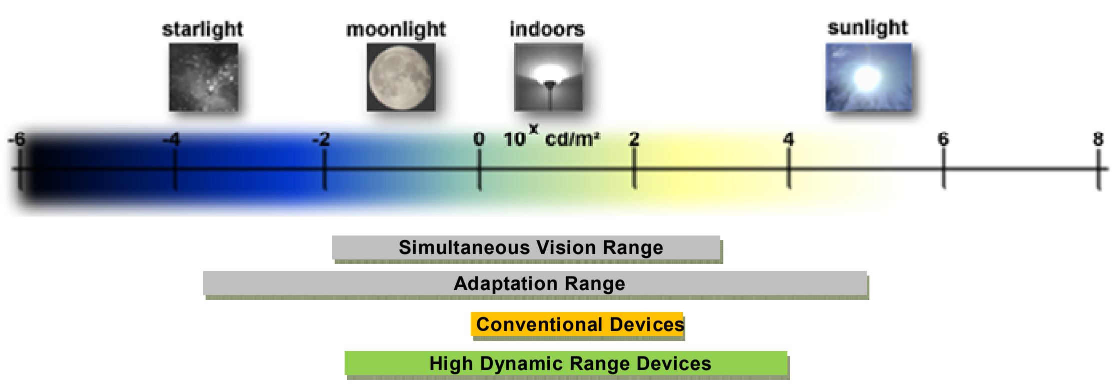

Figure 4: Capability of the HVS compared to conventional and HDR devices.

3.1 Dynamic Range

The world around us features a tremendous dynamic range from a star-lit night to the bright reflections of the sun. Our retina addresses the demands of a wide environment range with two types of receptors: rods and cones. Rods are highly sensitive and thus used primarily in dark environments between $10^{-1}cd/m^2$ to $10^{-6}cd/m^2$ (*scotopic range*). Though capable of resolving very small luminance differences in this range, rods do not provide colour vision and offer very limited acuity. This limitation is countered by cones capable of colour vision in the *photopic range* of $10^{1}cd/m^2$ to $10^{8}cd/m^2$. Both receptors are stimulated in the *mesopic range* of $10^{-1}cd/m^2$ to $10^{1}cd/m^2$.

Maintaining sensitivity across this wide range of luminance requires a complex adaptation system. In general, our retinal sensitivity isn't constant and decreases rapidly with increasing light intensity. This effect, known as *response compression*, creates a smooth ramp-off prior to saturation of the receptor and limits the instantaneous receptor range to approximately 3 to 4 orders of magnitude. Adaptation mechanisms, such as pigment depletion and pupil changes, ensure that the receptor input is appropriately scaled for the ambient environment. The timescale of these mechanisms can vary from minutes to near instantaneous. The faster adaptation mechanisms allow the eye to operate near-simultaneously over a dynamic range of 5 to 6 orders of magnitude. No conventional display technology can cover a dynamic range anywhere close to this capability and the bar for HDR devices is therefore set quite high.[8]

[8] A good overview of visual adaptation can be found in (Hood & Finkelstein, 1986) with more display device oriented examples give in (Boff & Lincoln, 1988).

3.2 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 visual angles). Different researchers report different values for the threshold past which we cannot resolve 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.

This inherent limitation can be explained locally by the optical imperfections of the eye including scattering in the cornea, lens and retina, and diffraction in the coherent cell structures on the outer radial areas of the lens. These effects are responsible for the "bloom" and "flare lines" seen around bright objects. The diffraction effect also causes a lenticular halo.

Veiling luminance (often termed "disability glare" or "bloom") is the result of light scattering in the ocular media with roughly equally impact of cornea, crystalline lens and retina scattering. Figure 5 illustrates this process. Light L_a from source A scatters inside the eye onto the same receptors as light L_b coming from source B. This adds luminance L_s to the receptor region that would normally correspond only to source B. The effective contrast ratio of the boundary between A and B is therefore $L_a/(L_b+L_s)$ and not L_a/L_b. In this way veiling luminance places a limit on the perceivable contrast in a small region. The magnitude of L_s depends on the angle of separation α and the luminance and solid angle of the source.

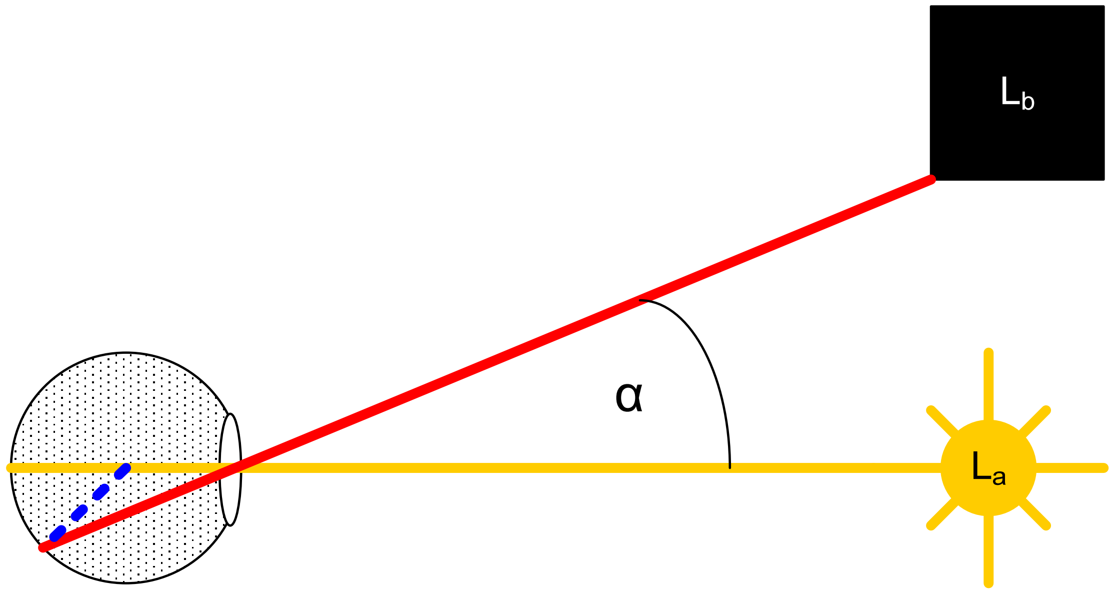

Figure 5: Veiling luminance effect.

Empirical psychophysics research (Barten, 1992) led to a Point Spread Function (PSF) P(α) for the bloom effect given by

$$P(\alpha) = \eta\delta(\alpha) + \frac{c}{f(\alpha)} \tag{2}$$

The constant η represents the fraction of the light that is not scattered and c is an empirically determined calibration constant. The function f(α) give the actually scattering light component as a function of the angle α between the light source and the dark region of interest. It has been successfully modeled to very high precision with a first order term of $f(\alpha)= \alpha^2$ for a wide range of observers[9]

Moon & Spencer's original work on glare (Moon & Spencer, 1945) confirms that any high contrast boundary scatters at least 4% of its energy on the retina to the darker

[9] The full equation can be found in (DICOM, 2003)

side of the boundary, obscuring the visibility of the edge and details within a few degrees of it. If the contrast of an edge is 25:1, then details on the darker side are 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, high contrast content clearly has some effect. An observer will notice when one region is much brighter than another, both by the challenge it creates in viewing the boundary, and by accommodation 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 – a familiar experience for any driver during night-time. 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 HDR displays can uniquely reproduce.

3.3 Just Noticeable Difference Steps

Given the complex adaptation mechanism of the HVS, display designers often ask how many distinct input/output levels are necessary to cover the desired range without banding or similar quantization artefacts. 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 display response 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 illumination. Psychophysical research in the area of human threshold versus intensity (t.v.i.) can address this issue in the context of HDR displays. The t.v.i. curve has been measured by vision researchers at different adaptation levels (Lubin & Pica, 1991) (Ferwerda, Pattanaik, Shirley, & Greenberg, 1996).

The combined t.v.i. curve for rods and cones yields values for Just Noticeable Difference (JND) steps as shown in Figure 6. A JND defines the smallest luminance step at a particular luminance value that yields a visible change. 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.

Based on Barten's original work, an analytical formula for JNDs was derived for the DICOM standard (Barten, 1992)(Barten, 1993) (DICOM, 2003). For conventional displays with a limited luminance range of typically 500cd/m^2 to just under 1cd/m^2 it is often sufficient to use a non-linear 8-bit or 10-bit modulator as the required number of

JNDs for this luminance range is only approximately 500 (Muka & Reiker, 2002). However, as the range of displayable luminances increases, so does the number of JNDs required to cover this range. Any high luminance HDR display should therefore have the goal to reproduce at least as many steps as predicted by this model for the luminance range of the device.

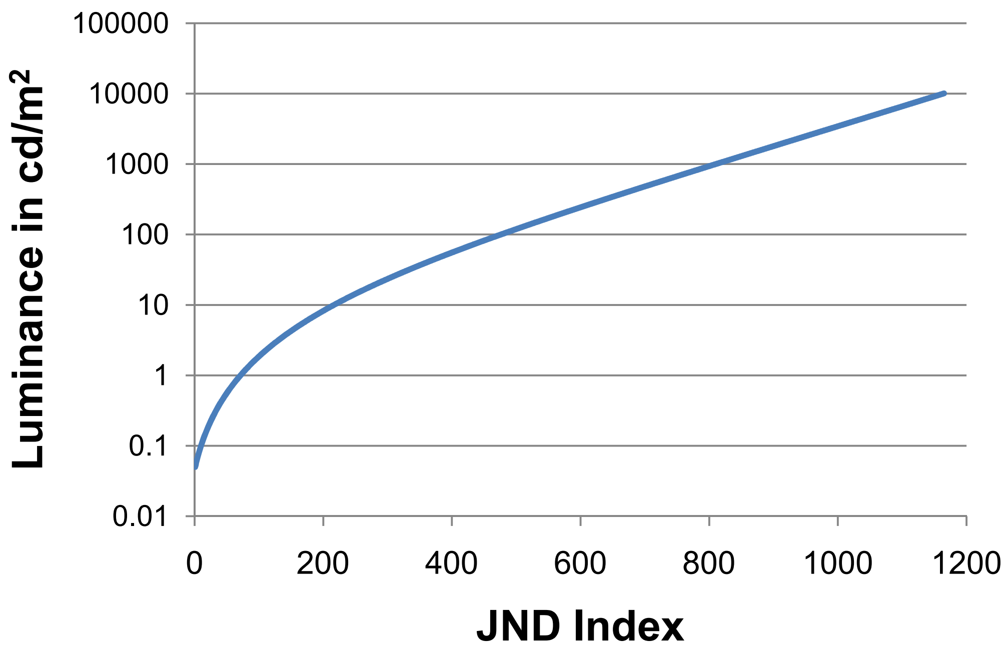

Figure 6: Number of JNDs for different maximum intensities according to the Barten model used in (DICOM, 2003).

3.4 Colour Perception

Colour perception is both one of the most basic sensory experiences in our daily lives as well as one of the most complex mechanisms of the HVS. It represents the interaction of physical stimuli with measurable attributes, and psychophysical elements that are only known to us through inference. The formal definition of colour states that "perceived colour depends on the spectral distribution of the colour stimulus, on the size, shape, structure, and surround of the stimulus area, on the state of adaptation of the observer's visual system, and on the observer's experience of the prevailing and similar situations of observation (Fairchild, 2005)". Colour perceptions are defined by their various attributes such as lightness, brightness, colourfulness, chroma, saturation, and hue. A full description of these attributes and their interaction can be found in the colour science literature and the following focuses on key aspects pertaining to display systems.

As discussed previously in section 2.3, the display industry generally reduced colour gamut specifications to the position of the corners of the device primaries in the chromaticity space. Often, the specification is even further reduced to the ratio of the area of the triangle created by these corners versus a reference triangle such as ITU-B Recommendation BT.709. The corners of the triangle represent the interaction of the emission spectra of each primary with the cone photoreceptors of the HVS. In principle, the light emitted by each primary sub-pixel has a consistent spectrum defined by the physical arrangement of the display (e.g. the spectrum of the light source and, if applicable, the spectral modulation of the display colour filters, etc). This spectrum causes excitement of one or more of the three retinal cones according to their response curve. Finally, in very simplistic terms, the neural signal from one or more cones causes the appropriate sensation of colour. In this fashion the chromaticities of display primaries correspond roughly to hue and saturation at least in the sense that increasing saturation causes an increase in the chromaticity triangle size.

However, the basic chromaticity triangle does not capture the full complexity of colour perception. First, the chromaticity triangle does not indicate the luminance of the display. Relative and absolute luminance correlates with perceived brightness and lightness and has a significant impact on colour perception throughout the visible range. The impact is particularly strong in the mesopic or scotopic range where colour perception is very limited. It is therefore important to further specify the luminance of the display device as full white and full black (or at least the relative contrast of the display). Peak luminance is generally measured at the so-called “white point” of the device where all three primaries are driven at equal RGB signal levels. The relationship between the relative luminance levels of the three primaries at this point provides the white point position in the chromaticity diagram and allows for interpolation of the device response inside the triangle.

The relative colour gamut can therefore be defined by the primary chromaticities, white point chromaticity, black point chromaticity and the relative luminance of the two points. The best way to describe all these factors in a single representation is the use of a colour space. Colour spaces generally feature a third dimension related to luminance in addition to the two dimensional structure of a chromaticity diagram. The CIELAB space for example is defined in terms of lightness (correlated to luminance), chroma and hue (both correlated to chromaticity)[10]. This three dimensional volume provides comprehensive information about a display device and allows for more meaningful comparison of the full colour gamut of two devices.

A colour space like CIELAB captures the essential elements of a display device as they related to the perception of colour in the absence of ambient lighting (i.e. a completely dark room). Environmental factors such as ambient illumination and direct ambient lighting (e.g. flares on the display surface) can have a substantial impact on the perception of colourfulness as well as general perceived image quality. These factors need to be considered for a full understanding of colour perception but are usually outside of the control of the display device developer.

[10] More information about CIELAB can be found at http://en.wikipedia.org/wiki/Lab_color_space

4 Viewer Preference

Having defined the physiological capabilities of the HVS in the previous chapter, the question remains whether viewers actually prefer to see such a wide dynamic range. For scientific imagery it would be possible to design HDR displays that accurately simulate the entire range of luminance outlined previously. In most cases we care more about the perceptual impact of images rather than their accuracy though. Good looking images trump accurate images any day. Or do they?

To explore the preferences of viewers, one must first identify key specifications of displays relevant to visual preference. Mechanical specifications such as display size and spatial resolution can be ignored since they scale with viewing distance. Any model of viewer preference developed for a particular viewing configuration should be transferable to other conditions if the angular ratios between viewer and display are maintained.

The most common non-mechanical specifications of a display are peak luminance, contrast ratio, amplitude resolution, temporal resolution, and colour. For newer display technologies, colour is arguably the closest to having reached the requirements of our visual system. Emerging wide colour gamut displays have broadened the range of presentable colors considerably. Moreover a comparison with real world chromaticity values obtained from multi-channel photography shows that even the more limited sRGB gamut (equivalent to ITU-R Recommendation BT.709) comes close to portraying most of the environment around us (see Figure 7). There is little doubt that the chromaticity capabilities of displays will soon encompass our perceptual requirements.

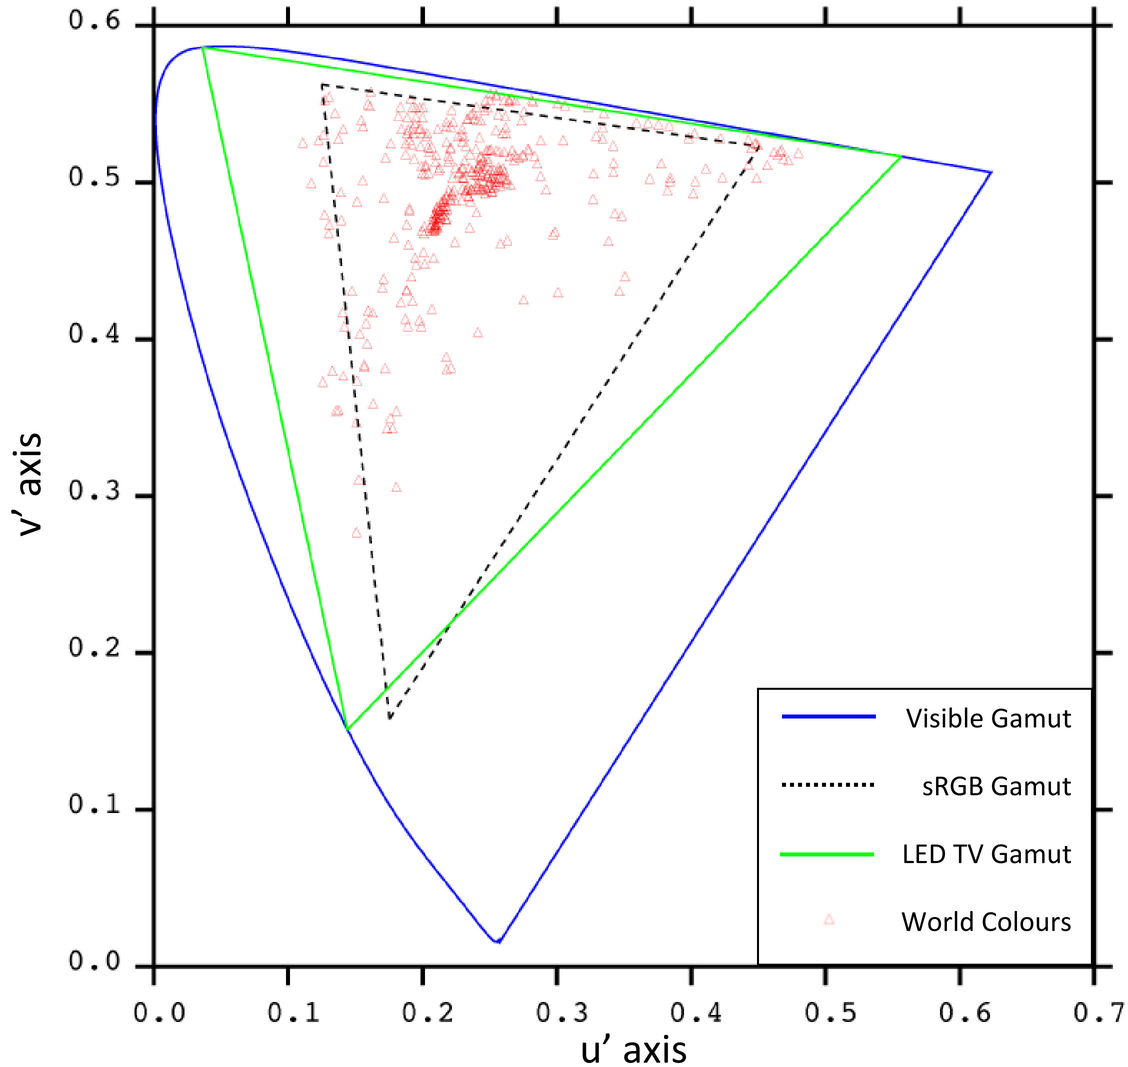

Figure 7: Real-world chromaticity values plotted in the perceptually uniform u' v' space. (Data taken from (Glassner, 1995).)

A similar argument can be made for temporal resolution. Liquid Crystal Displays have long left the 60Hz/16ms response time barrier behind, Digital Light Projection (DLP) chips are have recently reached refresh rates of 360Hz and Cathode Ray Tubes were fast enough all along.[11] While there remain issues such as motion blur for LCD and colour break-up for DLP, we have seen progress on all these fronts. More importantly, the gaps to be closed are relatively narrow and certainly not fundamental obstacles for any of the display technologies involved. One might debate whether today's display technologies offer adequate temporal resolution, but there is little doubt that tomorrow's will.

The story is quite different for the remaining three characteristics. The peak luminance, contrast ratio and amplitude resolution of most displays are significantly lower than the capabilities of the viewer (see chapter 3). Peak luminance is the easiest characteristic to measure and describes the highest luminance value attainable by the display. Contrast ratio, by the very definition of the word, measures the ratio between a bright and dark section of an image. Amplitude resolution describes the number of distinct steps of luminance that can be portrayed by a display. For digital devices amplitude resolution is usually provided in terms of bit depth. Virtually all displays use a gamma response curve (often a gamma with a shallow start in the black levels) to distribute those available steps in an optimized fashion for our visual system.

Three studies were conducted to investigate the relationship between peak luminance (hereinafter often abbreviated PL), contrast ratio (hereinafter often abbreviated CR) and amplitude resolution (hereinafter often abbreviated AR), and their impact on viewer preference.

[11] 120Hz LCD solutions are commonly on the market today while 240Hz LCD panels have been shown in commercial prototypes. DLP systems with 360Hz refresh were also commercially introduced recently.

4.1 Experimental Design

Each study addresses one of the display characteristics of interest and in each case the results of the previous study are taken into account to provide secondary confirmation. The participant pool varied during the study. The first study used 38 participants, the second 40 participants and the third only 12 participants. All participants were between 18 to 35 years old. For all studies approximately 1/3 of the participants were female and all had normal or corrected to normal vision including colour. All three studies were conducted on 18” dual modulation displays. A detailed description of these devices can be found in section 6.2 but for now it is sufficient to know that these displays use a conventional LCD in front of a dynamically adjustable matrix of light emitting diodes (LED) which allows for a much higher dynamic range of luminance than the LCD alone. Participants sat approximately 1m away from the screen. All participants were given 10 minutes to adjust to the ambient environment prior to the study and had a chance to see an introductory series of random images spanning the dynamic range of the study. This gave them an overview of the study and helped to normalize the semantic scales used in the first study.

The influence of ambient illumination on display luminance perception is significant. CIE recommendations for illuminance levels in living room environments are 50 to 120 Lux depending on the task. In this study an ambient illuminance of 100 Lux and a modestly reflective environment were used. The test room had diffuse medium grey walls and no specular surfaces. The average luminance on the wall behind the display was $20cd/m^2$ +/- $5cd/m^2$.

4.1.1 Luminance & Contrast Preference

The first study aims to establish a general overview of the impact of peak luminance and contrast ratio on viewer preference. For this purpose four basic representative scenes were selected. Sixteen test images were created by permuting four variations of peak luminance levels (1,600cd/m^2, 1,200cd/m^2, 800cd/m^2 and 400cd/m^2), with four variations of contrast level. Contrast adjustments were made around the center point of the encoding range. Since most images have an average luminance below the center point of their encoding range, a lower CR usually increases average image luminance even if PL stays constant. As a result the actual PL levels for different scenes very slightly. This was considered during data evaluation.

Each participant was exposed to two identical-looking 18” dual modulation displays as described above. A randomly selected image from the set above was shown on both displays. On one of the displays the image was rendered normally with the appropriate variations in the LED backlight matrix. On the reference displays the LCD image was the same but the backlight was uniformly lit for a peak luminance of 400cd/m^2. The rendering algorithm for the dual modulation display (see section 6.2.1) was further constrained to make the above PL and CR adjustments entirely on the LED matrix so that the reference image remained unchanged for all 16 combinations per scene. This setup ensured that colour, spatial information, screen reflectance and so forth were identical between the two displays. When the image appeared on each screen the participant was asked to rank the varying image in comparison to the reference image on a semantic scale employing four bi-polar adjectives: bright – dim, deep – flat, pleasant – unpleasant, realistic – unrealistic. A central mark on the scale indicates no perceived difference between the two displays. In this fashion each participant went through the images in all combinations. One of the four images was repeated with each combination to estimate learning and other such long term effects during the study. No statistically significant learning effect was observed.

A static display with a higher luminance was used for a high PL study since the PL of the dual modulation display used in the main study was limited to 1,600cd/m^2. Both LCDs were replaced by a calibrated stack of transparencies displaying the same series of LCD images. By replacing the LCD with transparencies the transmission of the system improved dramatically and PL levels of over 8,000cd/m^2 were achieved. All other aspects of the high PL study were the same as in the main study. The results of the two studies align well in the area of luminance intersection (top PL of the first study and bottom PL of the second are both 1,600cd/m^2).

Figure 8: Images used in the viewer preference study.

4.1.2 Contrast Preference

The first study provided a general overview of the PL & CR preference space. The rough division of PL and CR into only four choices each was necessary to maintain a manageable number of test images per participant. The second study focused on the relationship between PL and CR.

In a similar setup to the first study, each participant was exposed to a random image selected from a larger sample of 20 representative images. Each image was displayed at a randomly selected PL level and CR. Using the UP and DOWN keys of the keyboard the participant could adjust the CR of the image until it was most pleasing. The CR adjustment was designed to maintain the same average image luminance at all CR levels by adjusting contrast around the average point of the image data. Once the preferred CR setting was reached, the participant confirmed the selection and a new image was displayed. Each participant was shown all 20 images under 4 different random PL levels.

4.1.3 Amplitude Resolution Preference

The previous two studies provide a framework for PL and CR preference but assume that amplitude resolution of the displays is high enough at all PL and CR levels. This assumption is warranted since the test dual modulation display offers a full 16-bit depth with an effective spatial dither due to the analog nature of the LED in the backlight. Conventional displays do not have such a high amplitude resolution so it was desirable to investigate the minimum AR threshold per PL level. The goal of this third study was to investigate how many distinct luminance levels are necessary to present a visually smooth image for a fixed PL, and how many linear bits are needed to present these levels. These JND levels have been studied extensively by the psychophysics community with the two main models coming from Barten (Barten, 1992) and Ferwerda (Ferwerda, Pattanaik, Shirley, & Greenberg, 1996). The Barten model is used mostly in image critical applications such as medical imaging and predicts approximately 1,000 JND over a luminance range of 0.05cd/m^2 to 4,000cd/m^2. Ferwerda's study suggests a much smaller number of JND (approximately 250 in the same range). There is general agreement in the psychophysical community that for ordinary display applications, Ferwerda's estimate is on the low side, probably because his studies used a pulsing target and such a transient stimulus leads to higher perception thresholds. Regardless of the discrepancy between the two studies, both apply to abstract test environments rather than TV screens in a typical living room environment.

Figure 9: Example target with ambient ring image (target/background contrast is greatly exaggerated).

Using a single 18” dual modulation display each participant was shown a 3 degree target (1m viewing distance) within a uniform background luminance. Adjusting the UP and DOWN key allowed the participant to adjust the luminance of the target until it was barely distinguishable from the background. Once the target was visible the participant pressed ENTER and the background luminance was set to the current target luminance. In this fashion the participant traversed the entire luminance range of the dual modulation display in JND steps.

An initial pilot study with 6 subjects indicated that the luminance steps at the extreme low end of the range of the dual modulation display are larger than a single JND. This portion of the range (approximately below 1cd/m^2) has consequently been ignored in subsequent tests. The pilot also replaced the more conventional square transparent grating target with other geometric shapes to counter false positives resulting from a repetitive pattern. This was a concern because the study took 30 to 45 minutes per participant. The results of the pilot approximately matched the predictions of the Barten model to within 15% of predicted JND number and with a similar distribution for all participants. The target shape change and our experimental protocol are therefore appropriate for such a consumer experience oriented study.

In order to represent common display viewing conditions the main study added a ring of low spatial frequency image content at a distance of more than 6 degrees from the target. By adjusting the average luminance level of the outer ring the study simulated the impact of surrounding image content on the area of interest. The average surrounding luminance (ASL) level remained constant for 4 participants and then changed so that for the total 12 subjects 3 ring images were used. The ring images were taken from representative images adjusted for an ASL of 1,200cd/m^2, 800cd/m^2 and 400cd/m^2. The images were blurred strongly to avoid distracting spatial frequency content. Figure 9 shows an example of such a setup with an outer image ring, a constant background area and finally a geometric target.

4.2 Experimental Results

The first study covers a large multi-variable space. For data analysis all results are first corrected for individual participant variation using the duplicate images inserted in each series. Next, the results on the semantic scales are linearized using the Bright-Dim scale as a guideline for the other three scales. A pilot study showed that the Bright-Dim scale accurately matches the base three logarithm of the average luminance of each image. This comes as no surprise since this is a fair approximation of brightness perception in the given range of luminance range. The adjustment for non-linearity in the other scales is made by assuming consistency in the non-linearity of all scales. Finally, the three remaining scales are averaged into a single Perceived Image Quality (PIQ) scale. The results of all three scales are very similar and combining them greatly simplifies data presentation. The second and third studies are single variable designs and can therefore be used directly.

4.2.1 Luminance & Contrast Preference

Figure 10 shows the results of the first study including the high luminance data from the static display test (PL values above 1600cd/m^2). With the exception of two participants the results of the study were fairly consistent with at most 14% variation between the rankings of individual participants. The two outlying participants had generally far higher scores throughout and reached the top of the semantic scales prematurely at low CR and PL values. Since this saturated their results to the peak PIQ value their results have been discarded.

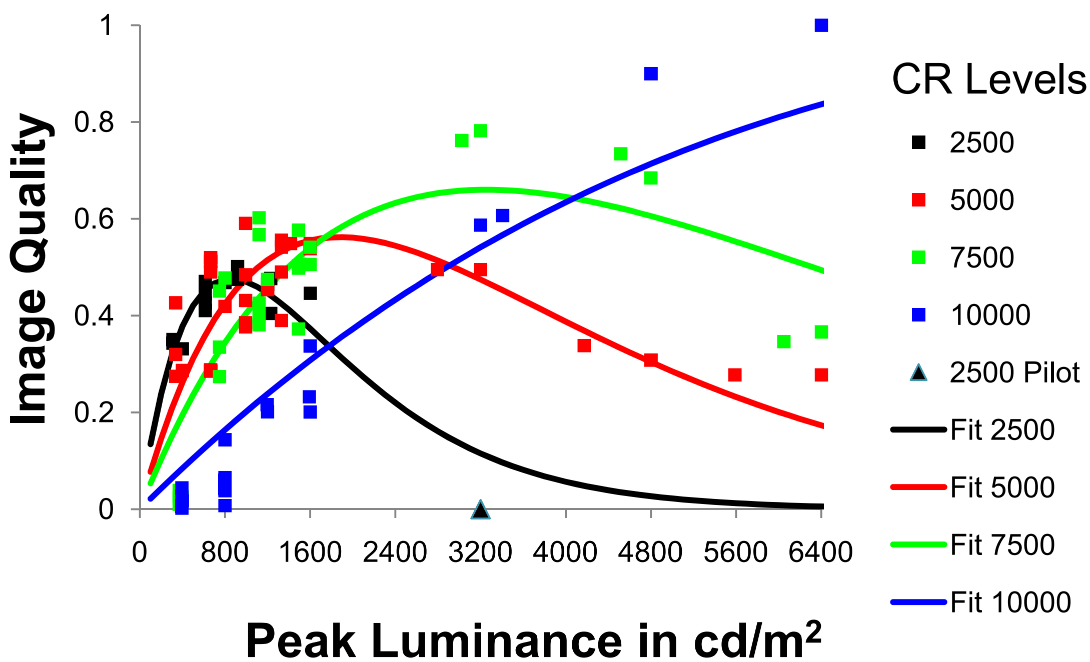

Figure 10: Viewer preference as a function of peak luminance at different contrast ratio levels (note that the PL 3,200cd/m^2 data point in the CR 2,500 series came from a pilot study).

In considering Figure 10, it appears that for each CR level PIQ increases with PL value up to a maximum after which it decreases. In other words for each value of PL, there is an optimum value of CR. At low PL the best result is obtained with the lowest CR used in the study (2,500). As PL increases this CR is insufficient to maintain the highest PIQ and higher CR provide better results. This relationship holds for rising PL and implies that high CR are not optimal for lower PL settings. The effect can be explained by considering the image creation process in section 4.1.1. The test images were scaled into the dynamic range between PL and PL/CR (the black level for this image) using a gamma of 2.5. At high CR and low PL this means that the bulk of the image content is shifted towards the dark region of the image – often to the point that shadow information is completely lost to the viewer. It might therefore be possible to gain some benefit from a higher CR at lower PL by adjusting the grey level distribution. A simple model can be used to describe PIQ as a function of PL for different CR levels according to

$$PIQ(PL) = a \cdot PL \cdot e^{-b \cdot PL} \quad \textbf{(3)}$$

A least-squares fit, adjusting the free parameters a and b, yields a reasonable fit with the data. For this fit, the optimal value of PL (PL_{opt}) and the corresponding maximal value of PIQ (PIQ_{max}) are given by

$$PL_{opt} = \frac{1}{b} \tag{4}$$

and

$$PIQ_{max} = e^{-1\frac{a}{b}} \tag{5}$$

These values are shown for the four different values of CR in Table 1.

CR Level	2500	5000	7500	10000
PLopt	856	1881	3256	12419
PIQmax	0.475	0.562	0.660	1

Table 1: Fit parameters for Figure 10.

4.2.2 Contrast Preference

With the general relationship between PL and CR indicated above, the results of the second study establish the optimal CR range for each PL level. The results have been combined into bins of 100 CR levels each and averaged over each bin, as shown in Figure 11. The standard deviation is modest for most bins though there are a small number of outliers for most bins as one would expect for perception data. The optimal contrast ratio can be related to peak luminance based on a least-squares fit as

$$OCR(PL) \cong 2862ln(PL) - 16283 \qquad \textbf{(6)}$$

The function also agrees with the two high PL results from static display data in the first study (The two highest points in Figure 10 correspond to PIQ_{max} for the 7,500 and 10,000 CR lines).

In view of the supporting data, it would be inappropriate to suggest that this value for optimal contrast function represents a sharply defined optimum; rather, it provides only a general guideline for the contrast value that achieves the optimal Perceived Image Quality for a given peak luminance. Figure 12 shows the results of the first study filtered by the optimal contrast function such that only results within 10% or 30% of the optimal contrast for that PL are considered. The resulting relationship between PL and PIQ with appropriate CR is logarithmic and well defined up to approximately 1,600cd/m^2. A pilot study with a much smaller number of images for PL values up to 12,000cd/m^2 provided higher luminance results. It appeared that subjective ratings of image quality declined for PL values above 7,000cd/m^2, regardless of contrast, suggesting that yet another effect is at work in this range. Likely, the problem in this range is simply discomfort glare given the modest ambient luminance level.

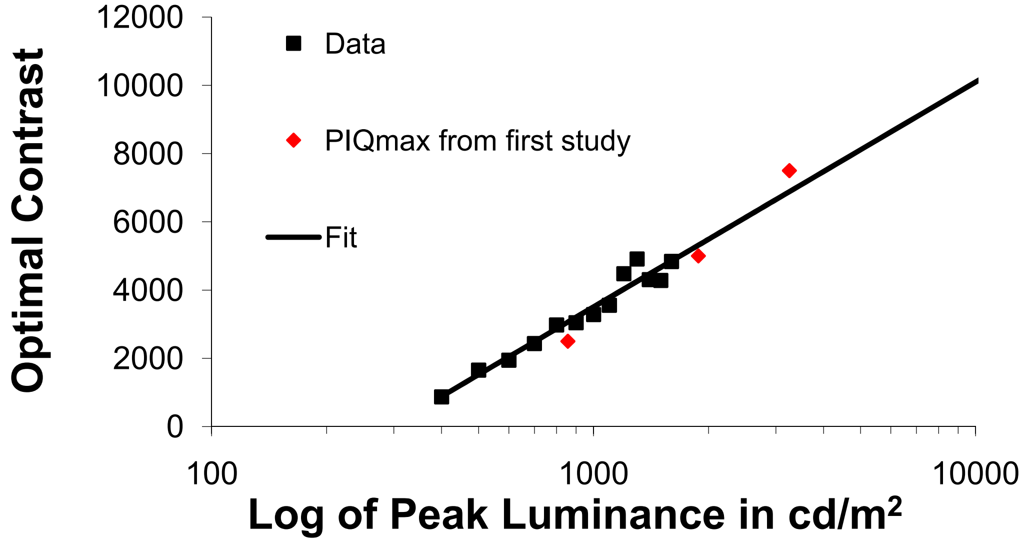

Figure 11: Log relationship between optimal contrast and peak luminance.

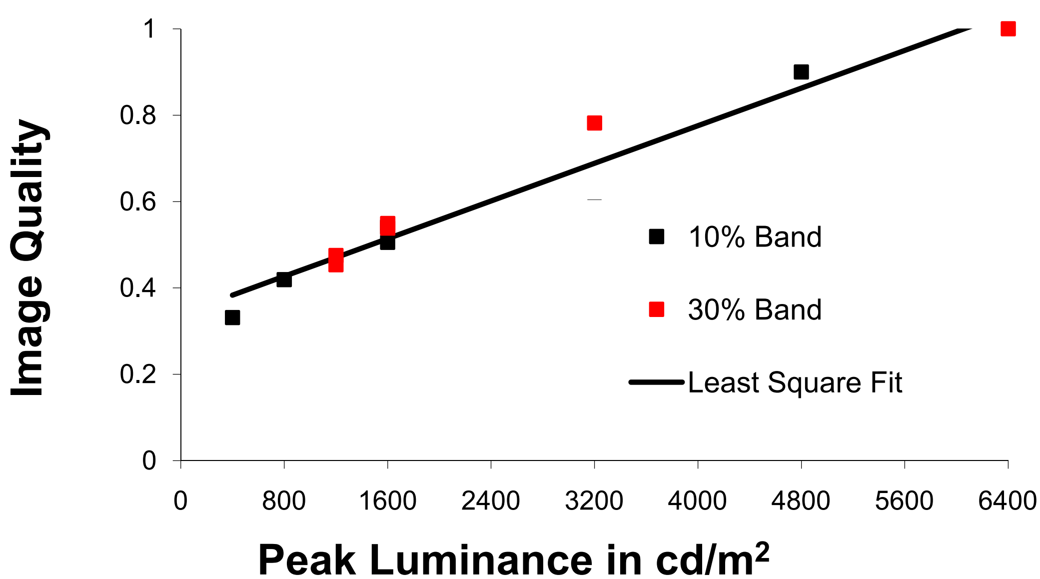

Figure 12: Viewer preference data filtered for optimal contrast.

These two models describe the relationship between PL, CR and PIQ under the given ambient conditions. As outlined in section 4.1.3, AR and grey level distribution play a critical role in this relationship. The AR of the dual modulation display is fortunately far beyond the requirements of the Barten model above $1cd/m^2$. It is therefore fair to assume that the results of both studies are unaffected by AR limitations. Yet this is not true for conventional displays which are usually limited to 8-bit or 10-bit depths. Figure 13 shows the results of the third study for each of the three ASL levels. The fourth (black) line is the Barten model which is the equivalent of an ASL of $0c/dm^2$. Table 2 summarizes JND counts found within the range of the experiment.

The ASL clearly has a significant impact in reducing the number of distinguishable steps. At the same time the perception model remains the same and the perception threshold stays approximately constant over the entire luminance range of the study. With the approximate ASL of specific application these values can then be used to estimate the required AR to remain below threshold. For television images for example ASL is usually between 20 to 30% of peak luminance.

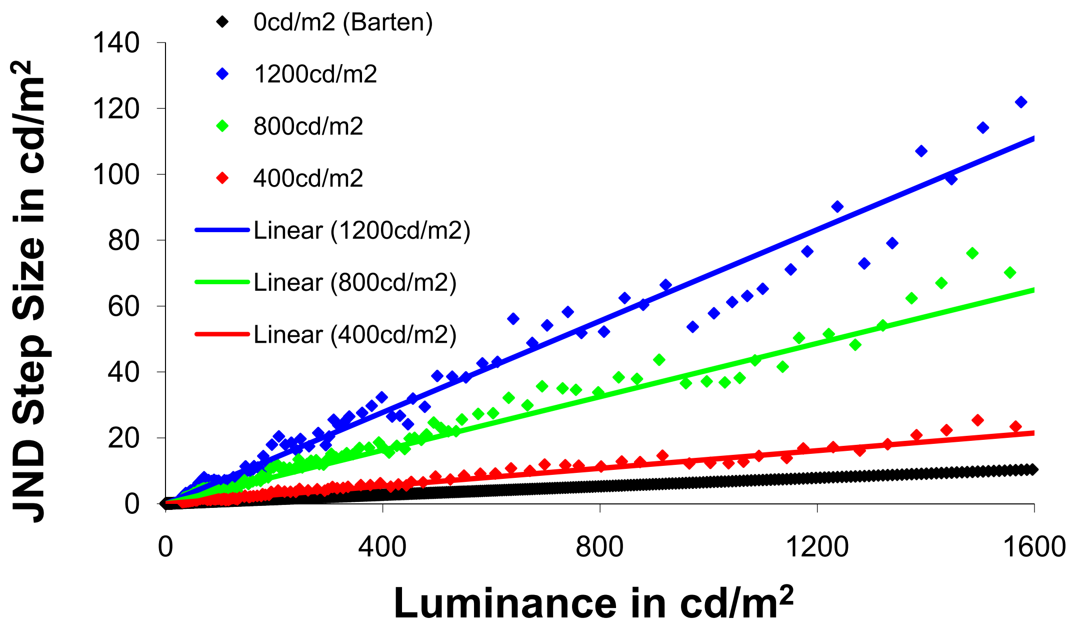

Figure 13: JND step size per luminance level under different average ring luminance conditions

Ring Luminance	0	400	800	1,200
JND (limited range)	845	579	296	114
Threshold	0.007	0.013	0.041	0.069

Table 2: JND steps and thresholds per ring image luminance.

5 New Goals

The display industry has undergone a major transition in the last decade. The declining era of Cathode Ray Tube (CRT) displays rarely saw luminance levels above 200cd/m^2 and early Plasma displays had a similar limit. Both technologies could in principle achieve high contrast but internal light leakage often prevents high contrast in all but full black screen test images. Neither technology has the ability to significantly increase luminance levels or create a true HDR experience.

The emergence of flat panel LCD displays has removed this low luminance barrier. Liquid Crystal Displays are passive modulators so their luminance level can be arbitrarily scaled by increasing the light output of their backlight. As a result the luminance levels of modern LCDs are higher with 500cd/m^2 becoming the de facto standard for televisions. Higher luminance versions in the range of 750cd/m^2 to 1,500cd/m^2 are available for outdoor viewing. These luminance levels are impressively high compared to older CRT technology but still short of the goals for HDR imaging. Improvements in contrast have also been made over the years. Liquid Crystal Displays now commonly feature ANSI contrast levels of 1,500:1. This allows higher end LCD to reach a 500cd/m^2 peak luminance while maintaining a reasonably black level below 0.5cd/m^2. Thus, LCD is the solution of the middle way. They offer higher luminance than CRT and Plasma but are still an order of magnitude lower than required for HDR viewing. The black level of LCDs is reasonably dark but still an order of magnitude brighter than ultimately needed.

In more general terms, display systems are based on one of two image creation principles: Light emission or light modulation. Emissive displays create light at each pixel and project it directly to the viewer. Examples of emissive displays include Organic Light Emitting Diode (OLED) displays, Plasma displays and the older CRT displays. Modulation displays are passive and require a separate light source to illuminate them. Once illuminated, the modulator can reduce the amount of light reflected or transmitted by each pixel and thereby adjust the intensity of each pixel. Digital Mirror Devices, LCD and their cousins Liquid Crystal On Silicon (LCoS) displays all fall into this category. Achieving the dynamic range and colour gamut targets set forth in the previous chapter is impossible with commercially available display systems or image modulator of either type.

Emissive displays such as OLED or Plasma displays can in theory achieve a very high contrast ratio but currently cannot reach the luminance level required for HDR imagery. While it is commonly acknowledged that Plasma and Cathode Ray Tube (CRT) displays will never reach high luminance, OLED is at least in principle still a contender for high luminance output. The materials for OLED devices have high theoretical quantum efficiencies but suffer from significant degradation over time. This limitation has so far prevented OLED displays from reaching the luminance level of today's displays, much less the required brightness for HDR imaging.

Modulation displays such as LCD or DLP have in principle arbitrarily scalable light sources so that the luminance targets are theoretically achievable. Unfortunately, their limited contrast prevents them from achieving the required dynamic range for true HDR performance. Neither LCD nor DLP can achieve complete extinction of light. As a result any increase in light source intensity also raises the black level of the display. This increase is linear and effectively eliminates any hope of achieving HDR performance with single modulation devices.

As seen previously, viewer's preference for luminance can be very high if the contrast of the display system is equally high and the image is appropriately adjusted. Reduced to display terms this means that small highlights in the image are extremely bright, black is black and the rest of the image is adjusted such that the overall luminance level is not glaring (i.e. skin tones, normal clothing and other diffuse surfaces don't glow). Unsurprisingly, this is exactly the type of viewing environment that we encounter every day in the real world. Light sources and specular reflections thereof are very bright. Yet, most of our ambient environment is about 1 to 2 orders of magnitude dimmer than those bright lights and dark regions are often very dark. These ratios occur naturally outside. Indoors we ask interior architects to maintain the same ratios.

Even in indoor environments this means that a display needs to deliver thousands of cd/m^2 at peak levels to re-create the impression of small light sources or reflections. The range of 3,000cd/m^2 to 6,000cd/m^2 is supported by the results for viewer preference shown in chapter 4. Further indirect support comes for our normal non-display related preferences. For decades interior spaces have been lit by a mixture of diffuse area illuminators and small spot lights to create interesting lighting environments. From metalwork and mirrors, over stainless steel appliances to glass chandeliers, we try to incorporate sparkle into our environment. The luminance of these spot lights and reflections are by design in the same luminance range of 3,000cd/m^2 to 6,000cd/m^2.

Simultaneously the display needs to anchor its black level at the point where human perception cannot resolve differences anymore. In the worst case of a dimmed viewing environment this level is below 0.1cd/m^2 and preferably closer to 0.01cd/m^2. Moreover, good amplitude resolution needs to be available at these low levels to show realistic shadow gradients, dark textures and other effects that we routinely encounter in the real environment. The required dynamic range for a true HDR experience is therefore of the order of 100,000:1. Precise limits for the upper and lower end are challenging to define due to the variations in ambient viewing environments and individual preferences.

Within this wide range of luminance the display should be capable of showing a smooth range of grey levels. Exact amplitude resolution requirements depend on the ambient environment and display response curve. As a point of reference, the Barten model cited in chapter 4 illustrates a need of about a thousand JNDs for the range of 0.05cd/m^2 to 4,000cd/m^2 with steps distributed along a very specific curve. To reach these JNDs with a normal display response curve requires a bit depth of at least 14-bit.

Finally, a full HDR experience would ideally also mimic our real world in terms of size. Of course all display solutions strive to be as large as possible to make the viewing experience as impressive and engaging as possible. In the case of high dynamic range there is an additional motivation. As noted above, many of the requirements for high dynamic range viewing are influenced by the ambient environment. High dynamic range viewing tries to make maximal use of our visual capability and the same capability adjusts with input through adaptation. The ultimate HDR experience is therefore always subject to the adaptation level of the viewer. There are two solutions to this problem. One would be to carefully manage the image content according the ambient environment surrounding the display. That information could be used to predict the adaptation state of the viewer using psychophysical models and adjust the image content to maintain the desired HDR experience. This requires the use of sophisticated sensor technology and colour appearance models[12].

The second option is to scale the HDR image to the point that it dominates the viewer's field of view. The display would in this case effectively control the adaptation level of the view since adaptation is based on the light input to the retina. This simplifies the image adjustment process considerably and in fact probably eliminates it completely. For the ultimate HDR viewing experience the display would show image luminances comparable to our desired real world experience (scaled for an indoor environment if needed). Those luminance levels would set the viewers adaptation level to the level encountered in the real world for similar scenes and thus the image content would look as desired – the eye, not the display does the work.

In summary, a true HDR experience requires high luminance (ideally 3,000 – 6,000cd/m^2), high contrast (100,000:1 or more), high amplitude resolution (14-bit or more) and ideally a large screen that substantially fills the field of view. No conventional display solution available today comes close to these requirements.

---

[12] A good overview of colour management systems for displays can be found in (International Color Consortium (ICC), 2004).

6 Dual Modulation Display Technology

A unique solution is needed if neither modulation nor emission displays are likely to reach the goals of the previous chapter. A single modulator might have insufficient contrast for HDR applications but two modulators in series can achieve the target if their contrast ranges are optically multiplied. Dual-modulation HDR displays combine two modulators in series to achieve high dynamic range. The choice of modulator can be any combination of passive or emissive modulators as long as appropriate optical coupling is possible. For example, two LCD panels can be stacked to effectively multiply the contrast of the two panels. The result is effectively a very high contrast modulation system with arbitrarily scalable luminance and thus HDR capability.

The following sections describe several implementation variants of this dual modulation concept.

6.1 Basic Dual Modulation Displays

The basic dual modulation display provides the easiest path to HDR imaging using only conventional component. A conventional LCD uses two polarizing layers and a liquid crystal layer to modulate light coming from a uniform backlight. The backlight is typically a fluorescent tube assembly[13]. The light is polarized by the first polarizer and transmitted through the liquid crystal layer where the polarization is rotated in accordance with the control voltages applied to each pixel. The light then exits the LCD panel through the second polarizer which absorbs a portion of light depending on the alignment of polarization. The luminance level of the light emitted at each pixel is controlled by the difference between the rotated polarization due to the liquid crystal layer with respect to the fixed polarizing layers. The LCD process is not perfect and the stack of polarisers cannot completely prevent light transmission. Even at the darkest state of a pixel, light is emitted. 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 1,500:1, with monochromatic specialty LCDs going up to 4,000:1 (e.g. medical imaging displays). The luminance range of the display can easily be adjusted by controlling the brightness of the backlight, but the dynamic range ratio remains the limiting factor. In order to maintain a reasonable 'black' level of below 0.5cd/m^2, the LCD is thus limited to a maximum brightness of about 500cd/m^2.

[13] A good overview of polarization systems for displays can be found in (Robinson, Chen, & Sharp, 2005).

The basic modification introduced by the HDR technology involves inserting a second light modulator and increasing the brightness of the backlight. The 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 most basic dual modulation display design, the backlight and the first modulator are combined into a single DLP projector using a Digital Mirror Device. The three central components of the basic dual modulation display are then the projector, the LCD and the optics that couple the two. Using these components, each image on the display is the result of modulated light coming from the projector which is directed onto the rear of the transmissive LCD by the optical system, modulated a second time by the LCD, and properly diffused for viewing.

A simple prototype of this design can be build by placing a commercial DLP projector behind a LCD panel. Other projector solutions such as LCD or LCoS can of course also be used. To reduce unnecessary light loss the colour wheel (if present) can be removed from the projector, resulting in a monochrome display system with a threefold increase in brightness due to the absence of the colour filter. New control electronics has to be integrated into the commercially available projector to re-synchronize it in absence of this colour wheel.

To allow direct transmission of the light from the projector, the LCD panel needs to be separated from the conventional backlight and all of the optical layers behind the display have to be removed to create a transmissive image modulator.

The optical system for such a basic configuration can be simple. The conventional projection lens can be used if the throw length is sufficiently short to be able to focus on the back of the LCD panel. A Fresnel lens should be placed directly behind the LCD panel to collimate the projected light into a narrow viewing angle for maximum brightness of the dual modulation display and to avoid colour distortion due to diverging light passing through the colour filter of the LCD. Finally, a standard LCD diffuser should be used to redistribute the collimated light into a reasonable viewing angle. The combination of Fresnel lens and diffuser also eliminates hot spot and parallax artefacts.

All three components should be 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 (digital or optical shift mechanisms). 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 artefacts associated with even a minor misalignment, the projector image can be deliberately blurred. Appropriate algorithms can compensate for that blur in the LCD image as described in the following section.

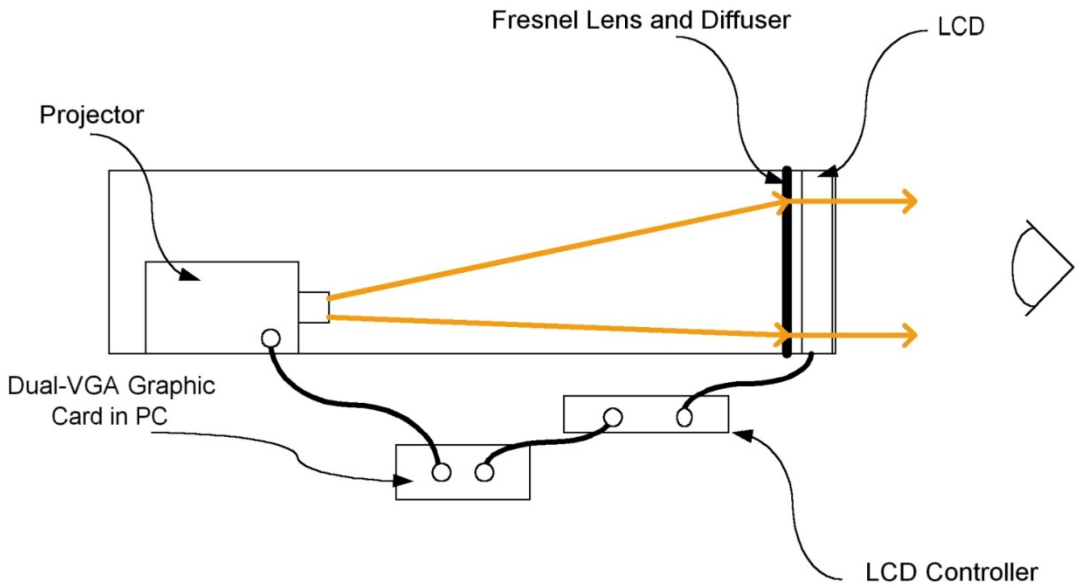

Figure 14: Schematic diagram of a basic dual modulation system.

Using this configuration, the light output of each pixel of the dual modulation 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). For example, if the DLP has a dynamic range of 800:1 and the LCD a dynamic range of 1,000:1, the theoretical dynamic range of the dual modulation display is 800,000:1. Imperfections in the optical path introduce noise that reduces the dynamic range. In actual prototypes of this design, dynamic ranges of over 100,000:1 have been measured. The luminance values matching these boundaries are a result of the brightness of the projector and the transmission of the LCD.

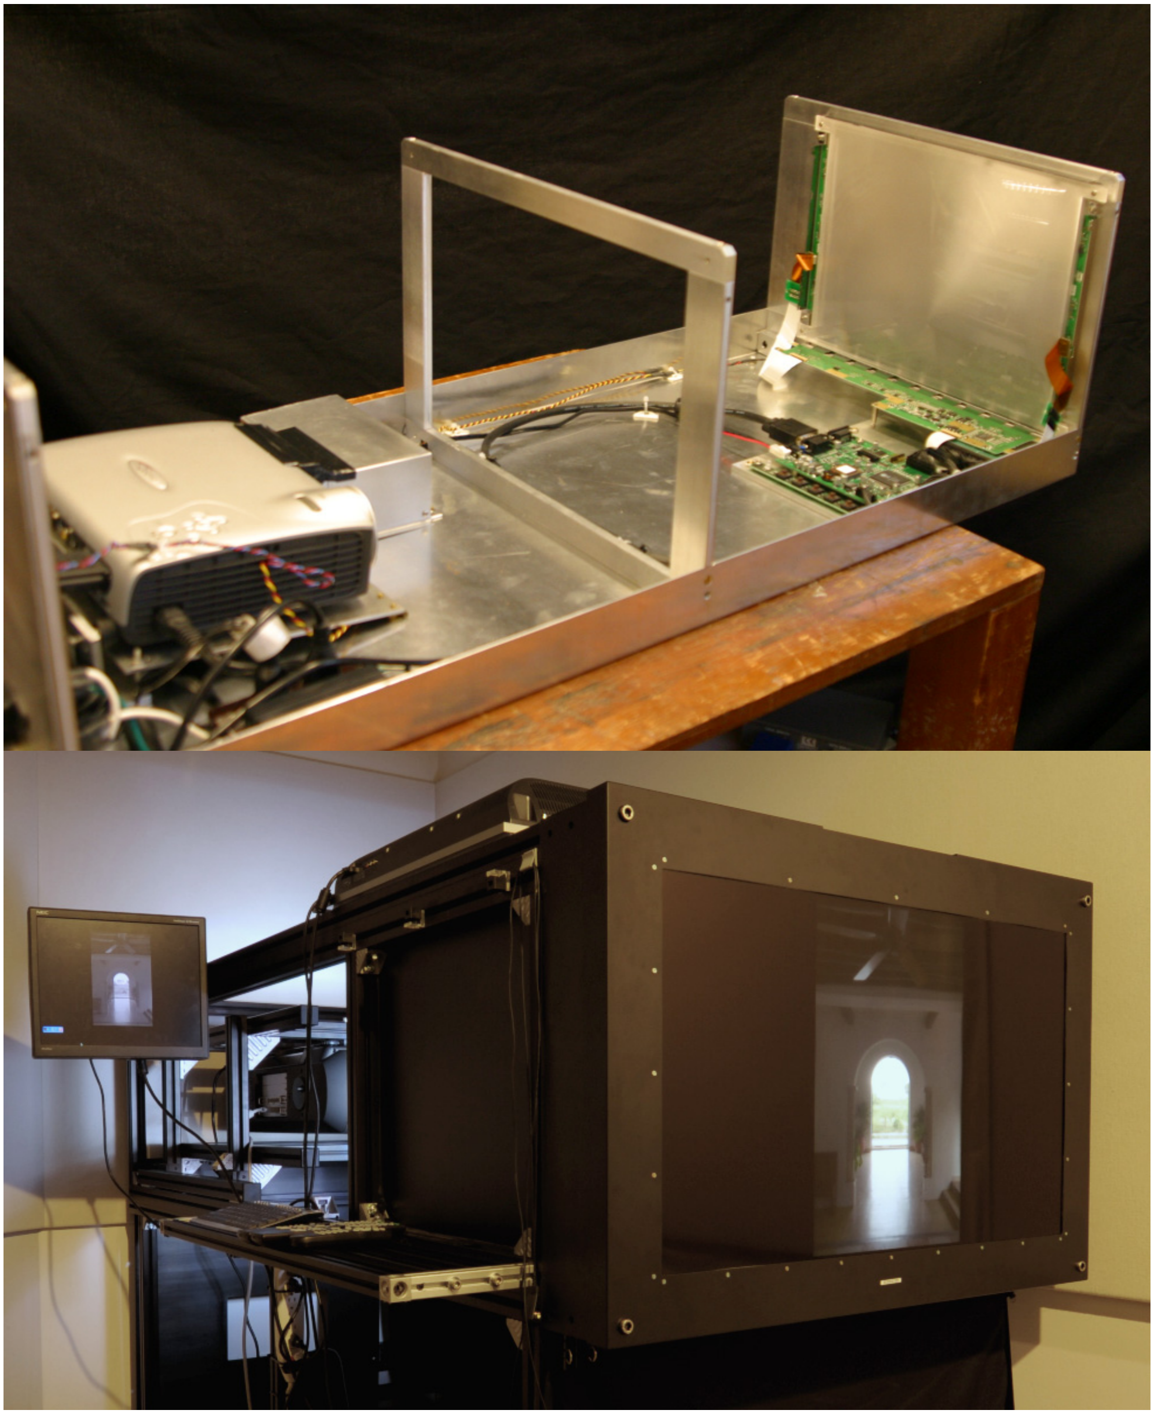

Figure 15: Examples of basic dual modulation designs. Top: Smaller second generation 15.4” prototype developed by the University of British Columbia. Bottom: Large third generation 40” prototype with similar configuration developed by Dolby Canada Corp.

Figure 15 shows photographs of different prototypes built on the basis of this design. The smallest version uses an Optoma EzPro737 DLP projector rated at 1,200 Lumens, or approximately 2,400 Lumens once the RGB colour filters are removed (Each filter for red, green and blue eliminates approximately 2/3 of the incoming light but DLP projectors often use a fourth "white" filter to boost light output)[14]. The Sharp 15.4" LCD panel used in this prototype 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 colour LCD without any losses is only 16% due to the light reduction of 50% at the polarizer and another 66% due to the RGB colour filter. Assuming that the light emitted by the dual modulation display is diffused across a solid angle ω, the maximum luminance is then given by

$$L_{\max} = \frac{\Phi_{\max}}{A\omega} \tag{7}$$

where A is the area of the LCD and Φ_{max} is the maximum outgoing flux. In the dual modulation display prototype, the flux is approximately 182 Lumens (2,400 Lumens ×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 3,956cd/m^2. Actual measured luminance of this prototype is 2,700cd/m^2. The theoretical minimum luminance is less than 0.01cd/m^2, while measurements show values around 0.05cd/m^2. Such differences between actual and theoretical luminance levels are the result of optical losses and scattering in the basic optical system.

[14] The majority of the smaller prototypes were built by the Structured Surface Physics Laboratory of the University of British Columbia Department of Physics and Astronomy.

The larger display variant shown in Figure 15 uses a Christie Digital Cinema LX-1500 1024x768 projection system with over 12,000 Lumens output and a Sony Bravia KDL-40XBR4 1920x1080 40” LCD panel[15]. The prototype features a lower peak luminance of 900cd/m^2 due to the large size of the LCD but a very low black level of less than the resolution of even high end luminance meters (at minimum below 0.005cd/m^2).

Alternative prototype configurations are of course possible. Peak luminance levels can be adjusted with brighter lamps since even a one order of magnitude increase of the maximum luminance would not significantly reduce the quality of the ‘black’ state (0.005 or 0.05 are both still very satisfying “blacks” in most viewing conditions).

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 65526. Over 17,000 of these combinations would be 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 full range of JND steps necessary to provide all visible and distinguishable luminance steps in the luminance range of these systems.

6.1.1 Video Processing Algorithm for Basic Dual Modulation Displays

The image formation process of this system requires careful consideration. In the simplest case both the projector and the LCD panel are perfectly linear, and both have the

[15] This prototype was built by Dolby Canada Corp in 2007.

same dynamic range. Under these assumptions and ignoring the blurring of the projector, the target intensity can be achieved by normalizing the intensity range of the HDR image to 0...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 artefacts are relatively larger for small values.

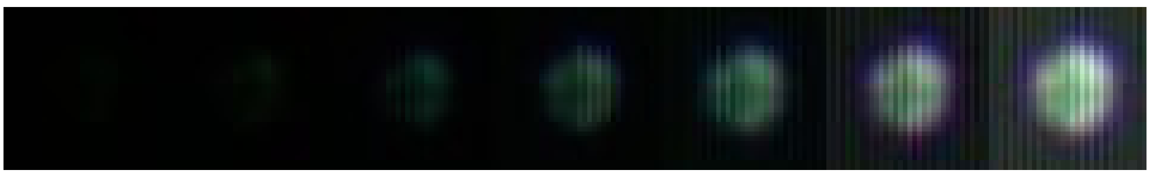

Figure 16: Point spread function of basic dual modulation prototype at different exposure levels.

In reality, neither the projector nor the LCD has a linear response, and compensation is also needed for the blurring of the projector image. This can be done in the following way: First, a simple target estimate of the projector intensity can be made based on the original image. This target estimate is then adjusted by the inverse response function of the projector to ensure that the output of the projector is linear. The projector image also needs to be blurred according to the measured PSF of the projector to account for the physical blur of the system. Finally, the pixel values of the LCD panel can be chosen such that they compensate for the blur effect.

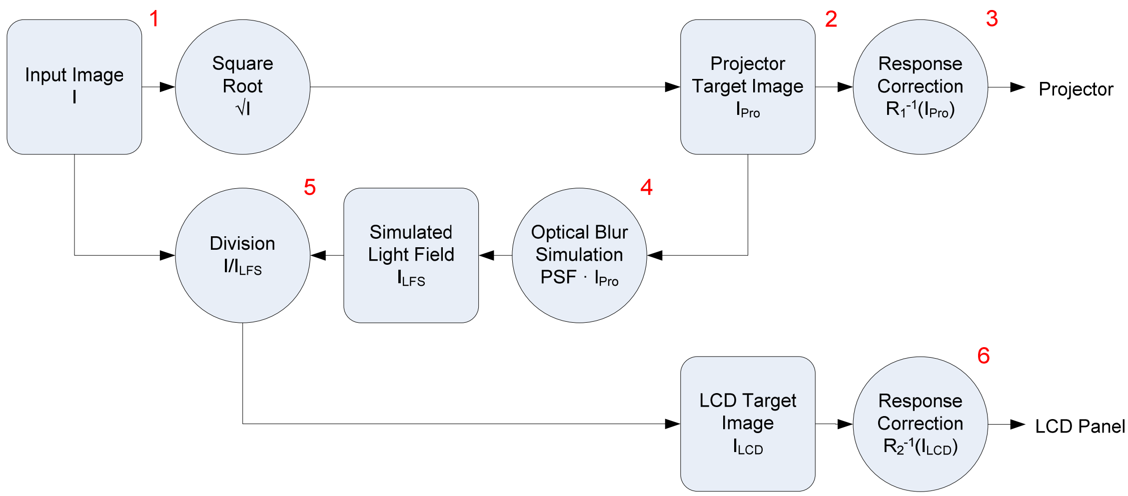

Figure 17: Video processing algorithm for basic dual modulation design.

The complete rendering algorithm therefore works as follows (also see Figure 17): The original HDR input image I (1) is adjusted by a square root or similar operations to yield a projector target image I_{Pro} (2). These intensities are mapped into projector pixel values by applying the inverse of the projector's response function R_1 (3) to ensure that the actual optical output of the projector is linear. To simulate the blurring of the projector image due to optical de-focus, the projector image I_{Pro} is convolved with the PSF of the projector (4) and the result is divided out from the original HDR image to get the target LCD transparency (5). For the final pixel values of the LCD, the inverse of the panel's response function R_2 (6) is applied. An exposure sequence of the projector PSF is depicted in Figure 16. Note the vertical lines visible in images with larger exposure times. These are the RGB sub pixels of the LCD panel. In order to speed up the computation, the measured PSF is replaced by a tensor-product Gaussian fit. The focus of the projector is usually set such that the fitted Gaussian has a standard deviation of about 2 to 3.5 pixels so that a 2D separable filter of width 13 can be used for the convolution.

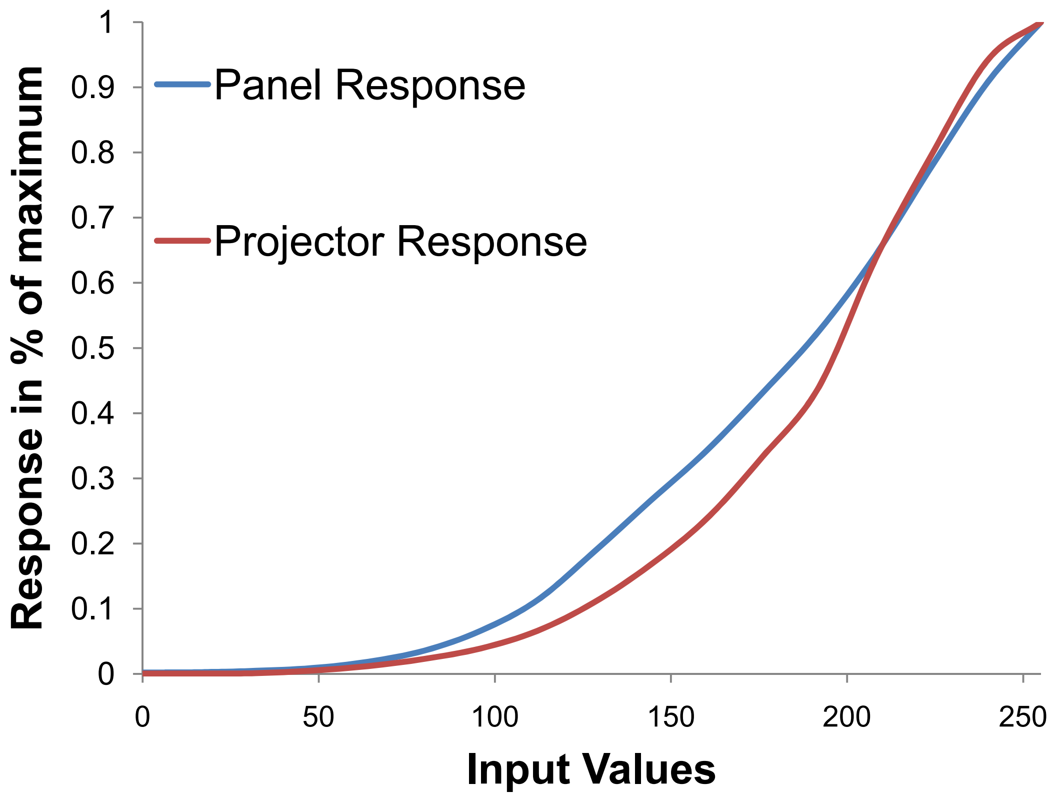

Figure 18: Response function of the DLP projector and LCD used in the smaller prototype (R_1 and R_2 respectively).

These image processing steps can be executed in software, programmable graphics hardware or other digital signal processing solution such as field programmable gate arrays.

Figure 19 shows the results of this image factorization on a portion of Paul Debevec's Stanford Memorial Church HDR photograph[16]. The top right shows a grey-scale image that corresponds to the square root of the original intensity values. Convolving that image with the PSF of the projector yields the bottom left image. This is the predicted image produced by the projector. Finally, the bottom right image is the colour 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 image processing algorithm is very similar in principle to a local tone mapping operator. This means that the LCD panel image is almost a tone mapped version of the original HDR image, although this method is clearly not designed for that purpose.

[16] The original photograph and luminance measurements were taken by Paul Debevec of the Institute for Creative Technology at the University of Southern California.

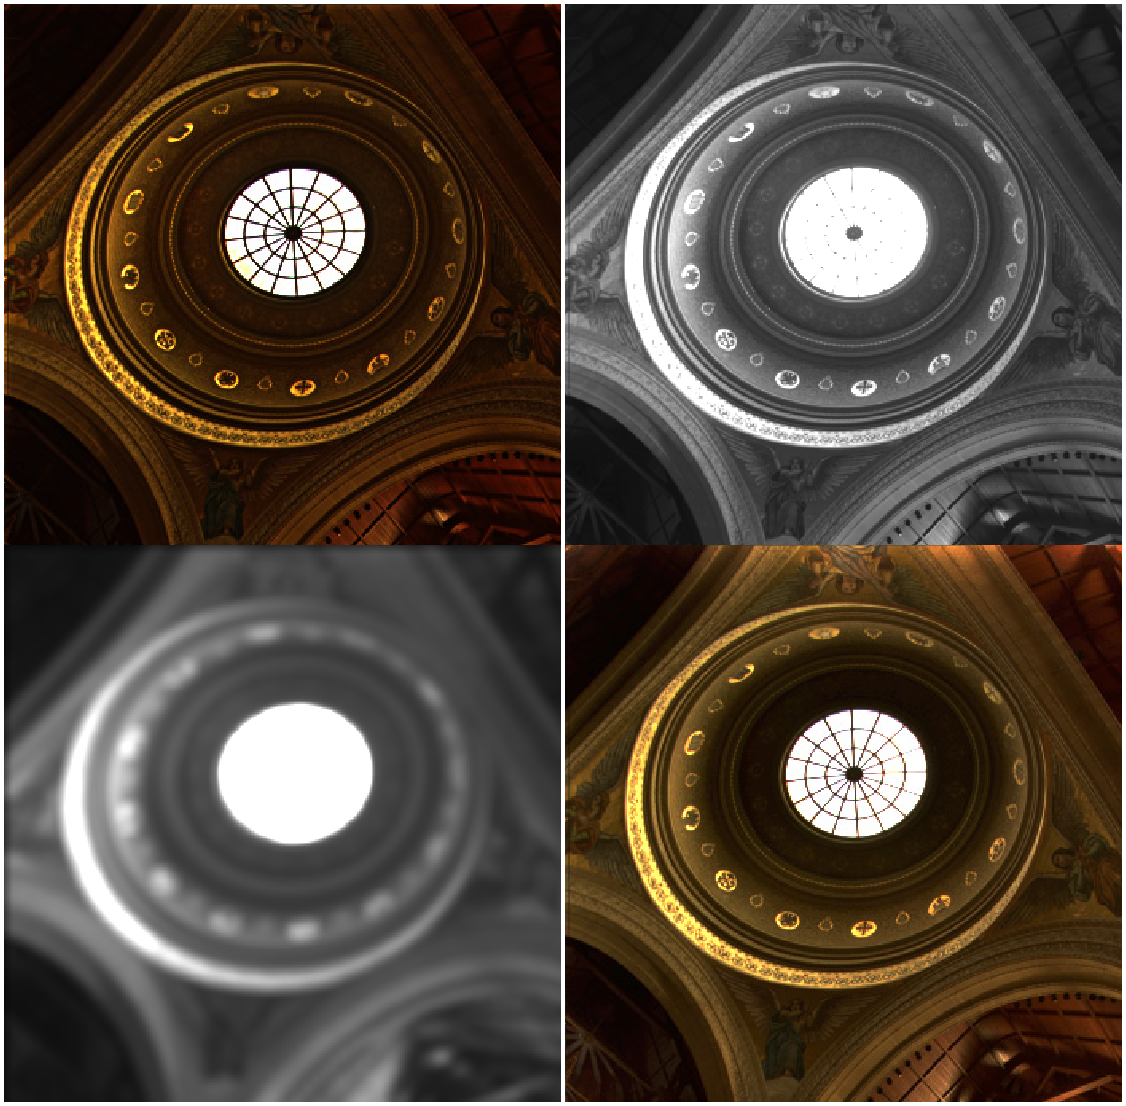

Figure 19: Example of decomposition of a HDR image. Top left: Original image section. Top right: Square root of the intensity. Bottom left: Blurred image which is predicted to be the image generated by the de-focused projector. Bottom right: Edge enhanced LCD panel image that corrects for the blurriness of the projector.

6.1.2 Performance and Limitations of Basic Dual Modulation Displays

This basic configuration achieves the HDR performance targets but has several drawbacks. In addition 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 those of 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 CRT or Plasma 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 therefore independent of the image and always very high to ensure that the lamp produces enough light for a full white image. In the dual modulation 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 dual modulation display. To achieve 10,000cd/m^2 on a 15" screen an outgoing flux of approximately 500 Lumens would be needed. Even with a very high transmission LCD this requires at least 5,000 Lumens to be emitted from the projector. In the prototype presented in section 4.1 the colour wheel 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 over 10,000 Lumens. Yet, in almost all HDR images the area that is actually at such a high brightness of 10,000cd/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 highest third of the luminance range (above 3,000cd/m^2) and that the average luminance over all images was less than 800cd/m^2 for indoor scenes and 2,100cd/m^2 for outdoor scenes. The projector dual modulation display consequently creates a factor varying between 12.5 and 4.75 too much light at any given time.

The projector also has an unnecessarily high bandwidth requirement. 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 dual modulation 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 graphics card (or a custom field programmable gate solution with twice the memory/gates) 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 dual modulation 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.

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 do not diminish the actual image quality, the basic dual modulation design provides researchers with a relatively simple to build solution with very high image quality sufficient for accurate representation of real scenes (Ledda, Chalmers, & Seetzen, High Dynamic Range Displays: a Validation Against Reality, October 2004).

6.2 Advanced Dual Modulation Displays[17]

The basic dual modulation concept described in the previous section delivers HDR performance but falls short in almost all other categories of consumer displays. It is bulky, very inefficient and challenging to scale to the large screen sizes popular in the display industry today. A different approach is necessary to maintain HDR performance and to overcome the design limitations of the basic system. As seen in section 6.1.1, software correction can compensate for a low resolution of the rear image of the dual modulation 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. This condition can be exploited by replacing the secondary modulator (the projector in the basic design) with an array of low resolution light sources. These light sources can be actively controlled to provide a blurry approximation of the image and act effectively as the secondary "modulator" in the dual modulation system.

All other elements of the system remain the same and the advanced dual modulation display has consequently three major components: a backlight with a low resolution array of light sources and appropriate optical structures, a high resolution LCD panel, and an embedded image processing algorithm to drive both modulators. The primary function of the light source array is to provide a smooth light distribution at high speed (at minimum the frame rate of the LCD unless advanced temporal dampening

[17] Dolby Laboratories Inc. offers a commercial licensing package for advanced dual modulation displays branded as Dolby Vision™ and Dolby Contrast™. See:
http://www.dolby.com/professional/video/dolby-vision.html
http://www.dolby.com/consumer/technology/dolby-contrast.html

algorithms are employed). Light Emitting Diodes are a good choice to achieve this function. As solid state devices they offer a fast optical response, stable performance and the point light source configuration necessary for dual modulation displays. Smoothness of the backlight light distribution can be achieved with a PSF of the light emitted by a single LED that overlaps significantly with the PSF of neighbouring LED. This blurs the light field created by the LED backlight sufficiently to remove any high spatial frequency information which could interfere with the high resolution pixel structure of the LCD. It also ensures that the backlight produces an approximately uniform light field if many LEDs are driven at the same level. The advanced dual modulation display can therefore operate like a conventional LCD display with a constant (and uniform) backlight if HDR image quality is not desired for a particular application.

Computationally it is easiest to model the PSF as a radial function. This matches the generally radial symmetry of the LED emission pattern. Best packaging can be achieved by arranging the LEDs into a hexagonal matrix. The spacing between LEDs depends on the desired specifications of the final display as well as the optical design of the system. For normal TV applications a spacing of 10 to 40mm provides a good compromise between high contrast performance and moderate number of LEDs (which reduces control electronic cost).

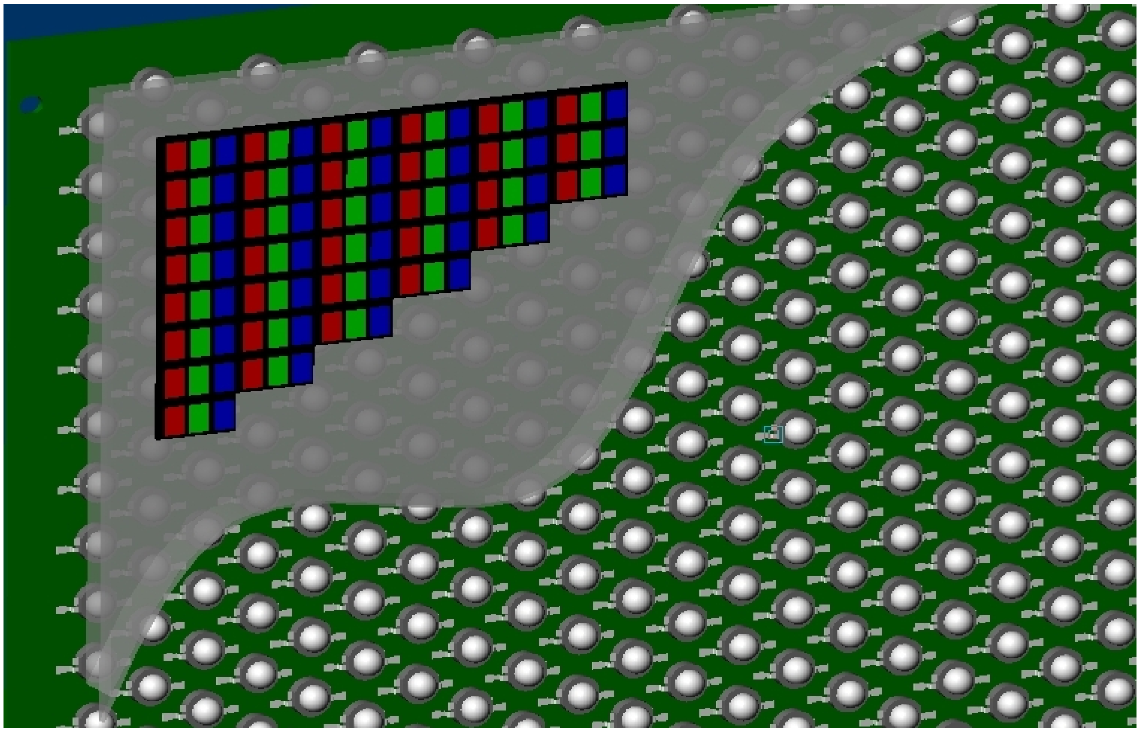

Figure 20: Pictorial representation of the LED array on a printed circuit board behind the LCD panel with RGB sub-pixel layout.

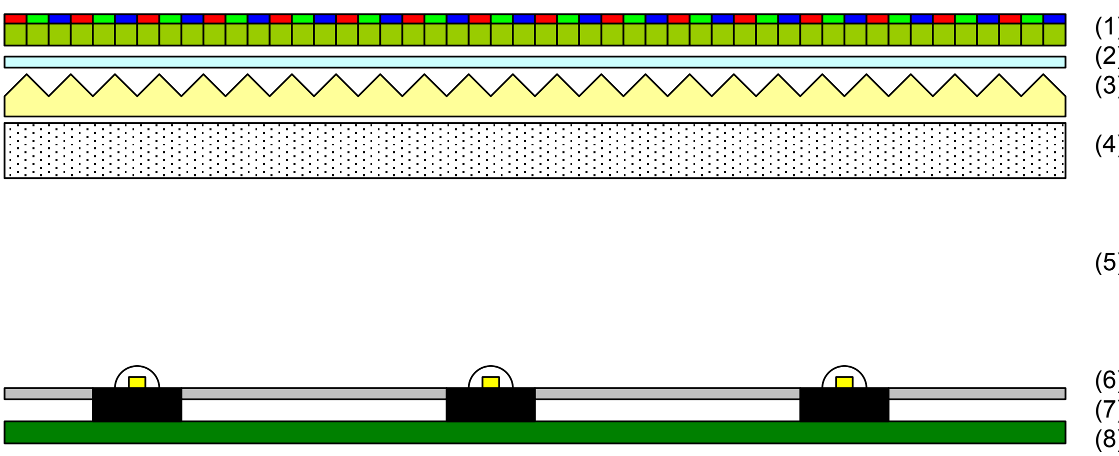

Figure 21: Optical package of the advanced dual modulation display. LCD with colour filter (1), optical films such as a reflective polarizer film (2) or Brightness Enhancement Film (3), diffuser (4), open cavity (5), reflector film (6), LED (7) and circuit board with heat sink or alternative thermal management solution (8).

This type of PSF can be achieved through a number of optical configurations. The easiest solution is to place the array of LEDs into an open cavity as shown in Figure 21. This solution makes use of conventional LCD optical films such as Brightness Enhancement Film[18], reflective polarizer film[19] and various diffuser sheets to shape the light distribution leaving the displays. The functions of these films vary but all reflect some portion of the incoming light back towards the backlight. This reflected light can then be bounced back to the LCD by an additional reflector film. Each bounce in this cavity allows some light to escape through the LCD toward the viewer. The light that doesn't escape spreads laterally in the cavity and passes through the LCD at a later bounce. In this fashion the PSF of the LED spreads spatially and can be controlled by the spacing between the two reflective layers.

---

[18] Brightness Enhancement Film is a linear prism microstructure film commercially supplied by 3M Company and other vendors. More information at http://solutions.3m.com/wps/portal/3M/en_CN/vikuiti/home/ProdInfo/Product/BEF/

[19] Reflective polarizer films are used in LCD displays to recycle light of the wrong polarization for the LCD panel. More information at: http://solutions.3m.com/wps/portal/3M/en_HK/vikuiti/home/ProdInfo/Product/DBEF/

6.2.1 Video Processing Algorithm for Advanced Dual Modulation Displays

The principal rendering algorithm for the LED-based dual modulation system is quite similar to that for the basic dual modulation display described in section 6.1. The primary difference between the two display systems from a video processing perspective is that the PSF of an LED has a much wider support than the one for a pixel of the projector in the basic configuration. This requires compensation for the low resolution of the LED array by the higher resolution LCD image.

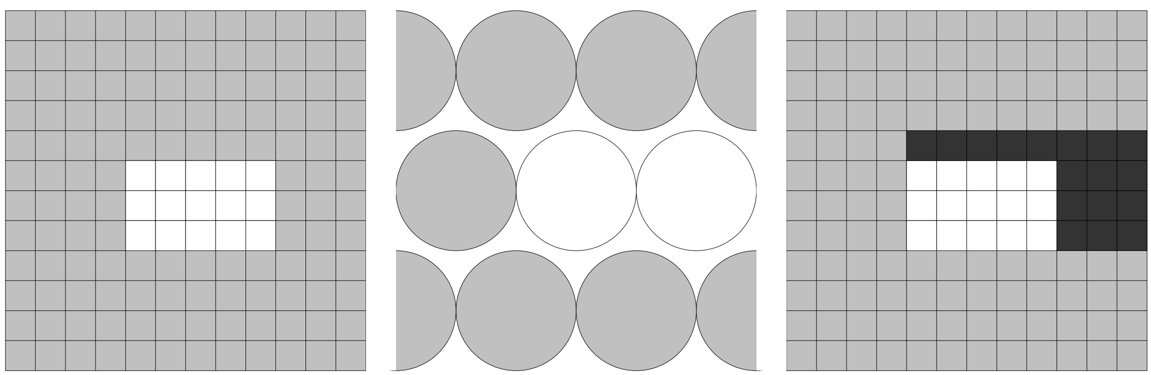

Figure 22: Schematic representation of the advanced modulation concept. The original image (left) features a bright white region and a darker grey surround. This causes the hexagonal LED array (middle) to provide high luminance in the central region and low luminance in the surrounding region. The LCD image (right) compensates for the excess luminance provided in the transition region as a result of the low LED resolution and the final image therefore matches the desired original. The actual size ratio between LCD pixel and LEDs is exaggerated for clarity.

Figure 22, shows this process for a picture of a bright rectangle upon a dull grey background. On a regular monitor, the rectangle would be a cluster of white pixels, and the grey region would be a cluster of grey pixels. When using the advanced dual modulation technology, the boundary between the white and the grey falls on several 4x4 pixel clusters, each corresponding to a single LED. For ease of visualization in this example the size ratio of LCD pixel to LED has been set to 4:1. The LED behind the cluster has to be set to maximum brightness to make the white as bright as possible (assuming that the white square is as bright as the maximum output of the display). Conversely, the LEDs behind the grey border have to be set to a low output because the border isn't very bright. The challenging region is the overlap of grey and white on top of a single LED. To achieve high brightness in the white region the LED has to be driven at maximum brightness and the grey region in that area therefore looks brighter than the grey in all other 4x4 pixel groups. To counter this effect the system sets these apparently grey pixels to a significantly darker shade of grey on the front LCD display. This reduces the final output to the same grey as the one in the neighbouring 4x4 pixel groups. Basically, a low-light LED is modulated by a medium transmissive LCD pixel resulting in light grey in the all grey 4x4 pixel groups. The 4x4 pixel groups with some part of the white square in it instead have a bright light LED modulated by a very weakly transmissive LCD pixel resulting once again in light grey.

This example illustrates the basic video processing architecture but a full implementation requires additional concepts. Unlike the example, the actual ratio between LCD and LED resolution is not 4:1 but often higher than 1,000:1. The LEDs are usually also arranged on a hexagonal grid rather than a rectangular grid to achieve a better fill factor. 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 than shown in the example. Since the supports of the PSFs for neighbouring 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 has to be implemented differently from that described in section 6.1.1.

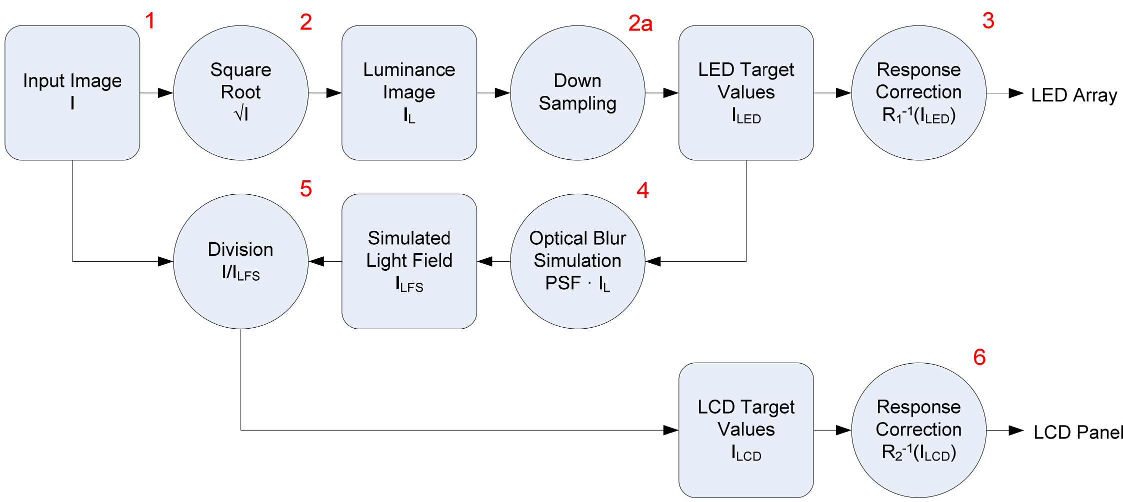

Figure 23: Video processing algorithm for advanced dual modulation display.

Figure 23 shows an expanded video processing algorithm for advanced dual modulation displays which addresses both these issues[20]. The first step in this process is to reduce the resolution of the incoming image to approximately that of the LED array. Since the LED array usually has a hexagonal layout this involves some asymmetric down-sampling. Different methods can be used for this process depending on the capabilities of the processor. Once the image has been reduced to the LED layout, the algorithm decides what portion of the total image is presented by the LED backlight. The final image is the product of the LED and LCD image so that an approximately equal

[20] More details about the dual modulation algorithm can be found in (Trentacoste, 2006) and (Trentacoste, Heidrich, Whitehead, Seetzen, & Ward, 2007)

distribution can be achieved by taking the square root (2) of the reduced resolution image. Different powers adjust the amount of dynamic range carried by the LED backlight versus the LCD panel. In some case it is advisable to decrease the modulation of the backlight to avoid artefacts. This is especially true if the LCD has a low native contrast and cannot fully compensate for an aggressively driven LED backlight. Likewise, non-linear response of the LCD and other factors such as thermal management characteristics influence this decision.

To address the wide support of the PSF, the process to determine the target intensities I_L for every individual LED is more involved (2a). Solving for LED values while taking overlapping PSF into account 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. Instead, the solution can be approximated with a single Gauss-Seidel iteration over neighbouring LED pixels. This amounts to a local weighted average of neighbouring LED target values, where some of the weights are negative. Once derived, the LED values are sent to the LED array with appropriate linearization where necessary (3).

Other solver mechanisms are possible and many variants have been proposed to address specific display designs and suppress motion artefacts. Fortunately, the LED solver step is loosely constrained because many choices for the LED values yield acceptable results. The LCD image compensates for most of the variation in the LED backlight. The same reason allows generous application of common image enhancement techniques to the LED image such as sharpening filters and other techniques designed to optimize the energy efficiency of the display.

The next step is to calculate the anticipated light field (4) generated by the LED by summing up the contributions of each LED to the total light field. To achieve this the PSF of each LED is adjusted by the drive value of the LED, the response function of the LED, and any other necessary factor such as thermal or lifetime calibration. The adjusted PSFs for each LED are then summed up over the entire image according to their geometric position. It is important that the PSF used in this process accurately reflects the effective light distribution of one LED, including any interaction with the optics and films used in the package. The result of this summation is a simulation of the light field of the LED backlight for this input image.

Finally, the LCD image can be generated by dividing each pixel of the input image by the corresponding pixel of the light field simulation (5). If the LED backlight uses LEDs with different spectral components then these need to be simulated as well and the division can occur at the colour channel level. For a white LED backlight the normally RGB input image is divided by a monochrome light field. In both cases the result of the division are adjusted by the inverse response function of the LCD.

At this point the image processing algorithm is finished and the display can show a complete frame. The LED is driven to the appropriate intensity and the LED backlight creates a luminance distribution that is similar to the light field simulation. The LCD modulates this luminance distribution such that the final output matches the input image.

Figure 24: Factorization of a HDR image. Left: Light field simulation of LED array. Right: LCD compensation image (notice the dark region around window).

In principle the algorithm can accurately reproduce any input image within the range of the product of the dynamic ranges of the LED and LCD elements. In actual devices there are limits to this approach because of the discreet step size of the LCD and the LED, the rounding errors in the simulation and other precision factors. Specifically, if both the LED and LCD are 8-bit devices with 255 discrete steps each then the total modulated range is 0 to 65025 but several steps inside that range are not achievable. For example, levels between 64770 (= 255 * 254) and 65025 (= 255 * 255) are unreachable. As a result there are luminance levels that the advanced dual modulation display cannot reproduce accurately. Fortunately, the magnitude of these gaps in the luminance range of the dual modulation display is very small at the low end of the range and increases towards the high end. The relative gap size is therefore very low everywhere along the range. This issue is generic to all dual modulation systems and discussed in more detail in section 8.2.

A further limitation of the advanced dual modulation design occurs at very high contrast boundaries. Such boundaries require the LED located directly under it to be very bright, leaving only the LCD to adjust the image. The light field from the LED is approximately constant in the region very close to the boundary so that the highest local contrast is the dynamic range of the LCD. As a result the bright side of the boundary is displayed accurately but the dark side is slightly grey near the boundary. Further away from the boundary full black is achievable because a different LED contributes to that area. The display is therefore capable of modulation over the dynamic range of the product of the LED and LCD ranges globally but only over that of the LCD locally.

Fortunately, this limitation is not relevant in conventional imaging. While our visual system is extremely good at dealing with high contrast images, section 3.2 showed that our local contrast perception is limited by veiling glare.

Using the veiling luminance model and the average constants outlined in section 3.2, it is possible to approximate the perceived luminance pattern corresponding to each image on the dual modulation display. In particular, the model provides a description of the perceived veiling luminance blur at each high contrast boundary. No degradation is perceived as long as the veiling blur is more significant than the image degradation of the dual modulation display caused by the lack of LCD dynamic range.

6.2.2 Psychophysics Based Design of Advanced Dual Modulation Displays

In order to validate the predictions of the veiling luminance masking effect, two studies was carried out with 20 participants each. All participants were between 19 and 35 years of age and had normal or corrected to normal vision. Fourteen participants were male. The studies included two comparison of real and test images. The images used was Paul Debevec's Stanford Memorial Church HDR image used previously in this chapter and a greyscale test image designed to show all possible boundaries between 16 luminance levels, each twice as high as the last. All images were shown as pairs on a basic dual modulation display.

The first study was designed to validate the general claim that HDR images appear more realistic and pleasant than low dynamic range images shown on conventional displays. For this comparison, the test image and the real scene were shown side by side in a random arrangement of low and high dynamic range settings. The low dynamic range image was presented on the basic dual modulation display by setting the rear projector modulator (i.e. the low resolution image plane) to a uniform grey level of the same brightness as a conventional LCD backlight. This leaves only the dynamic range of the front image plane for modulation of the light and thus emulates the display capabilities of a conventional LCD.

The second study was designed to provide empirical data for the degree of discomfort and unrealism, if any, associated with the blur introduced by the rear modulator. In order to vary the degree of blur in the rear modulator a basic dual modulation display design from section 6.1 was used. Optically the secondary modulator in the basic design is the equivalent of the LED array but with the benefit of control over the resolution of the rear modulator. In this study, the participants were exposed to two adjacent HDR images of the same scene (either the real scene or the test image). One of the two images was randomly chosen to be the reference image featuring a resolution match between the front and rear modulator (i.e. the same high resolution at both the LCD and the projection with pixel by pixel alignment of both images). The other side of the test image presented the same scene but with a varying degree of blur in the rear modulator. The blur was created by blurring the image data for the projector. The appropriate blur correction image was then displayed on the LCD.

A 15” basic dual modulation display was used with a viewing distance of 50cm. In this configuration a 5mm simulated LED patch corresponds to approximately 3 degrees half-angle for the veiling luminance model. The study included 4 different sizes of the low-resolution ‘pixel’ (2.5mm, 5mm, 10mm and 15mm). For each set of images, the participants were asked to provide ratings on 5 semantic differential bipolar adjective pairs (bright - dim, interesting - monotonous, sharp-smooth, pleasant – unpleasant, realistic - unrealistic). The last scale (realistic-unrealistic) was omitted for the test image. In general, the expectation is that HDR images would be considered brighter, less uniform, more interesting and more realistic than corresponding low dynamic range images. In the comparison of blurred and reference HDR image the blurred images should be perceived as progressively smoother and potentially less pleasant and less realistic with increasing blur. For the comparison of blurred and reference images there could be a small difference in the perception of brightness. The artefact halos introduced by the blurred image at high contrast boundaries could trigger a perception of brightness as our visual system considered halos of this kind to be indicators of bright areas.

The first study of the dual modulation display quality test provided the anticipated results. Low dynamic range (i.e. 8-bit) images were perceived as significantly less bright, less interesting and somewhat less pleasant and less realistic. Perception of the sharpness of the image was unaffected by the reduction from 16-bit to 8-bit as one would expect given that both images where shown at the same spatial resolution. These results align well with the larger more general study set described in chapter 4.

The comparison of non-blurred and increasingly blurred HDR images indicated that no degradation of the image was perceived even with significant blur of the rear display. In particular, the result shows no decrease in the perception of sharpness of the image in the range of blur sizes used in the test (2.5mm to 15mm). Instead, the blurred images were consistently observed as sharper than the non-blurred HDR image. This is likely the result of the blur compensation features found in the front display which might slightly overcompensate for the blur in the rear display. Such overcompensation could lead to very slight dark edges around bright areas and slightly lighter edges around dark areas. This effect is unnoticeable during close inspection of any particular area but might lead to a crisper overall appearance of the image.

All other scales (brightness, interest, pleasantness and realism) followed approximately equal trends and consequently all four scales were treated as a general quality scale in the following. The blurred test pattern was perceived to be of equal quality as the non-blurred test pattern through the entire range of increasing blur from

2.5mm to 15mm. This result is consistent with the psychological model of intraocular scattering and the assumption that even very large blur size will not lead to perceptible degradations even in fairly artificial scenes composed entirely of sharp high contrast boundaries.

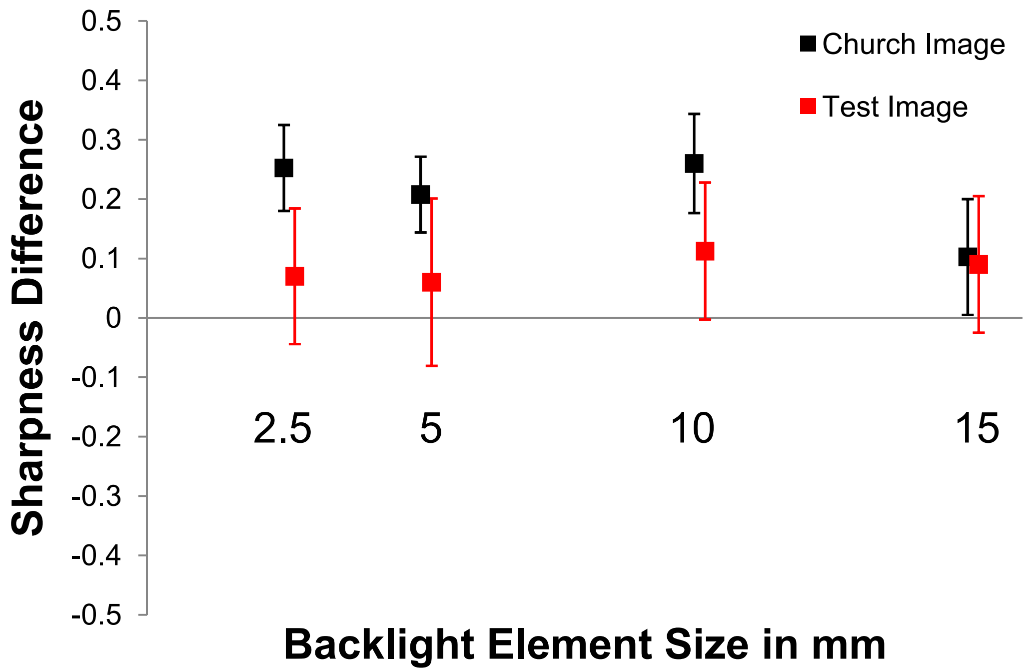

Figure 25: Sharpness difference ratings at increasing blur of the rear display. A rating below 0 (up to -0.5) indicates that the 8-bit image was perceived as less sharp than the non-blurred 16-bit image and vice versa for ratings above 0 (up to 0.5).

At small blur sizes, the Memorial Church image was perceived to have higher general quality than the non-blurred version of the image. This higher quality perception diminished with increasing blur size and at 15mm both the blurred and non-blurred images were perceived to be of approximately equal quality. The higher quality perception at small blur sizes is likely the result of sub-pixel misalignment of the two display layers which would lead to a small loss of high spatial frequency contrast in the non-blurred image. This effect does not occur in any blurred image since the effective pixel size of the rear display is so much higher than the pixel size of the high resolution display that sub-pixel misalignment becomes insignificant.

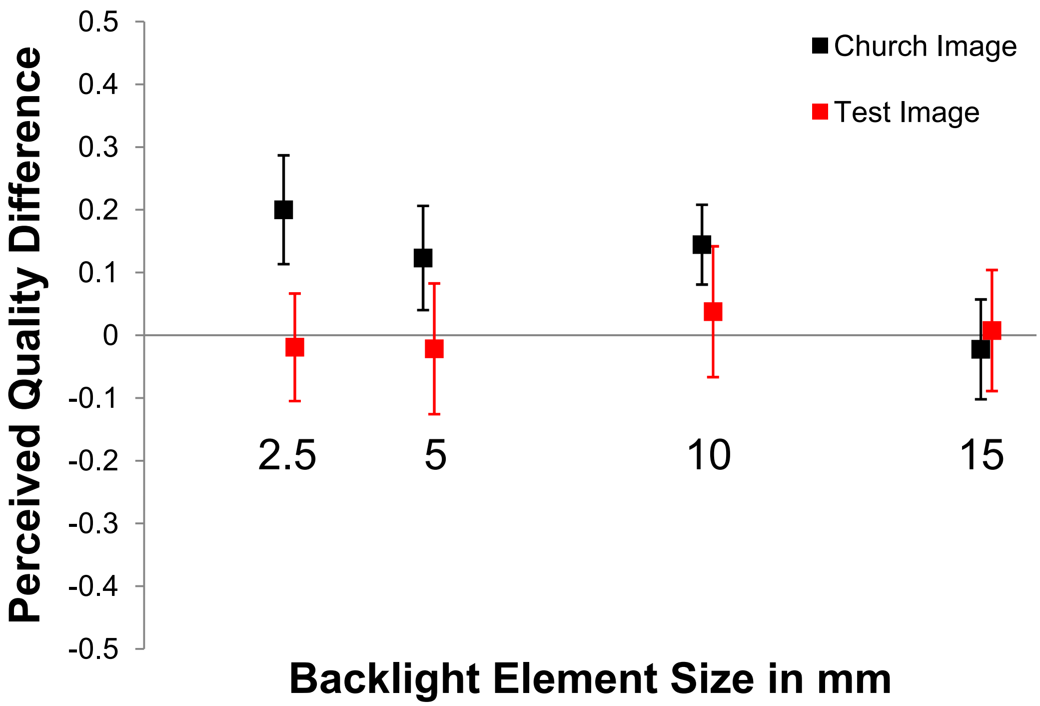

Figure 26: General quality results for experimental study.

The test results provide statistically significant support for the postulate that 2.5 to 15mm blur of the rear display does not degrade the perception of sharpness or general display quality. The 15mm upper limit provides a design target for all advanced dual modulation displays. At the viewing distance of the study this simulated LED spacing corresponds to approximately 8.5 degrees. Exact specifications are challenging to derive from this number because of the differences in PSF shape for different designs.

6.2.3 Performance and Limitations of Advanced Dual Modulation Displays

The advanced dual-modulation approach offers many improvements over conventional displays and its basic precursor. Like the basic approach, the dynamic range of the display is greatly increased and the contrast ratio can approach infinity if the light elements are allowed to turn off for dark regions. Amplitude resolution is likewise increased dramatically though the very low spatial resolution of the backlight array makes it challenging to predict the exact amplitude resolution of the overall system. Even though the light elements are usually controlled digitally and have distinct output levels, the high degree of spatial blur in the backlight means that the backlight light field is effectively a smooth analog distribution. The amplitude resolution of the advanced dual-modulation system is therefore best described as the combination of the digital control range of the LCD and the analog local multiplier provided by the backlight unit. This configuration is well suited for the requirements of the HVS and can result in perceptually quantization free images.

The unique architecture of the advanced dual-modulation system delivers further advantages. Dynamic drive of the light elements means that the power consumption of the display varies with image content as the backlight array only illuminates bright sections of the image. In comparison, a conventional static backlight consumes the same energy corresponding to the highest luminance of the display regardless of image content. This effect is described in more detail in section 8.3.

Figure 27: BrightSide Technologies DR37 high dynamic range display using advanced dual modulation display principles. The display offers a peak luminance of over 3,000cd/m^2 with a 37" 1920x1080 panel.

The use of colour LEDs (or integrated RGB LED packages) can significantly improve the colour gamut of the display by providing more saturated primaries. Some of these benefits could also be obtained with a non-modulated RGB LED backlight but the segmented control of the dual-modulation approach provides the ability to achieve even purer colours in larger regions. In a static backlight design a small portion of light always leaks through the LCD panel. For example the red primary of a static RGB LED backlight therefore always contains some contamination from the blue and green LED. A dual-modulation backlight would reduce or completely turn off the blue and green LED in a red region and thus achieve even better colour saturation.

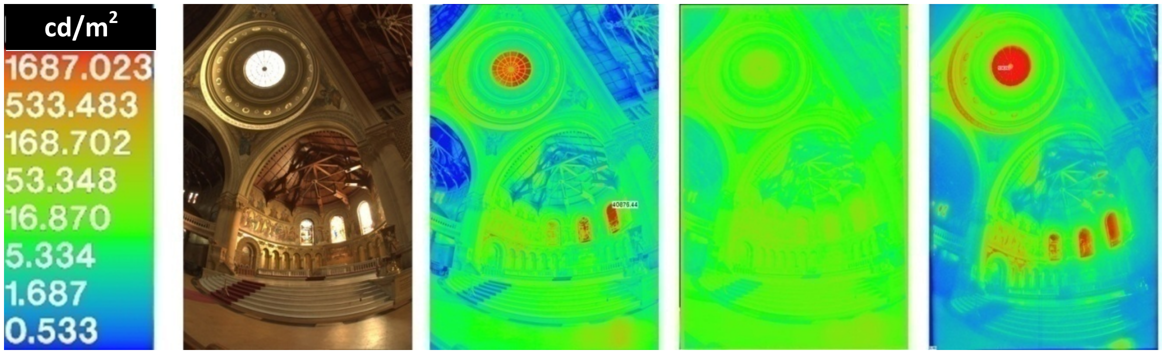

Figure 28: Achievable dynamic range on an advanced dual modulation display. Logarithmic false colour luminance scale shown on the left. The first two images from the left represent a photographic (tone mapped) representation of Stanford Memorial Church and the actual scene luminance levels in false colour. The third image from the left represents the luminance levels of this image shown on a conventional LCD display without local modulation. The image on the right is a false colour luminance map of the image shown on a dual modulation display using 768 LED behind an 18” 1280x1024 LCD panel.

The LED array can also be used to reduce perceived motion blur and compensate for other temporal artefacts. The LEDs have a very fast response time so that flashing them can increase the effective response time of the LCD (the LCD is a passive device and the effective light emissive period is defined only by the LED flash time). This creates a strobe-like effect similar to CRT technology and greatly reduces perceived motion blur because the LCD is not illuminated during transition times (Fisekovic, Nauta, Cornelissen, & Bruinink, 2001).

Overall, the advanced dual-modulation display design offers a compelling path to HDR imaging for front-view displays. The solution provides high contrast, potentially very high luminance and good amplitude resolution. At the same time the design is energy efficient and offers additional benefits in traditional LCD weak spots such as motion blur and colour gamut. Different advanced dual-modulation designs have been demonstrated including a 37” display with 1380 single-LED-elements and a peak luminance of over 4,000cd/m^2. Many variants of this design are currently entering the market and the technology is predicted to become a major part of the LCD consumer industry in the future.

7 Dual Modulation Projection Technology

The advanced dual modulation design described in the previous chapter achieves all but one of the objectives laid out in chapter 5. The display can deliver HDR imagery with both high luminance and contrast, good amplitude resolution and many other compelling image quality features. Even better, it does so within the envelope of a conventional LCD with respect to cost, physical size and in fact improves energy efficiency compared to its non-modulated peers. Unfortunately, the use of a conventional LCD panel also prevents it from scaling to very large screen sizes. While LCD panels are increasing in size each year, there are currently devices available in a size class that would provide a very large screen immersive experience. The only solution to this final barrier to the ultimate HDR experience is a projection design.

7.1 Basic Dual Modulation Projection Systems

A direct transfer of the advanced display principle to projection is not practical due to the limits on miniaturization of small high power light source arrays. A projection system illuminating a 40” diagonal screen would require the same number of LEDs with the same or higher light output as a 40” advanced dual modulation display – but on a 1” diagonal surface of the projector chip! Current LED based projectors struggle to achieve even a modest fraction of the output flux of conventional Ultra-High-Pressure lamps even without the complexity of local modulation of the LED source. Since the target for high dynamic range is to achieve much higher than conventional image luminance, the required LED output would need to increase by more 20 times. As LED performance increases further, it might be possible to manufacture a high density, high power LED array but this is unlikely to happen in the short term.

If light source arrays cannot be miniaturized then the alternative is to return to the basic dual modulation design described in section 6.1. A conventional projection system uses a light source followed by appropriate optics to guide the emitted light onto a transmissive or reflective image modulator. The modulator adjusts the intensity of light of each image pixel. The modulated light is then sent through projection lenses onto a screen. The dynamic range of such a system is given by the modulation range of the image modulator. Scatting in the optics of the light engine often reduces the modulation range further.

Similar to the basic display design, the basic dual modulation projector augments the conventional light engine with a second image modulator. The light from the lamp is now modulated by both the original and the secondary image modulator. If the contrast ratios of the original and secondary modulator are c_1 and c_2, respectively, then the effective dynamic range of the output image will be $c_1 * c_2$. Because of the unique image processing required in a dual modulator design with different spatial resolution among the modulators, the full dynamic range of $c_1 * c_2$ is not always available for some images. This limitation and related solutions are explained later in this section.

The basic dual-modulation projection system can be designed with a variety of components. The two modulators can be LCD panels, LCoS panels, digital mirror devices (DMD) or most combinations thereof. Independent of the combination, one of the modulation stages can be of low resolution to increase the optical efficiency of the system for the reasons already discussed in section 6.1.

At this point the similarities between the basic display and projection design end. In the display configuration the main modulator is a large LCD panel. Optical coupling of the two modulators is therefore very easy and light can be diffused for viewing directly at the main modulator. For the projector, the main modulator needs to be approximately the same size as the low resolution modulator as dictated by the overall optical system of the projector. Light passing through both modulators also needs to be further channelled through a projection lens. This imposes a requirement of very low beam spread as the light passes from the low resolution modulator to the main modulator. A diffuser is therefore not a practical choice to achieve the optical blur necessary for dual modulation. The solution is to place the low resolution modulator slightly out of the optical focal plane of the projection lens system. This often means that it is easier to place the low resolution modulator after rather than before the main modulator in the optical path. While this appears counter-intuitive, the lack of diffusion in the system means that the physical order of the two modulators is irrelevant from an optical imaging perspective.

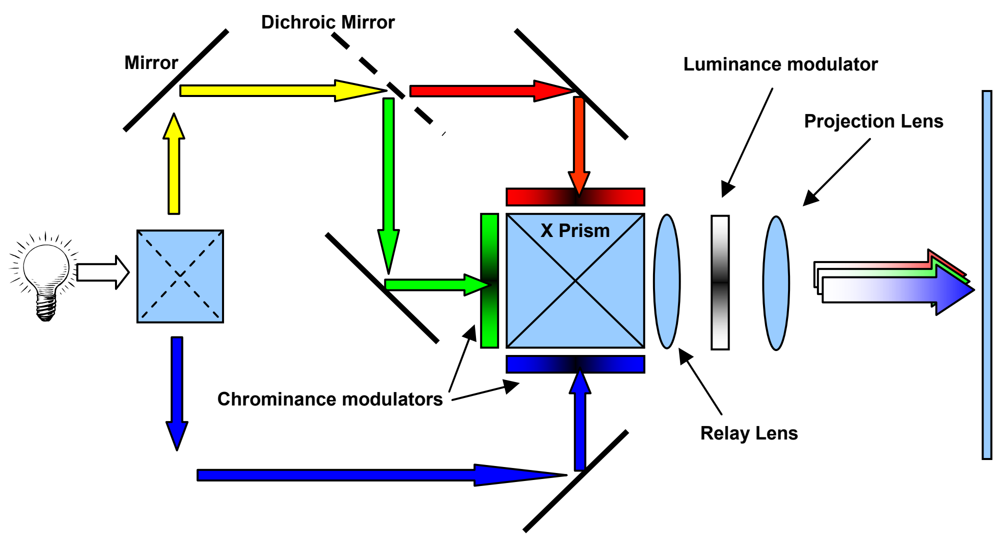

Figure 29: Implementation of the basic dual modulation projection design. This variant uses a standard 3-LCD projection system and a secondary low resolution luminance modulator.

The secondary modulator can be of very low resolution, thus hardly reducing system efficiency (e.g. as few as 1,000 pixels for a 1080p projector). A modulator of this kind would have to be customized, but can have a fill factor close to 99%. In the case of a projector with three separate colour channel modulators, it can be placed before or after the recombination of the three high resolution modulators for each of the red, green and the blue channels. With the exception of the optical design differences outlined above, this is the closest equivalent to the basic dual modulation display design.

A variety of combinations of different modulators for chrominance and luminance is possible. Many conventional projectors today use three modulators of a signal colour each rather than a single three colour modulator. Such a design can be augmented with a new additional luminance modulator (effectively a 4^{th} modulator affect all three original colour modulators). Alternatively, a new low resolution modulator can be added independent to each colour path so that the final projector has effectively 6 modulators. This design enables very pure colour reproduction. The low resolution panels can be placed directly next to their corresponding high resolution panel. This eliminates the need for additional optics. The amount of introduced blur can be controlled by adjusting the distance between each pair of panels. The low resolution panels can be either driven in parallel or controlled separately.

Another design uses binary modulators such as DMDs or digital LCoS imagers which both lack the ability to adjust the amplitude of each pixel. Light from such a pixel is either “on” (light leaving the projector towards the screen) or “off” (light going into an absorbing cavity inside the light engine). Rapid switching between these two states generates the greyscale of the projection system due to temporal integration in the viewer’s eye.

Combining two frequency modulation based devices in series does not yield the dual modulation effect described above. If the two modulators are in phase, then they act effectively as a single modulator. If they are out of phase, then no light is transmitted at all. Either way, no gain in dynamic range is achieved. Sophisticated drive schemes and coupling of such digital modulators can overcome these issues and allow dual modulation even with two stages of digital modulation. In particular, the lower spatial resolution of one of the modulators can be utilizes to enable this combination.

Another alternative is the coupling of a binary modulator with an amplitude modulating component such as an LCD or analog LCoS. Overall, the DMD/LCoS or DMD/LCD hybrid design is probably the least desirable of the combinations due to the necessity of adding a polarizing component to a non-polarized system. While a good choice of hybrid components can achieve acceptable efficiencies, the differences in the drive techniques of the two modulator types add unnecessary complexity to the design.

7.1.1 Video Processing Algorithm for Basic Dual Modulation Projection Systems

The image processing algorithm required to drive dual modulation displays is very similar for all implementations. In each case, the algorithm needs to consider two image modulators with different spatial resolutions and colour capabilities. The algorithm also needs to compensate for the optical blur of the low resolution modulator. The degree of blur can be measured and characterized by a PSF for each low resolution pixel which is then expressed in corresponding pixels of the higher resolution modulator. This PSF is one of the core parameter of the image processing algorithm together with the response curves and spatial resolution of the two modulators. These values vary for the different design implementations but the remainder of the algorithm implementation is design independent.

The luminance of the desired output image is distributed between the luminance and the chrominance modulators by applying a square root function, resulting in the drive values for the low resolution display. Given these values and the parameters mentioned above, the optical blur can be precisely simulated at high resolution. A division of the desired output image by this blurred image results in the necessary compensation mask that is displayed on the high resolution modulator. In a closed-loop feedback, the compensation image is analyzed for saturated regions in which the corresponding drive values on the low resolution panel is then locally enhanced to deliver enough light. Figure 30 shows the processing steps for this algorithm.

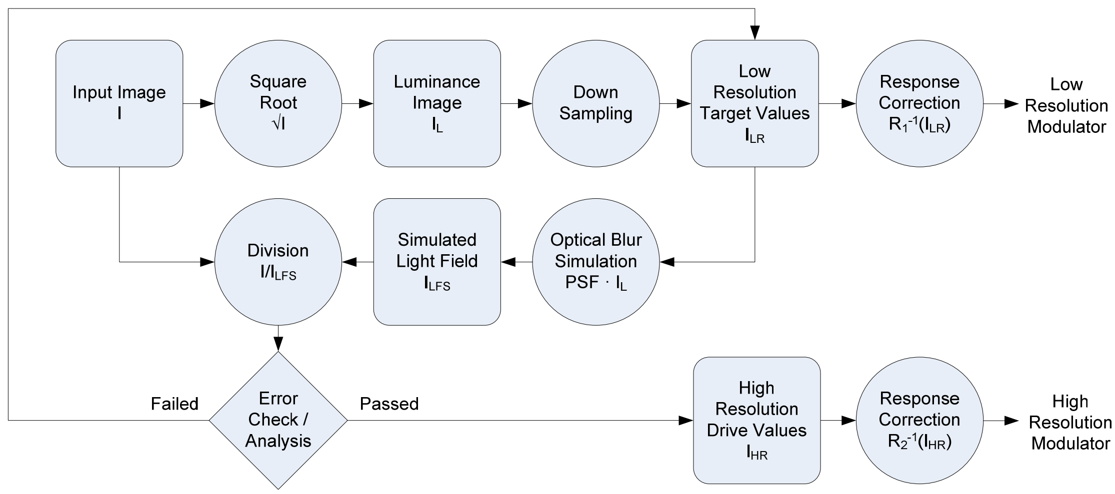

Figure 30: Video processing algorithm for basic dual modulation projector.

A careful hardware design in combination with precise modeling of the PSF ensures that artefacts are always smaller than the veiling luminance halo described in section 3.2. In a system with a higher resolution luminance modulator, the pixel ratio of the luminance to the chrominance modulator needs to be selected appropriately. Artefacts that cannot be compensated for are, at worst, marginally visible and only appear at boundaries between fully saturated regions at high spatial frequencies.

7.1.2 Performance and Limitations of Basic Dual Modulation Projection Systems

Figure 31 shows photographs of a basic dual modulation projector designed with the principles described in this section. The system provides a significant expansion of the dynamic range through serial combination of light modulators of different spatial resolution. The prototype is based on a typical projector that utilizes three transmissive LCD panels. A set of two panels per colour channel were driven in concert (i.e. a corresponding second low resolution modulator for each of the three transmissive panels). The simultaneous contrast ratio was improved by more than one order of magnitude over that of the original projector. The resulting simultaneous contrast matched the theoretical product of the low (18:1) and the high resolution modulator (155:1) within a 5% range for a total simultaneous 5x5 ANSI contrast of 2695:1.

The basic dual modulation projection design delivers the anticipated HDR image quality and can be scaled to large screen sizes. Unfortunately it inherits the same energy efficiency challenges previously encountered in the basic display design. Overall, the design is unlikely to be efficient enough for commercial large screen applications such as cinema.

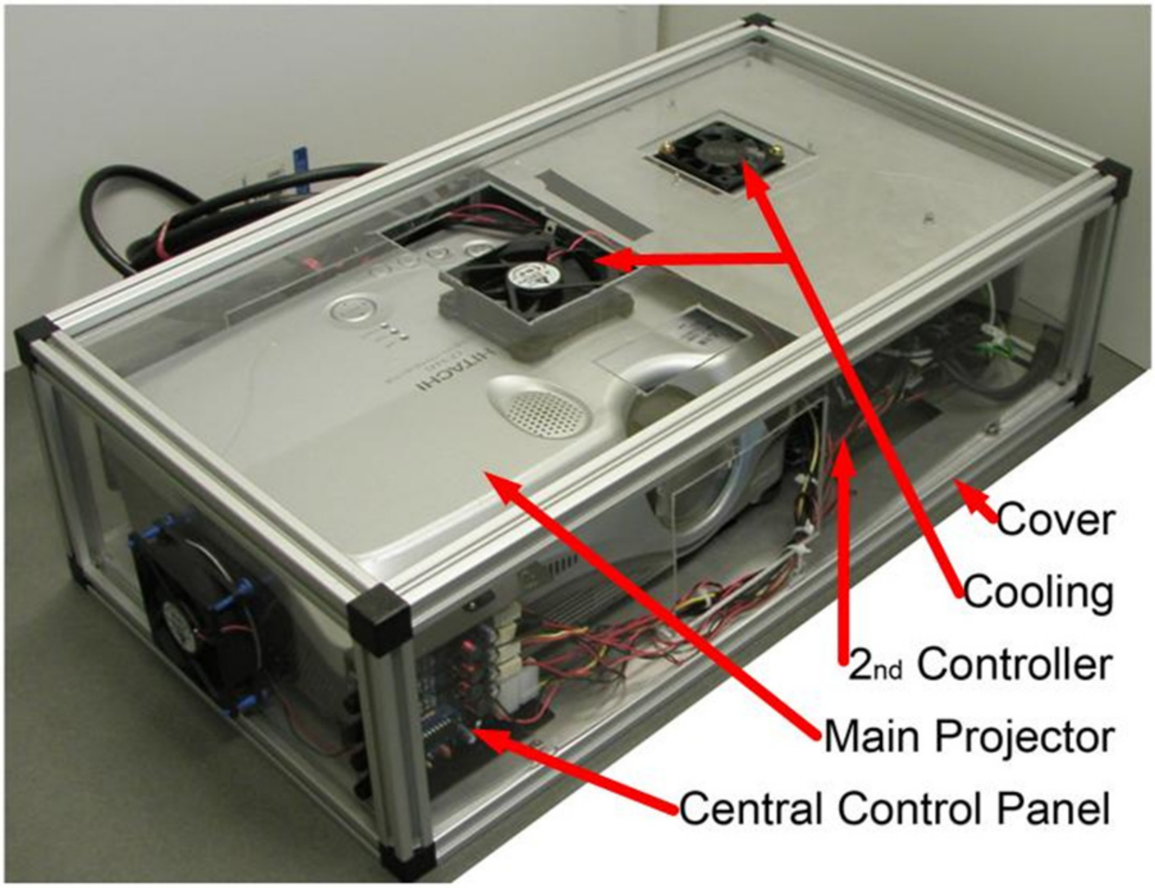

Figure 31: Second generation prototype of basic dual modulation projector system. A conventional Hitachi LCD projector is augmented with a secondary modulator inserted into the optical path of the main unit[21].

[21] The prototype was fabricated by Mike Kang and Gerwin Damberg of Dolby Canada Corp.

Figure 32: Comparison of a conventional and dual modulation projector image. The left image is created by a conventional 3 LCD projector and the right image by the same projector upgraded with an additional low resolution luminance panel. Note the increased contrast of the system and higher definition in the dark regions (Actual photographs).

7.2 Distributed Dual Modulation Projection Systems

As described in previous chapters, dual modulation systems offer significant visual advantages over conventional displays and can be successfully extended to front-projection devices. Nevertheless, the systems described so far fall short of the objective of a large scale HDR experience for different reasons. Front view display systems discussed in chapter 6 can achieve very high dynamic range and high luminance but cannot be scaled to the desired size. The basic projection configuration described in the previous chapter supports large screens but its comparable low efficiency makes it an unlikely candidate.

Conventional cinema projectors use very bright lamps to reach the comparably low luminance level of cinema on a large screen. The gating factor for such lamps is the thermal limit of the projection system. For a full HDR experience, the screen luminance needs to be increased at least fivefold which would also increase the thermal load in the projection system by the same multiplier. This situation is aggravated by the fact that the second modulator in a basic dual modulation projection system is placed within the main projector. That modulator absorbs an additional 10 to 30% of the light emitted by the projector lamp and thus contributes significantly to the heat generation in the system. The resulting increase in thermal load is unlikely to be acceptable for conventional projection systems which are already operating at the peak of their thermal envelop. This chapter describes a distributed dual modulation projection concept that reduces the thermal load of the system and therefore allows for large scale HDR projection.

The advanced dual modulation display configuration in section 6.2 overcomes the thermal and energy challenges of the basic display design by replacing the light source with a modulated array of light elements corresponding to different regions of the display. Such a design is energy efficient but for reasons already discussed not feasible for conventional projection system. Fortunately, a similar distributed light source effect can be achieved with a slight modification of the original concept.

The distributed dual modulation projection concept re-captures the efficiency benefits of the light element array but adapts it to front projection without the need to miniaturize the array. This is achieved by distribution the light elements over a large area and matching them up with equally distributed primary (non-emitting) modulators at each location. Conceptually, this makes the design similar to a tiled array of projectors but with some key differences.

The concept of tiling multiple projected images is well understood (Bordes, Bleha, & Pailthorpe, 2003) but suffers from the need to calibrate each sub-image to avoid non-uniformities at the seams of two tiles. Many techniques have been proposed to achieve this calibration but they all require a reduction of sub-image intensity (i.e. aligning all projectors to the lowest common intensity). Combined with the already low efficiency of high resolution projection systems this generally leads to very inefficient designs. Moreover, since the calibration methods inevitably reduce the dynamic range of each sub-image, the overall dynamic range of a tiled projection system is even lower than that of a conventional projector. The distributed projection system therefore needs to allow for easier calibration than conventional tiled projection systems. It also needs to support an increase rather than a decrease of overall dynamic range.

7.2.1 System Architecture of Distributed Dual Modulation Projection Systems

The distributed dual modulation projection system has three core components: A processing unit, a large number of small projector elements and a screen. The video processing steps required for this concept are significant and require a unique parallel processing arrangement to handle the large number of individual devices and the high overall resolution. For this purpose the system should contain a central processing unit, optional sub-node processors and finally the individual processors at each projector element. The screen can be a conventional cinema or home entertainment screen and does not influence the design of the system.

Each projector element contains a low number of single colour light sources (possibly just one) which can be addressed over a reasonable range of brightness. If the element is small enough then the light source can be single colour LED. A primary passive modulator is placed in front of the adjustable light sources. The modulator can be a LCD or digital mirror device. The modulator can have a low spatial resolution and can be monochrome. The easiest configuration is a single-primary (i.e. red) LED as the light source coupled with a simple low resolution LCD or LCoS modulator. In addition the projector element includes drive electronics to control the LED and primary modulator as well as simple processing electronics to establish the appropriate image data for LED and LCD. Optionally a lens design can be mounted onto the front of each projector element though as discussed later this is likely not necessary. The refresh rate of the LCD or digital mirror device could be lower than conventional as each projector element is only responsible for a single colour and field sequential frame rates are thus not required (compared to a conventional digital mirror projector using a colour wheel with three filters). In principle the LCD refresh rate can be as low as 24Hz (cinema film frame rate) with the LED flashing two or more times per cycle to achieve an effectively higher flicker frequency.

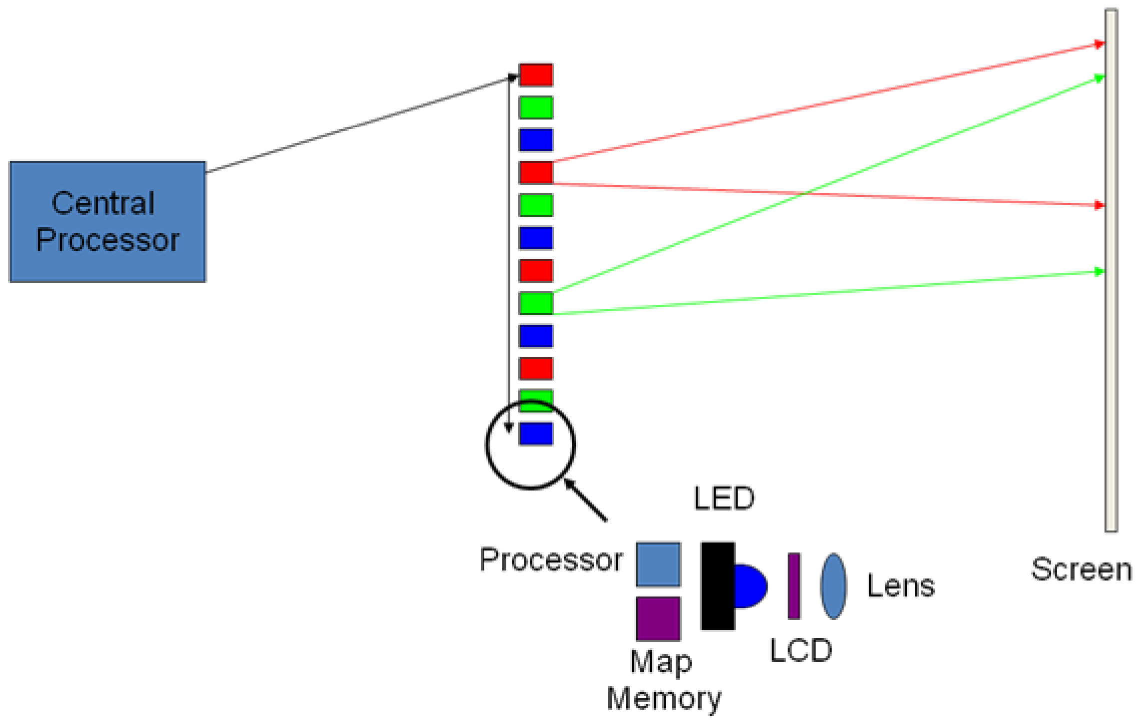

Figure 33: System architecture of distributed projection concept. The central processing unit drives an array of projector elements. Each element contains a small processor, memory for the responsibility map, LED light source and LCD modulator.

During operation each projector element adjusts the light output of the LED based on incoming image data. This light is further modulated on a pixel by pixel basis by the LCD. Transmission losses can be very low because of the low resolution and lack of colour filters on the LCD. The twice modulated light leaves the projector and creates an image on a small section of the overall projection screen (called an “image element” in the following to separate it from the overall screen image). This process is very similar to that of advanced dual modulation displays described in section 6.2 with the difference that each projector element only represents a single LED cell (corresponding to a single LED and the LCD pixel directly on top of it in an advanced dual modulation display design). Different projector elements should have different colour LED. At minimum red, green, blue elements are needed but additional colors such as yellow and cyan can be added to expand the colour gamut of the final image. Figure 33 shows a schematic representation of a projector element using a transmissive LCD modulator.

A large number of projector elements can be mounted together as a larger projector array. Each element creates a small image element on the main screen. Unlike conventional tiled projection images, these image elements need to overlap substantially so that each point on the screen is illuminated by light from at least one projector element of each colour in the overall set. More overlap is desirable and ideally each point on the screen should receive light from 5 to 15 projector elements. The orientation of each image element and the layout of all the images elements with respect to each other are not important. This significantly reduces the need for precision in the installation of the projector elements. To achieve this overlapping effect, a large number of projector elements replace a single conventional cinema projector. Fortunately, the projector element components are very inexpensive and even 10,000 projector elements are estimated to be cheaper than a single cinema-grade projector.

The overall image on the screen is now composed of thousands of image elements at different primaries. Overlap of the image elements ensures that each section has all required colour primaries. The number of primaries for the system is at this point an arbitrary design choice as it is easy to add additional projector elements with additional colour primaries (e.g. cyan LED, amber LED, etc). The overlapping pattern also ensures that seams between image elements are invisible. In a conventional tiled projection system the two image elements only overlap slightly at the boundary region which causes visible seams and transitions to appear unless the two projectors are perfectly calibrated. The massive overlap of the distributed dual modulation system ensures that no seams are visible. There are in fact thousands of very small randomly oriented "micro-seams" between different image elements. No larger seams exist so that the overall impression is that of a seamlessly blended image.

Similarly, the effective spatial resolution of the system benefits significantly from the overlap configuration. The effective resolution of the overall image is very high even if the resolution of each projector element is low. A high end digital cinema projector has a resolution of 2 million pixels (1920x1080 pixels). Assuming an array of 9,000 projector elements (3,000 red, 3,000 green, 3,000 blue), each with an 320x240 pixel LCD and on average an overlap of ten image elements per point, the overall system would achieve an effective resolution of 23 million pixel (=3,000 x 320 x 240 / 10). In reality, the actual resolution of the system is impossible to calculate because the layout of the image elements is random or near random. The calculation above assumes that the pixel structure of the image elements lands on a common grid on the screen. In reality each pixel is in a slightly different position than any pixel from an overlapping image element. This causes a type of spatial dither which effectively increases the smoothness of the image as the number of overlapping elements increases.

The system achieves high dynamic range in two stages: at each projector element and on the overall screen. Locally at each projector element the dynamic control of the LED ensures that the projector element itself can adjust from no light output to high output without sacrificing image fidelity. Since the element only has a single light source it cannot inherently produce high dynamic range but control of the LED ensures that the element never emits more light than strictly necessary for the brightest feature in the small image element that it projects onto. At the global level this causes a dynamic range increase very similar to that found in advanced dual modulation displays. Where light is needed on the overall image the corresponding projector elements deliver bright light. In dark regions the projector elements remain dim. This split architecture requires sophisticated image processing and calibration. The following section provides more detail for these processes.

7.2.2 Calibration and Setup of Distributed Dual Modulation Projection Systems

Video processing for the system is split into initial setup and ongoing operation. Initial setup includes a number of steps that are crucial to achieve good image quality. All setup operations can occur in non-real-time and do not need to be part of the overall system architecture. For example, the installation process of the system could include a manual setup stage using portable equipment that is shared by multiple projection installations or provided by the manufacturer of the device.

During the setup process it is necessary to gather information about each projector element, their relationships to other projectors and to the main screen. This information is stored after calibration in the central processing unit and portions of it at each projector element. The first set of calibration data can be obtained during manufacturing of the projector element and does not require local measurements. That data includes the relevant optical specifications of each projector such as LED response curve, LED spectrum, LCD response curve, LCD filter spectrum (if any), LCD resolution, PSF of the LED onto the LCD, and any other optical aspects that influence the performance of the projector element. This data should be stored locally on the projector element and is used by the element to ensure that it produces an appropriate image element for a given set of incoming image data. In essence this data is identical to the calibration data of any other dual modulation display with the difference that only a single light source needs to be considered.

The second calibration data set describes the spatial and directional relationship between the projector elements in the array. Fortunately, the actual physical relationship between the devices is irrelevant. Only the position and shape of the image elements influences the overall image. Regardless of the physical placement of the projector elements, the calibration process is therefore the same. The projector array is physically installed and the elements are oriented approximately towards the screen so that all image elements appear on the main screen and the main screen surface is evenly covered by image elements of each colour primary type. A high resolution still image camera is placed centrally in the viewing environment and adjusted to capture the entire screen.

With the complete array and camera in place, one projector element at a time is turned on with a full white image. This produces a corresponding white image element on the main screen. For each element the camera captures an image of the entire main screen. The resulting image is predominantly black with just a small white region corresponding to the image element. The white region is an arbitrary quadrilateral as the projector element is most likely not normal to the main screen. The luminance distribution in the white region is also likely non-uniform due to both the geometric relationship between the projector element and the screen, as well as the likely non-uniform PSF of the LED in the projector element.

This completes the calibration of the system. The camera images can be linearized to yield a map of the illumination capability of each projector element with respect to the main screen. These maps describe the portions of the main image to which the projector element contributes. Offline processing can now be used to create a responsibility map for each projector element and store it in local memory.

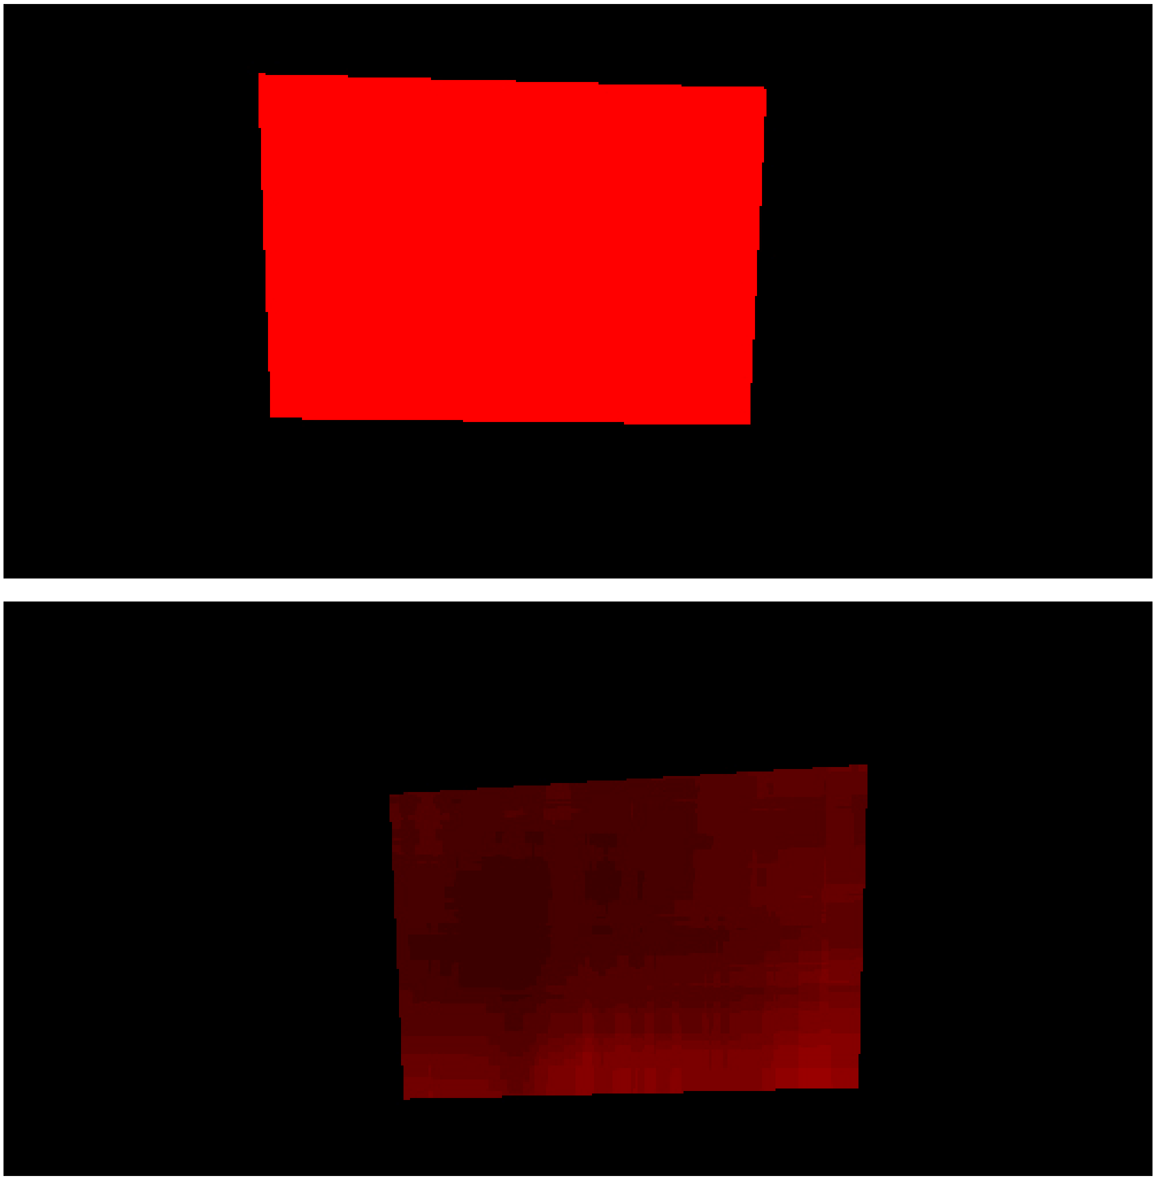

Figure 34: Example responsibility map for a single red projector element. Stylized map shown on top (no overlap and therefore uniform responsibility). The lower image shows an actual map with variation resulting from the overlap of over 100 projector elements.

In black regions the image values are zero and thus the projector element has no responsibility for image content in those regions. In the small white region the projector element can be assigned responsibility based on the intensity of the white region per pixel compared to the intensity of all other projector image elements that illuminate the same pixel. For example, if three project elements create image elements that overlap on the same point of the main screen and the intensity of the camera image for the three projectors are 50, 100 and 50 then the first projector element would be responsible for 25% (= 50 / (50 + 100 + 50)), the second for 50% and the third for 25% again. This process generates a map for each projector element that can be expressed as an array of responsibility values between 0 and 1. For the sake of explanation it is easiest to think of this map as an image at the resolution of the overall projection system stored at each projector element. In a real implementation the resolution of the map would likely be higher in the region of interest to leverage some of the resolution benefits described earlier and much lower resolution in the very large area without responsibility for each projector. A vector based memory format or some basic form of compression of the zero-region can create a very efficient local memory representation of this data.

An optional step is to add an overall adjustment factor or map to the main image and multiply it directly by each of the projector element responsibility maps. This reduces computational complexity during operation. For example, some projection systems use a compensation mechanism that dims the central portion of the image to reduce vignetting (screen luminance fall-off towards the edge of the screen). Such a compensation map can be overlaid onto the overall main screen before local responsibility maps are computed. Instead of applying the map to each incoming main image frame in the central processing unit, the adjustment now occurs automatically without any computational cost.

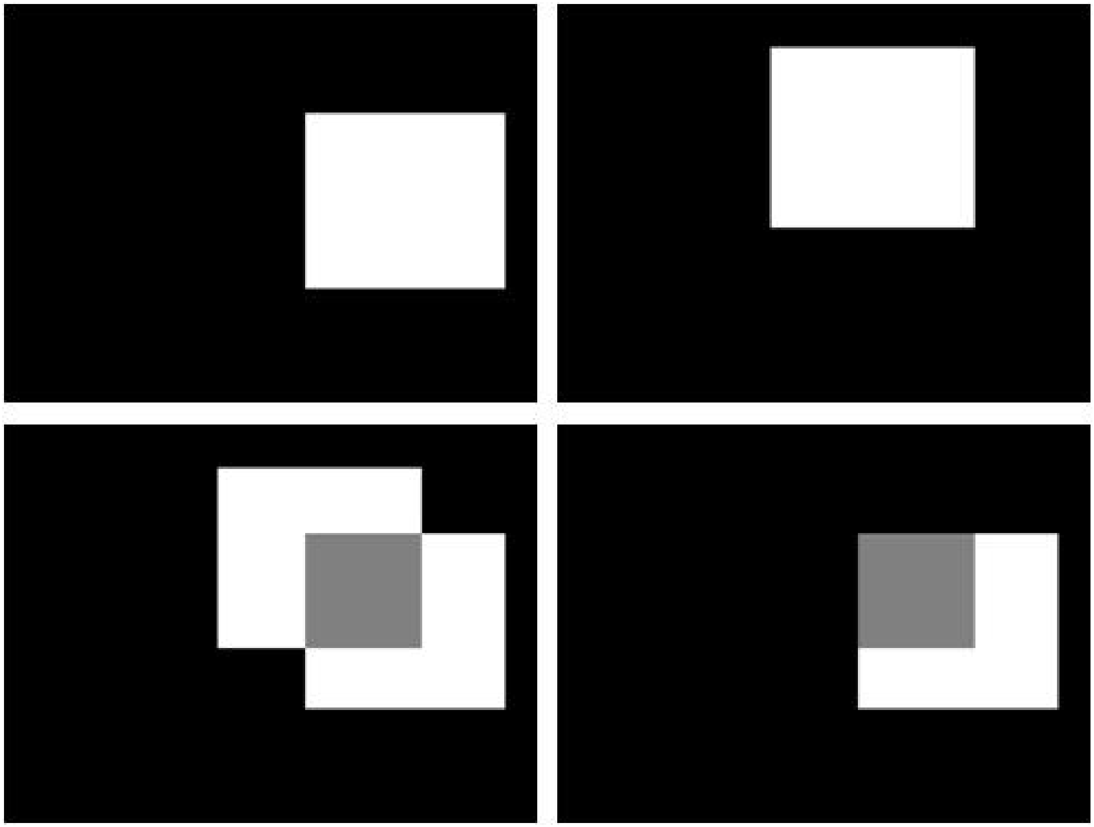

Figure 35: Schematic representation of responsibility map calculation. Top images show image element for two projector elements. Bottom left shows physical overlap on the main screen. Bottom right shows effective responsibility map of the first projector element (100% in white region, 50% in grey overlap region).

The final calibration step is to calculate warping matrices for each projector element. The LCD modulator in the projector element is generally a rectangular array of pixels while the corresponding image element on the main screen depends on the geometric relationship between the projector element and the main screen. For example, if the projector were positioned normal to the screen but pointing slightly downward then the image element would be a trapezoidal shape. In a real installation the random alignment of the projectors creates arbitrary quadrilaterals. In theory it would be possible to compensate for this warping of the image element by calculating the physical movement of the projector element away from the main screen normal in a three dimensional geometry and then compensating for each change. In reality such an operation is complex once the order of the movements is lost. Fortunately, knowledge of the geometry is irrelevant. The only goal is to ensure that the quadrilateral image element can be computationally warped onto the rectangular LCD. This can be achieved using simple warping matrices[22]. Once established, the warping matrix coefficients can be stored in the projector element.

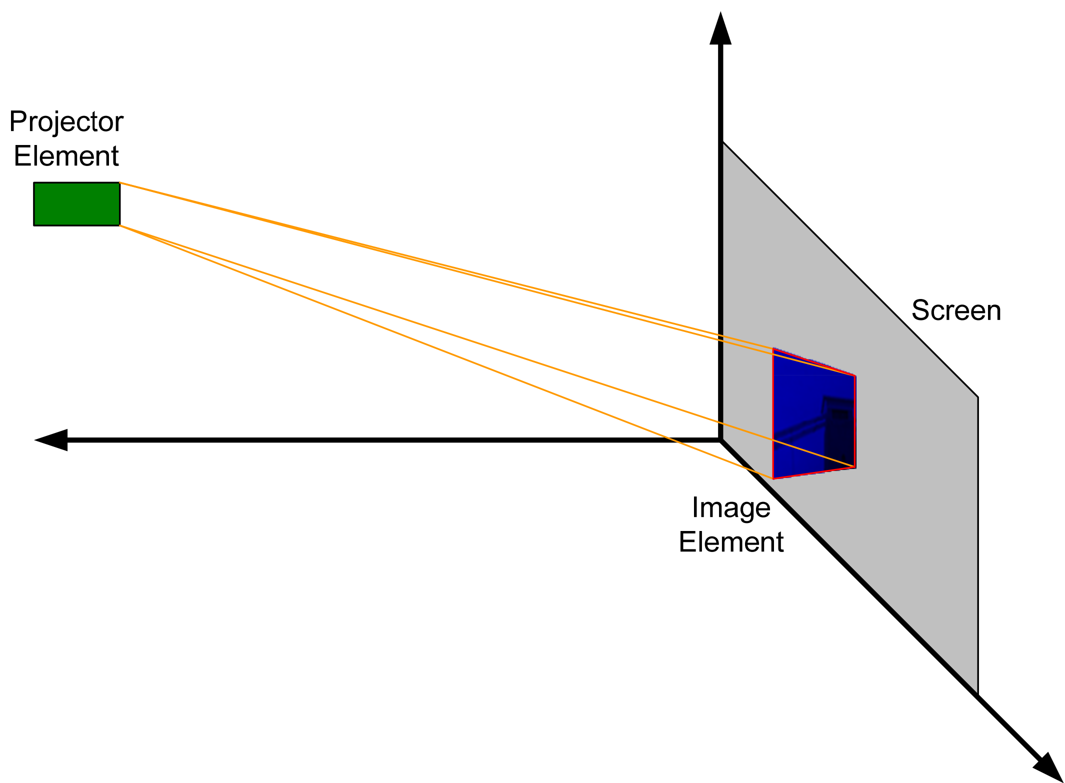

Figure 36: Warping geometry of a single projection element relative to main screen.

[22] A tutorial on warping matrices can be found at http://www.vision.caltech.edu/bouguetj/calib_doc/index.html

At this point the system is calibrated and the camera system can be removed. Any physical change to the system such as addition of further projector elements, failure of elements or changes to the geometry will require re-calibration. Fortunately, failure is unlikely to occur with solid state devices and the setup of projection venues tends to be fixed for long time periods.

7.2.3 Video Processing for Distributed Dual Modulation Projection Systems

Once calibration of the system is completed the main video processing steps are relatively simple. Similar to other dual modulation systems, the final image on the main screen is effectively the optical composite from different controllable elements. The task of the video processing algorithm is to separate the initial video frame into image elements using the algorithmic inverse of the optical re-combination. The first step in this process is to receive a standard video signal from a playback device. The resolution of the video signal is not important but higher resolution is desirable to leverage the high resolution capability of the distributed system. The central processing unit performs any require standard video processing operations (e.g. dithering, colour adjustments, low dynamic range to high dynamic range extension if needed, etc). Once those operations are completed the video frame should faithfully represent the desired main screen image since the remainder of the process is completely transparent in terms of image adjustment. The video frame is now transmitted to the projector elements. This transmission can take many forms. For the sake of illustration the simplest mechanisms is a massively serial transmission where each projector element receives the entire main video frame. A more complex and efficient distribution mechanism is described later on.

At each projector element the incoming video frame is multiplied by the responsibility map of the element. Since all areas for which the projector element is not responsible are black (zero value), only the relevant portion of the image for the particular projector element remains non-zero. That portion is scaled at each pixel according to the relative responsibility of the projector element. Finally, the image is warped according to the transformation matrix between projector modulator and the image element on the screen (the corresponding matrix and basis weights are established as part of the calibration process in the previous section). The result of these operations is a much smaller image with the spatial resolution of the projector element LCD. This image becomes the input for a classic dual modulation algorithm with a single LED (see section 6.2).

Figure 37: Pictorial representation of the separation process for a distributed projection system. Top left: initial main screen image. Top right: stylized responsibility map for a single projector element. Bottom left: section of the main image that is cut out by the responsibility map. Bottom right: LCD image on the low resolution LCD in the projector element. For visualization purposes the projector element is simulated without overlap and covering a larger than normal portion of the main image.

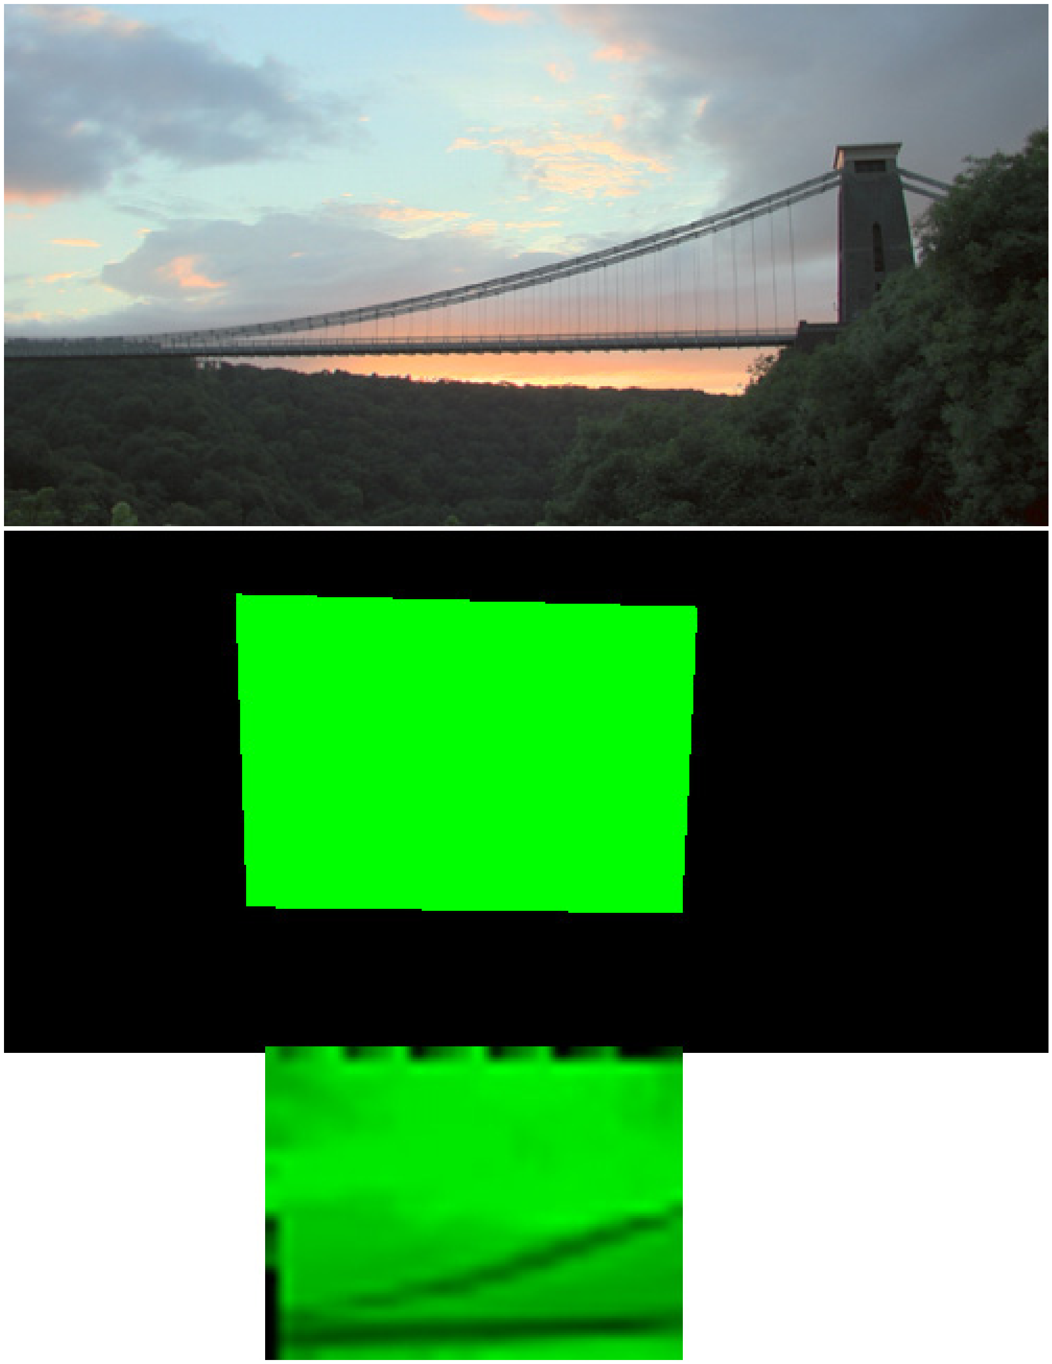

Figure 38: Simulation example of distributed projection process. Similar to Figure 37 the example is stylized with a larger than usual responsibility map and no overlap. Top: original main screen image. Middle: responsibility map. Bottom: LCD image. The effect of the warping matrix calculation can be soon on the top and left edges of the LCD image where black regions appear as a result of the rotation at low resolution.

The drive value for the projector LED is established by the average or adjusted average of the small image. The drive value of the LED provides a scale factor for the PSF of the element which in turn establishes the light field on the LCD modulator. Unlike other dual modulation system the projector element has only one light source and thus no overlap simulation is necessary. The incoming small image is divided by the light field to obtain the drive image for the LCD modulator. After appropriate adjustments by the inverse of the LCD response curve this image is sent to the LCD panel.

The computational complexity of this algorithm is extremely low due to the lack of overlap of several light sources found in any other dual modulation design. Each projector element processor only needs to execute one multiplication at main screen resolution (responsibility map x video frame) which can be further reduced by intelligent management of the large zero-area in the responsibility map. All other operations occur at the very low resolution of the LCD element and have a very low level of complexity. Of course every projector element needs to perform this operation independently so the system computation is comparable to a normal dual modulation system. The compute load has simply been spread across a massively parallelized grid.

Figure 37 shows the three relevant sub-images for a single projector element during video processing. For illustration purposes the main video frame is only split into a small number of projector elements. With such a small number of projectors each projector element is largely responsible for a region of the screen and thus the source image is mostly mapped directly on the projector modulator (in the appropriate colour channel). Once more projectors are added the modulator image becomes virtually unrecognisable until all image elements optically recombined on the screen.

7.2.4 Optimized Data Architecture for Distributed Dual Modulation Projection Systems

The naïve data architecture described above is unlikely to be efficient. It requires that the entire main screen image is transmitted at high resolution to every single projector element despite the fact that each element only requires a very small portion of the image. This might be possible in a smaller contained system such as single printed circuit board but is completely bandwidth-prohibitive in a larger array. A better solution is to use a staged separation of the video frame data. Figure 39 shows such an architecture where the main screen video frame is split into a number of sections which are then provided to the projection elements within those sections.

The assignment of projector elements to screen sections can occur during the manufacturing process. The easiest way to achieve this is by creating pre-fabricated projector arrays which are mounted onto the ceiling or other convenient location of the projection environment. If the projectors on the array are globally aligned (while maintaining local random overlap) then they can be assigned to a specific main screen section. Alternatively, the assignment to screen sections can occur during the calibration process by assigning each randomly placed projector element a section identifier based on the position of its image element.

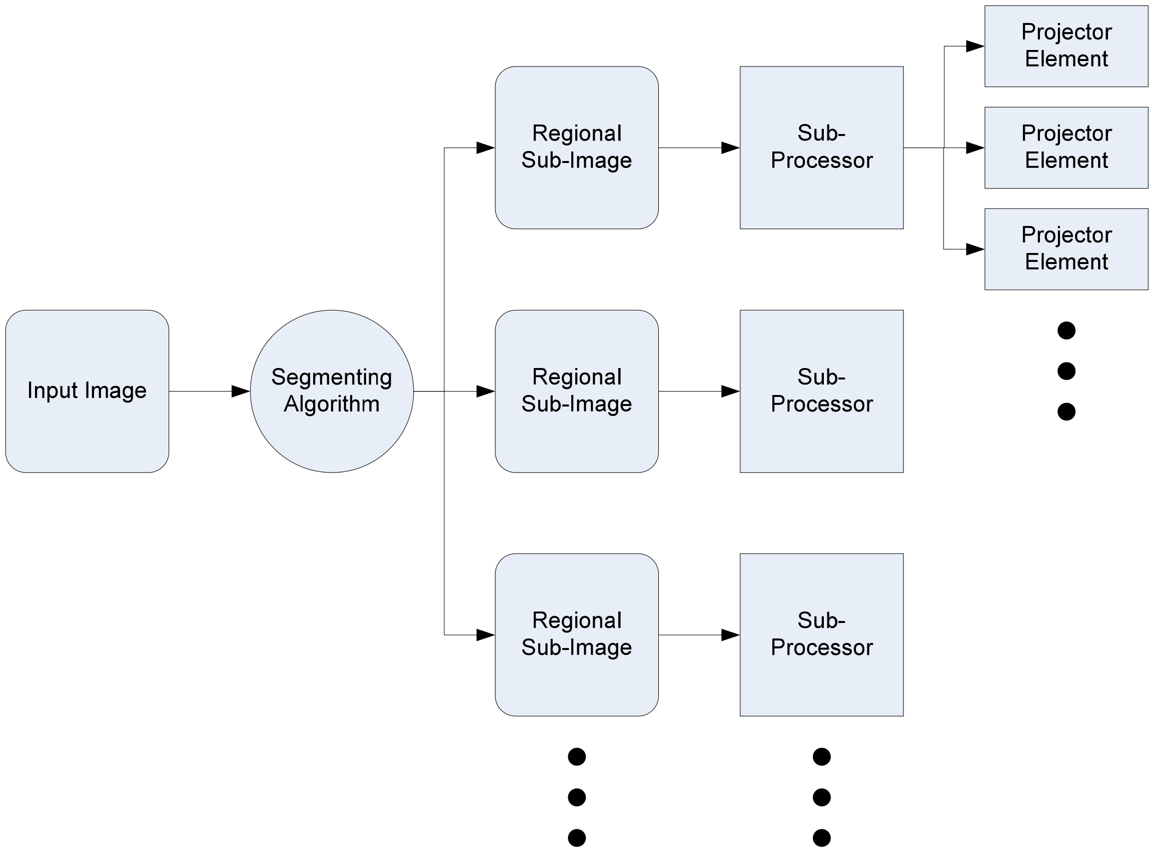

Figure 39: Advanced data flow for distributed projection system. The incoming image is segmented into several regional sub-images with appropriate boundary overlap. Each sub-image is sent to a sub-processor which performs the necessary image processing steps for a small group of projector elements.

Due to the random orientation and overlap of the image elements it is impossible to find clean boundary lines between sections. Instead, each section needs to carry an additional border of image information to ensure that projectors with image elements that are only partially within the section still receive all relevant information. The width of this extra information border is determined by the horizontal and vertical size of the largest image element in the section if each section border is scaled independently or the size of the largest element on the whole screen if it is preferable to define a single border size. Even though the inclusion of such a border means that the same information is sent redundantly to multiple sections, the parallelization effect ensures that the bandwidth required in each section is still very low.

This concept can be further refined by transmitting the segmented video frame not as raw data but in an encoded format such as MPEG. Cheap low resolution MPEG decoders are readily available and could be added to each projector element. Rather than decoding the video frame at the central processing unit and re-encoding sections for distribution it would then be possible to exploit the slice architecture of MPEG-4 and segment the incoming video signal without decoding. MPEG-4 slices can be separated without decoding the overall signal in a process that is widely used for partial modifications of MPEG transmissions (e.g. replacing the channel logo at the bottom of the screen in a life transmission of a sporting event). Such a technique can be used to segment the incoming video signal into many still encoded MPEG-4 slices which would then be transmitted to a set of projection elements. Decoding of the slice segments can occur at each projection element or in a sub-processing unit if all elements of a section are mounted closely together.

Using a sub-processing unit can also reduce the computational load for the element-level calculations. If the unit is mounted directly on the projector array then it can store all responsibility maps and calibration data for those projectors locally. The sub-processing unit can then perform the required computations for each projector element and send direct LED and LCD drive signals to the projectors.

7.2.5 Performance and Limitations of Distributed Dual Modulation Projection Systems

Many of the benefits of the distributed system derive directly from the large number of projector elements. While mass-produced elements would be comparably inexpensive or at least cheaper than a conventional cinema-grade projection system, low volume design of such units is completely cost-prohibitive. As such, the validation strategy for this system is based on a smaller scale physical system and a computational simulation of a larger array.

The physical system uses three commercially available pocket projectors[23] that have been modified to act as projection elements of a single primary only. Figure 40 shows the physical setup where each mini projector is mounted on an adjustable arm that allows movement along six degrees of freedom. The three mini projectors are randomly oriented and then manually adjusted such that they substantially overlap in a central region on the screen. This mimics a small segment of a full distributed system. Each projector is driven by a separate video output from a PC with enough graphics cards to support all three projectors. Computation of the projector images and drive signals is executed on the common PC to avoid the need for a micro-processor or field programmable gate implementation.

---

[23] Mitsubishi Pocket Projector PK20 SVGA DLP

Figure 40: Physical setup of distributed dual modulation projection system. The three Mitsubishi pocket projectors are mounted onto a common mechanical system. The calibration camera and screen complete the setup.

The system can be used on two modes. In the first mode each projector element uses a different single primary (i.e. red, green and blue) which allows an investigation of colour mixing and alignment aspects of the design. The second mode uses the full RGB mode of each projector element to simulate nine projector elements overlapping (of which three each are constraint to the same geometry). Figure 41 shows an image generated by the first mode. In the second mode the system achieves a measured dynamic range gain of 2.8 times that of a conventional mini-projector if the light sources are not directly controlled. With direct control of the light source the effective dynamic range of the system because entirely image content dependent and hard to evaluate with a three element unit.

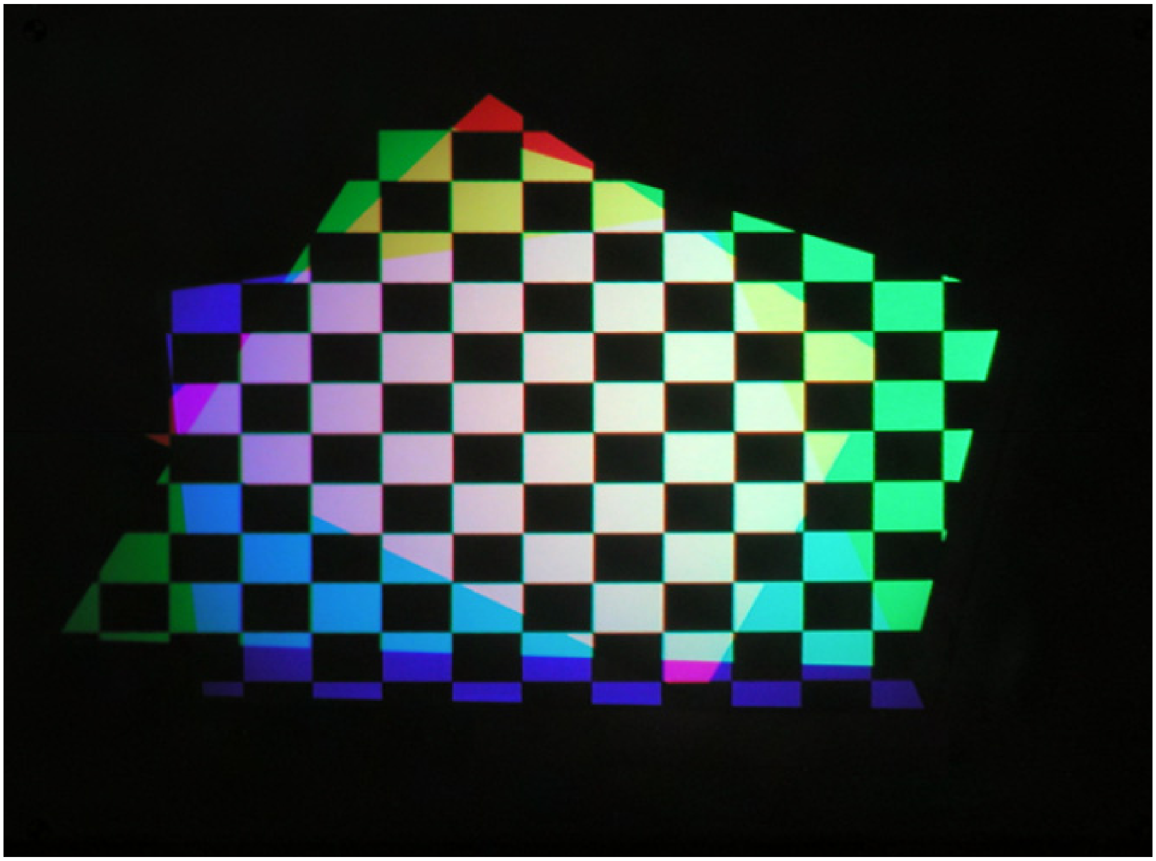

Figure 41: Photograph of distributed system prototype image. The three projector elements show a red, blue and green image element with overlap in the centre. Precise alignment can be achieved in the central region.

The software simulation provides a better tool to evaluate the performance characteristics of the distributed system. The simulation uses a 3D model of a projection environment with the bottom left corner of the projection screen as the origin. The model supports arbitrary screen dimensions as well as specification of a zone on the ceiling in which projector elements can be mounted. The projector elements are distributed inside that zone and oriented towards the screen. By default this arrangement is random but can also be specified to a regular grid layout. The projectors are randomly oriented towards the screen (a maximum rotation angle can be specified).

The overall image creation process is simulated using three independent software components:

- Physical Setup Simulation
- Forward Process Simulation
- Backward Process Simulation

The Physical Setup Simulation establishes all parameters defining a multi-element projector system. This includes generating a physical layout of all the projector elements and creating calibration data such as individual efficiency of each projector. All calibration data is stored in a format comparable to that found in the actual device (i.e. compartmentalized for each projector element).

The Forward Process Simulation executes all steps necessary to go from an incoming video signal to the final LED and LCD signal of each projector element according the to algorithm described in section 6.2.1. Apart from the execution on a PC platform, the algorithm is implemented identically to the way a real system would use it.

The Backward Process Simulation simulates the response of the projector elements according to the signals provided by the Forward Process Simulation. As the name implies, this simulation reverses the process of the forward simulation to obtain the final image output on screen via optical projection. The simulation includes appropriate random variation steps to simulate calibration errors, LED lifetime decay and other factors that influence the output of the actual system.

Once the physical setup and forward simulation is completed, the software forecasts the appearance of the final image on the screen by reversing the projection process. Each projector element projects its image element onto the main screen according to the orientation of the projector. On the screen the image element joins all the other images from other projector elements to create the final main screen image. As long as each projector element generates the appropriate image element (as established by the responsibility map and the source image in the forward simulation) the sum of all images elements adds up to the original source image (though at higher resolution on the screen due to the beneficial resolution increase described earlier).

The backward simulation allows also simulates specify rotational and efficiency error margins to reflect real-world conditions. The physical location of each projector element is assumed to be very stable but its rotational alignment could vary as a result of loose mounting screws or ground vibrations. With these variations taken into account the backward simulation returns the final on-screen image. A user study with 15 participants was conducted to investigate the effect of these alignment errors on image quality. The study used a simulated system with 1,000 projector elements. Participants were 2/3 male in the age range of 19 to 35 and all had normal or corrected to normal vision. The simulated images were shown on a 47” 1920x1080 advanced dual modulation display at a viewing distance of 1m. The results show that variations in projector intensity above 5% cause visible artefacts in either colour separation or luminance uniformity. Physical alignment changes of more than 3% of the initial side of the image element likewise cause visible artefacts. This represents a very substantial safety margin as mechanical tolerance and stability levels are likely to be significantly higher than 3%.

Figure 42: Distributed dual modulation projector system simulation. Top: result with low physical variation (<3%). Bottom: high variation for visualization of artefacts in print (>20%).

Overall, the distributed dual modulation projection concepts can achieve all the requirements of HDR viewing in an energy efficient manner. The concept also introduces an opportunity for better implementation characteristics such as better thermal performance, easier maintenance, upgradeability and an easy path to wider colour gamut designs. These benefits are also the root of the sole disadvantage of the concept which is its introduction of a paradigm shift for projection installations. Instead of a single source of modulated light, the new projection installation hosts a large number of smaller image sources. In the long term this feature is attractive for the reasons given above but in the short term it might pose implementation problems.

7.3 Active Screen Dual Modulation Projection Systems

Previous chapters have applied the dual modulation concept to a variety of display devices. So far all variants had in common that the two modulators were placed inside the device make it a compact unit independent of the surrounding environment. This is essential for front view display devices and projection devices in mobile environments (e.g. business projectors). The primary disadvantage of this approach is the need for modification of the core device. In particular, the distributed dual modulation projection concept requires the installation of a completely novel type of projection system and an overhaul of the entire cinema projection paradigm of a single central projection unit. Ultimately the benefits of the distributed design such as thermal benefits, solid-state installation and the HDR viewing experience might overcome this major barrier but in the near term an alternative path to HDR viewing would be welcome. A backwards compatible HDR projection system could achieve this goal by delivering the full HDR experience without changing the core projection system.

To achieve backwards compatibility with current projectors it is necessary to place the second modulator outside of the projection device. A simple solution would be to retrofit a secondary modulator directly onto the main projector lens and effectively create the same optical system as the basic dual modulation projector described in section 7.1. Unfortunately, this approach also carries the energy efficiency and thermal disadvantages of the basic system and is therefore unlikely to be suitable.

An alternative design is to separate the two modulators physically by making the screen itself one of the modulators. A normal projection screen is a passive component that reflects light without any modulation but it can be replaced with an actively modulating layer. The advanced dual modulator concept introduced the option of reducing the spatial resolution of one of the modulators and compensating for some degree of blur of the secondary modulator. This makes it possible to use a low resolution reflective display instead of a normal un-modulated screen. Such screens could either be custom-made display systems or could leverage existing reflective display technology such as E Ink's electronic paper. The normally high cost of these components would be significantly reduced by the extremely low resolution of the overall screen.

7.3.1 System Architecture for Active Screen Dual Modulation Projection Systems

The active screen dual modulation projection system uses a conventional high output projector as an integrated light source and primary modulator. Added to the projector is a modulated viewing screen which acts as a secondary modulator at low spatial resolution. Similar to dual modulation concepts described earlier, the primary modulator creates an image which captures the high resolution colour features of the input image and includes a correction pattern to compensate for the low resolution blur of the secondary modulator.

Figure 43 shows this basic architecture including the projector and the active screen. The active screen is a reflective surface whose reflectance can be smoothly adjusted over the entire screen. Good candidates for such a screen are emerging "electronic paper" display technologies. These displays use a reflective layer that is natively in one reflective state (usually black or white) and can switch to the opposite state upon application of an electric potential. A good example is E Ink's electronic paper material. E Ink fabricates a film which contains an array of micro-capsules[24]. The sealed capsules are filled with black and white particles that carry opposite electric charges. In a normal display application such as an electronic reader the film is sandwiched between a transparent layer of conductive coating and a segmented layer of individually addressable pixels. The pixelated layer is usually a conventional thin film transistor layer on glass. For image creation the thin film transistor is energized appropriately at each pixel and the resulting electric potential difference between the pixel and the common electrode

[24] E Ink VizPlex™ film.

creates an electric field in the capsule directly above the pixel. Particles in the capsule now rise to the surface of the capsule to form one of the two image states (black or white for the corresponding particles depending on the polarity of the particle charges and the thin film transistor arrangement). The opposite image state can be achieved by removing the potential difference and allowing natural diffusion to reduce the excess number of particles of one type on the surface of the capsule[25]. Other electronic paper materials use different mechanisms for image creation by retaining the same principle of a reflective layer with adjustable reflectance resulting from the adjustment of an electric potential per pixel.

Basic electronic paper film is available from several sources in large sheets and could be fabricated seamlessly for large scale cinema applications where the distance from the screen to viewer is very large. The scale challenge for electronic paper arises from the need of a high resolution thin film transistor layers. Those cannot be manufactured in larger sizes without significant cost. Nor can they be tiled effectively because at least one side of the drive layer needs space for other non-pixel control components which would create a visible seam.

While this is a barrier for the use of electronic paper in large scale direct display, it isn't a factor for the active screen dual modulation concept. The primary modulator in the projector delivers the required high resolution so the active screen can be of very low resolution. As a result, the screen drive layer can be either a very low resolution thin film transistor or more likely a simpler drive architecture altogether. As seen in previous chapters, the pixel count of the low resolution modulator can be in the order of hundreds

---

[25] See www.eink.com for more details on E Ink electronic paper and application examples.

or low thousands. For a large cinema screen this could be achieved even with direct wiring leading to large metal electrodes.

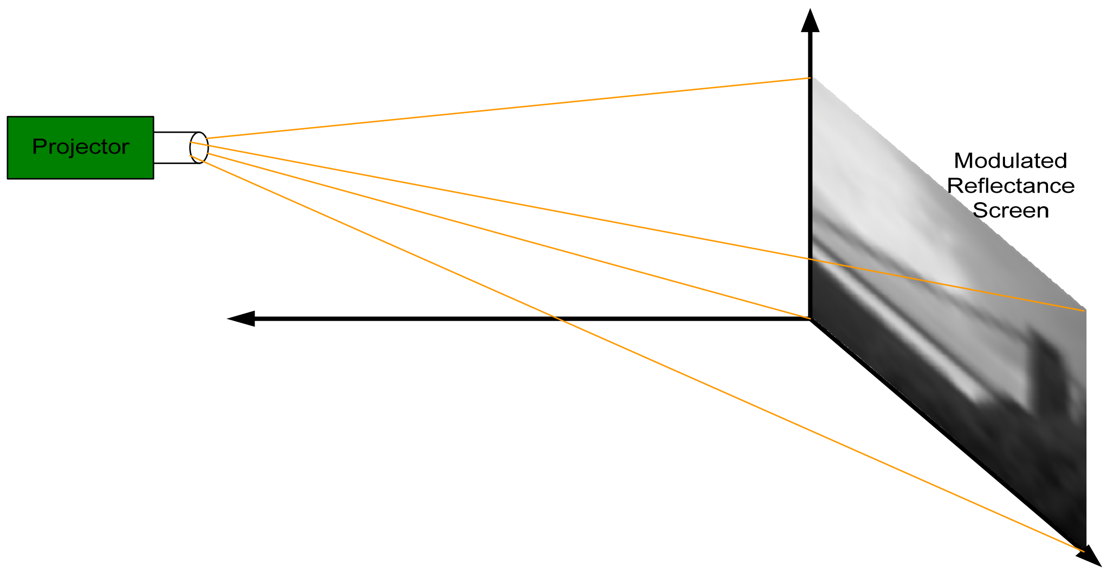

Figure 43: Active screen dual modulation architecture. Conventional projector images high resolution data onto a low resolution screen with modulated reflectance[26].

Apart from low resolution, the screen must feature a smooth transition from one low resolution element to the next with some overlap. For all previously described dual modulation systems this was easily achieved by inserting a diffuser into the optical path or defocusing the optical coupling between first and second modulator. Neither option is available for the active screen concept because the low resolution modulator is not placed physically after the high resolution modulator. A diffuser placed on the active screen

---

[26] For a smaller implementation of this concept see (Bimber & Iwai, 2008).

would blur the high resolution image as well as the low resolution pattern and thus destroy the overall image.

A better alternative is to use non-optical means to smoothly change the reflectivity pattern on the screen. This can be achieved by "blurring" the electric potential of the screen drive layer rather than the front facing optics. The naïve path to a smoothly varying electric potential would be to increase the drive layer resolution again and then use the higher resolution pixel to spatially dither to the blurred lower resolution. Unfortunately, that solution is neither practical nor in line with the expected benefits of the active screen concept. Instead, the solution is to use material properties to create the electric potential "blur".

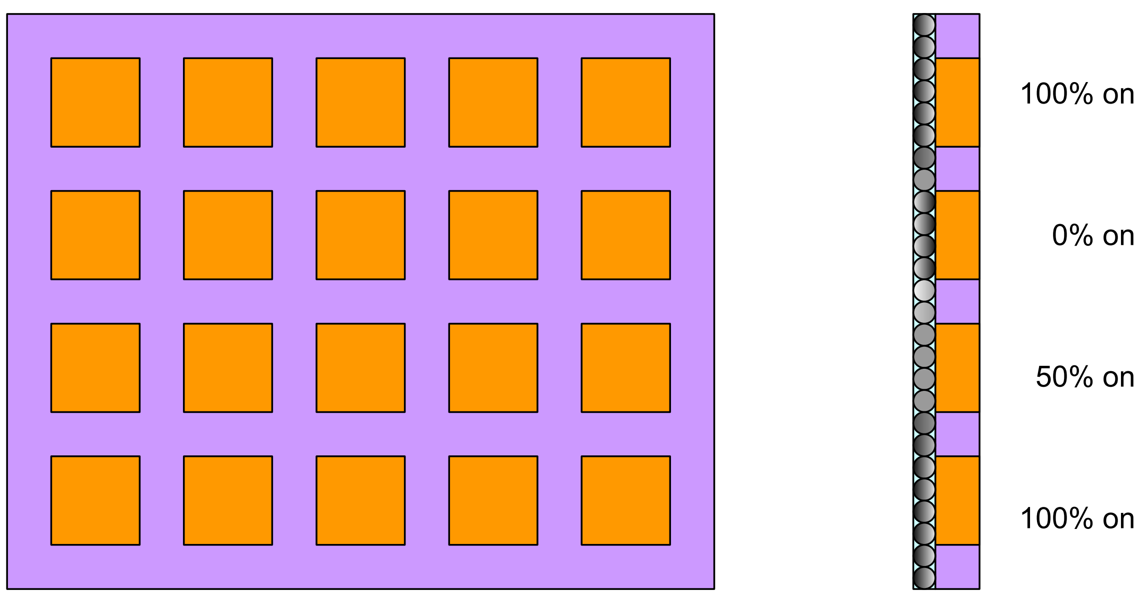

Figure 44: Active screen drive plane. Electrodes shown in orange, dielectric filler material in mauve. Cross-section on the left shows drive plane behind VizPlex™ layer with drive levels and corresponding capsule states indicated.

In a conventional drive layer the individual pixel electrodes are placed as close as possible together to avoid visible boundaries between the pixels. For this application the boundaries are desirable but need to be filled with an electric potential that varies smoothly from one pixel electrode to the next. This can be achieved by placing smaller conductive electrodes into a weak dielectric layer as shown in Figure 44. The dielectric material causes the electric potential to gradually reduce rather than cut off sharply at the pixel boundary. With proper choice of dielectric material and electrodes spacing, this technique can create a smooth electric potential between the drive layer and the common electrode on the other side of the electronic paper film. This varying potential creates a varying reflectance pattern on the screen as each capsule in the electronic paper film is exposed to a different potential.

Using this technique it is possible to create the required low resolution blur without reducing the sharpness of the primary image. The number of drive electrodes also remains very low. As a result the video processing steps for the active screen system are identical to those of the basic dual modulation projector in section 6.1. The only difference is the unique PSF of the active screen as a result of the slight differences in electrical versus optical blur. This requires only a parameter change in the algorithm.

As with most dual modulation concepts, the specification of the projector defines most of the overall system. In particular, the spatial resolution, refresh rate and colour gamut of the overall system are identical to that of the projector as long as the active screen has a reasonable temporal response and spectral reflectance profile. The low resolution modulated screen only affects the dynamic range related system specifications such as contrast (the product of the two modulator ranges) and brightness.

7.3.2 Performance and Limitations of Active Screen Dual Modulation Projection Systems

The primary benefits of the active screen concept compared to the basic dual modulation projector are thermal management and backward compatibility. While the second modulator absorbs 50% of the light from the source in both cases, the loss occurs at a different point. A basic dual modulation projector absorbs the light within the projector chassis where the thermal load is already very high due to the presence of the light source and primary modulator. The active screen concept on the other hand shifts the secondary absorption to the screen itself. Not only is the screen physically separate from the projector, it is also extremely large so that no cooling solution is required.

The backward compatibility advantage is compelling because no modification to the main projector is required. The installation of a new active projection screen is still required and it is a question of economic choice whether replacement of the screen or projection system is a better option. The image quality of the active screen projection system is comparable to that of a basic dual modulation projector with similar primary modulator. The contrast of the E-Ink or alternative active screen material is slightly lower than that of a low resolution LCD modulator but can be increased by increasing the voltage on the screen drive pixel. This is impractical in a normal thin film transistor design but very easy with the small number of very large metal electrodes used in this application. All other characteristics are comparable including the expected light loss of approximately 50% for a conventional solution and possible as little as 20% with advanced electronic paper films such as retro-reflective film solution (Mossman, Kotlicki, Whitehead, Biernath, & Rao, 2001).

Figure 45: IRex Iliad E-Reader using E Ink electronic paper display

Availability of larger sheets of E Ink or comparable material is very limited outside manufacturing environments so an IRex Iliad E-Reader[27] was used to mimic a smaller version of the active screen (see Figure 45). The Iliad uses E Ink electronic paper film couple to a thin film transistor drive layer. For the prototype setup the normal E-Reader operating system of the Iliad was modified to allow playback of full screen image frames. With that adjustment the Iliad can produce a blurred low resolution reflectance pattern. The pattern was computed using a simulation version of the dual modulation

[27] See www.irex.com for more details on the Iliad E-Reader.

algorithm with a 10x20 simulated screen resolution followed by a blur function similar in shape to the expected electrical field spread in the actual implementation. The simulated dielectric spacer layer was half the width of the drive electrodes. Figure 46 shows an Iliad unit with such a modulated low resolution reflectance pattern. A conventional LCD projector was placed in front of the Iliad and focused on the 9.1" electronic paper surface. Both image modulators were connected to the same PC (with appropriate override of the Iliad OS) to synchronise the two image streams. The measured contrast of the combined system was greater than 20,000:1 with the projector contrast being 2,000:1 and the Iliad delivering 15:1 independently.

The active screen concept also offers an opportunity to finally explore the full HDR viewing experience on a large screen. Figure 47 shows a static large scale prototype of the active screen concept. Static images were generated for the expected reflectance pattern on a large electronic paper screen and then printed onto 240x180cm front projection screen sheets (1.7 gain material with a 90 degree viewing angle). The printing process was calibrated to ensure that the reflectance pattern is comparable to a real electronic paper screen with the same image. The simulated resolution for the screen was 32x24 with the dielectric spacer width being half as wide as the drive electrode. Multiple screen images were generated in this fashion and placed on an interchangeable mounting system. The screen was illuminated by a Christie LX-1500 digital cinema projector with a 12,000 Lumen peak output. With the simulated peak reflectance of basic E Ink electronic paper at 40%, this resulted in a final image peak luminance of 1,050cd/m^2 and a dark level of approximately 0.06cd/m^2. The native contrast ratio of the

Christie projector is 1450:1 (spec of 2000:1) for the gamma and colour setting used. The native contrast of the simulated electronic paper is 15:1.

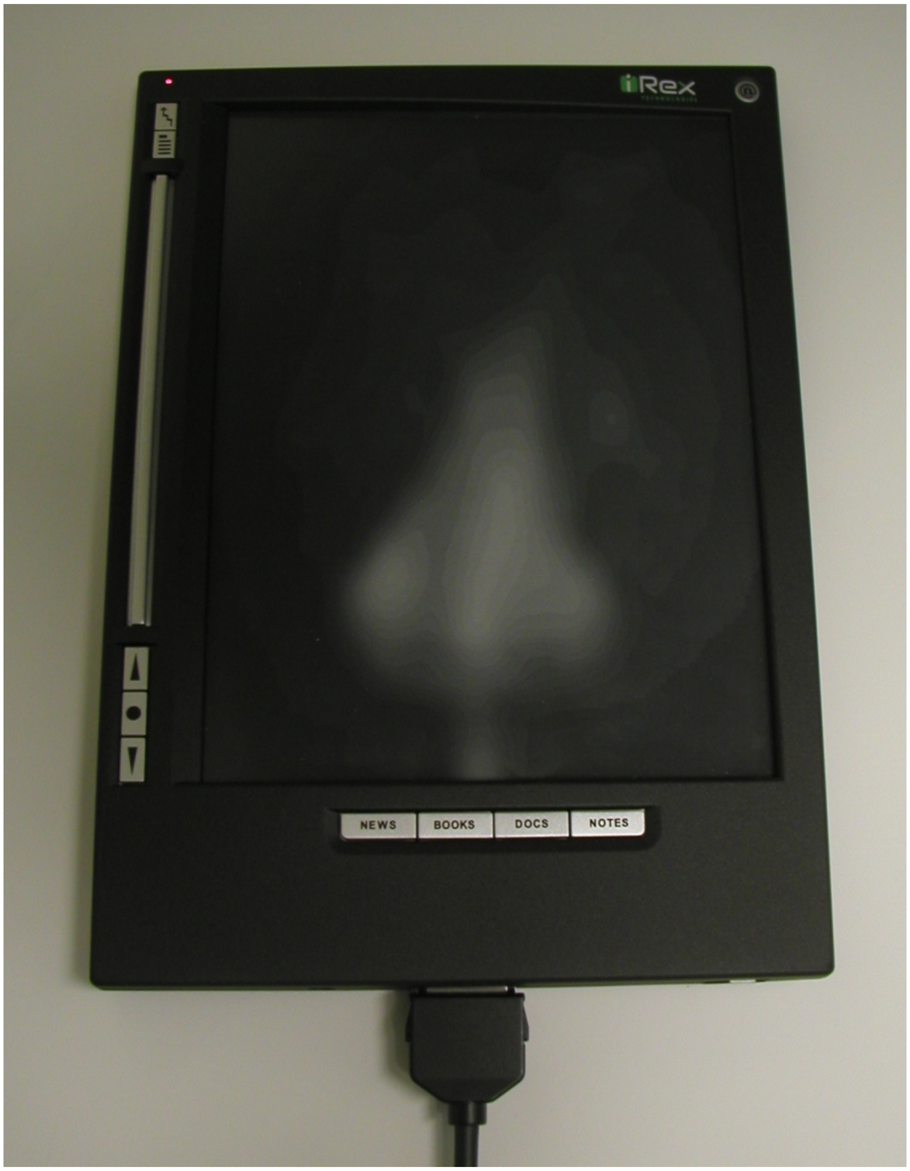

Figure 46: Iliad unit showing a low resolution reflectance pattern of a lit match-stick image. Notice that the background remains black while the matchstick features higher reflectance.

Figure 47: Photograph of static large active screen prototype (240cm x 180cm). The photographic capture is saturated on the top and bottom of the range (i.e. clipping in the windows and deep shadows). Those areas are filled with detail for the real world observer of the system.

Figure 48 shows five exposures of an image on the large active screen demonstrator (1 f-stop each). The active screen dual modulation system delivers true HDR performance in a backward compatible way without modification of the main imaging unit. This makes it a compelling solution for installations that show both low and HDR images in the same setting.

Figure 48: Five exposures of the large active screen system. Exposure varies by 1 f-stop per image.

8 Dual Modulation System Benefits

Regardless of the choice of basic or advanced implementation described in the past chapters, a well implemented dual modulation system offers several benefits over conventional solutions. First, the use of the second modulator greatly improves contrast, luminance and colour gamut. Second, several system specifications such as amplitude resolution and luminance uniformity are also enhanced. In the case of advanced dual modulation, further benefits can be obtained in the areas of energy efficiency and system lifetime.

Independent of these benefits, dual modulation systems also leverage the capabilities of the main modulator. For example, the spatial resolution of the system is identical to the resolution of the main modulator. Any improvement made to that modulator (e.g. a choice of a higher resolution LCD) directly improves the capabilities of the overall system. Unlike display solutions that enforce a choice between a new feature and the quality of existing features (e.g. 3D effect vs. spatial resolution for 3D displays), the dual modulation system combines high dynamic range with the benefits of conventional technology.

8.1 Contrast, Luminance and Colour

Dual modulation systems offer several perceptual advantages that are immediately obvious to a first-time viewer. Specifically, dual modulation yields better contrast, luminance and colour quality. The contrast of a dual modulation system is given by the product of the capabilities of the individual modulators. Unlike most contrast improvement techniques this is a multiplicative effect so that very high contrast ratios are achievable. For example, combining two modest 100:1 modulators yields an impressive 10,000:1 effective contrast for the combined system. Using a light source array as the first modulator in the advanced configuration improves the contrast ratio even further. In fact, the advanced designs practically eliminate contrast as a meaningful measure when the light sources are switched off in dark areas. Obviously, this high contrast enables much higher luminance ranges because the black level is effectively anchored at zero and independent of the peak luminance of the display.

The improvement of colour quality has two sources. First, the improved contrast and luminance of the dual modulation system creates a perceptual impression of higher colour saturation. For example, an image of a red rose shown at 1,000cd/m^2 looks significantly more saturated than the same image at 100cd/m^2 even if the actual chromaticity of the rose remains unchanged[28]. The second improvement of colour comes from the ability to use multi-primary light sources as the first modulator in the advanced variants (e.g. RGB LEDs). Such light sources are narrow spectral band emitters and produce very saturated display primaries. Figure 49 shows an example of the same LCD

[28] This effect of perceived increase of saturation at higher luminance is called the "Hunt Effect". More details can be found in (Hunt, 2005).

being backlit by a conventional white LED and a triad of RGB LEDs. The increase in saturation for the three primaries yields a significant overall colour gamut expansion.

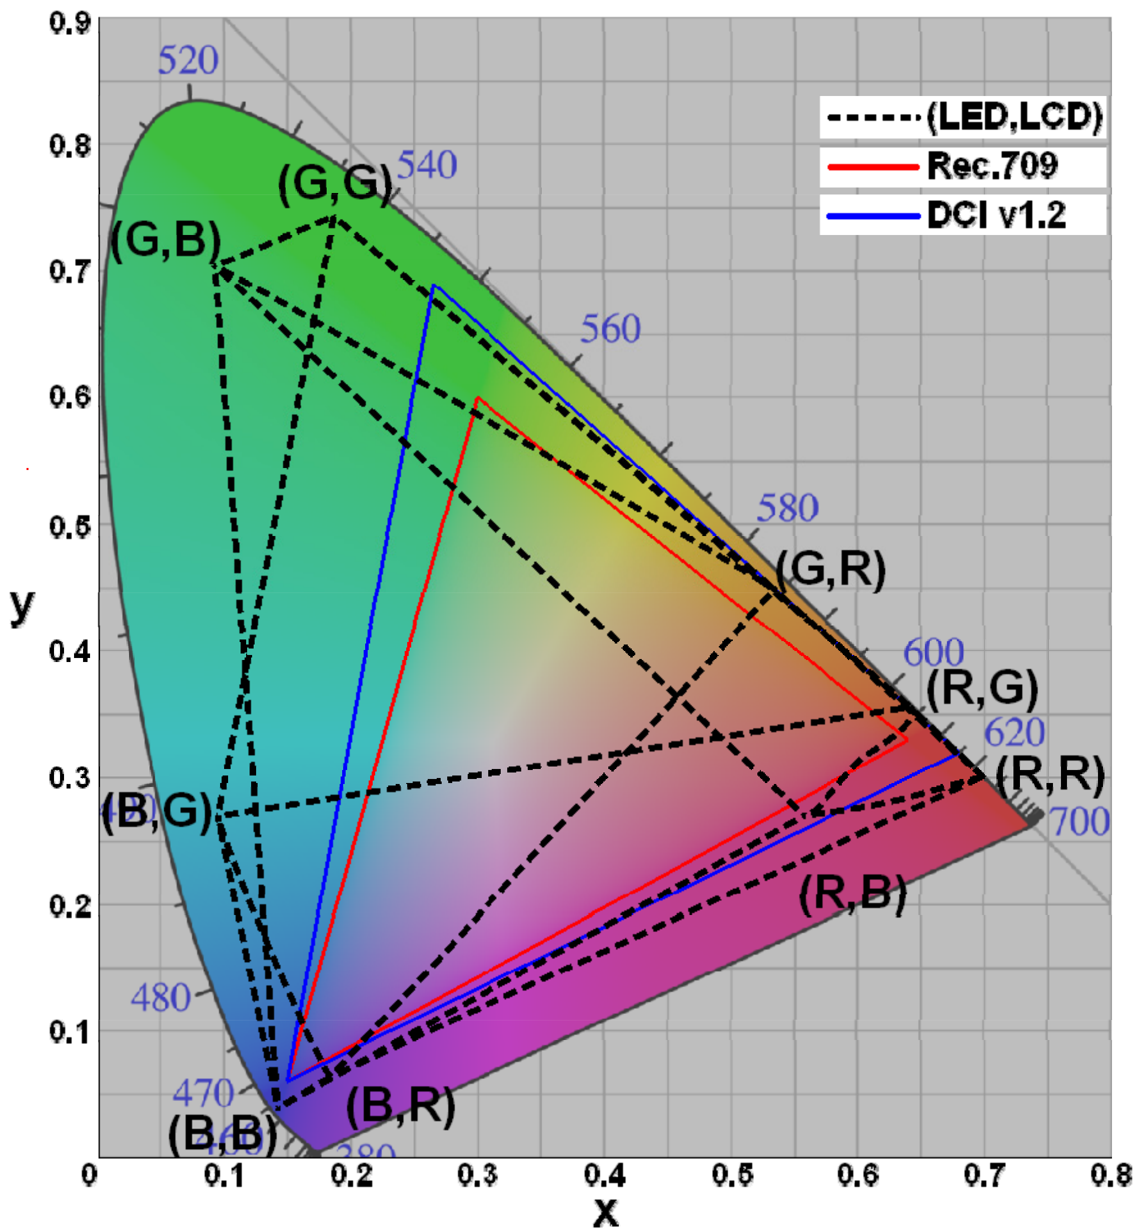

Figure 49: Effective multi-primary gamut of an advanced dual modulation display using RGB LED and RGB LCD. All combinations of the two modulators are indicated.

8.2 Amplitude Resolution

Dual-modulation has a strong impact on the amplitude resolution of a display. Similar to contrast, the amplitude resolution of a dual modulation system is given approximately by the product of the individual ranges. For any amplitude step of the first modulator there are many choices for the second modulator and even though some combinations are not unique, there are many more distinct amplitude steps than in either modulator alone. For example, if both modulators use a linear 8-bit range then the complete display has a 16-bit range with 65025 combinations of the two modulators. Out of those approximately 17,000 steps would be unique. These unique steps would be distributed non-linearly so that the relative step size is below 1% over the entire range. Figure 50 shows a comparison of a true 16-bit system and a dual modulation system with two linear 8-bit modulators to illustrate this effect. Notice that the relative step size of the dual modulation system remains below 1% even at the high end of the range.

In actual implementations, the amplitude resolution of dual modulation system is harder to characterise. The response curve of most modulators is non-linear and the number of unique combinations is therefore more challenging to estimate. An even bigger challenge comes from the fact that the secondary modulator has a much lower spatial resolution than the main modulator. This creates a spatial coupling between individual pixels and makes it impossible to accurately determine the effective amplitude resolution of the overall system. The blurred PSF of the secondary modulator is an analog distribution of light with continuous smooth transitions between main modulator pixels. Even if the drive system for the secondary modulator is digital, the effective light field generated by it is analog in nature. The dual modulation system is therefore actually a combination of a smoothly varying analog light field and a digital stepped main modulator.

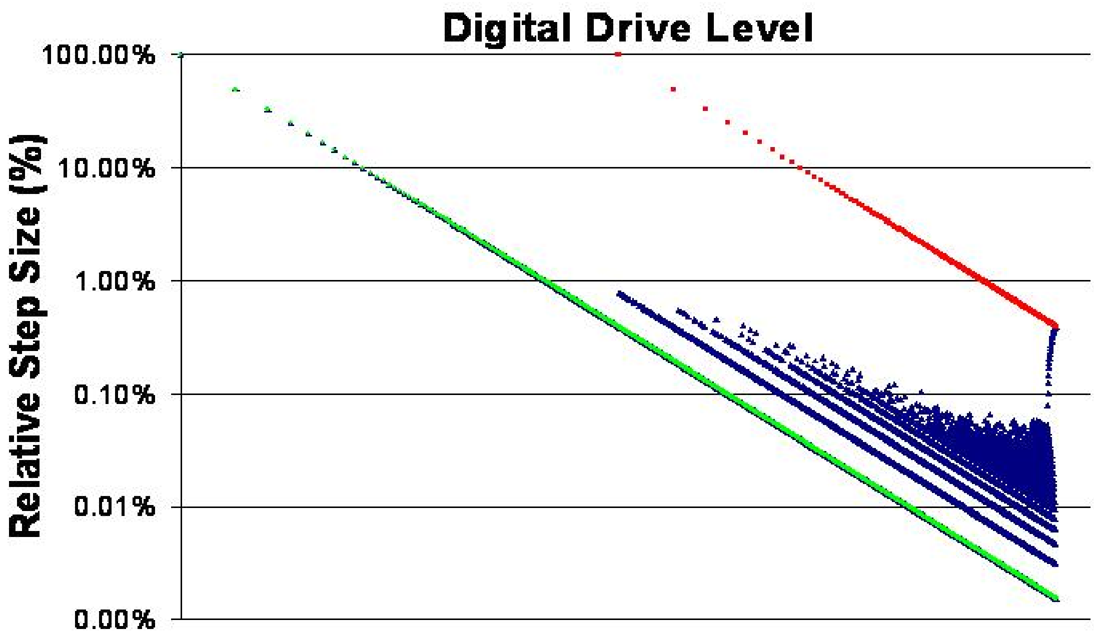

Figure 50: Comparison of amplitude resolution for conventional and dual modulation steps. Graphs show the relative step size over the range of a true 16-bit display (green), a dual modulation display using two 8-bit modulators (blue) and an 8-bit display (red).

Fortunately, this odd combination of analog and digital is very consistent with the limitations of our HVS. Veiling luminance and other limiting effects prevent our eye from resolving high contrast in a local region. However constant movement of the eye across the images provides the ability for rapid local adaptation to different areas. This means that a low spatial frequency high dynamic range luminance modulation with high spatial frequency lower contrast detail is very pleasing to the eye.

8.3 Energy Saving

Previous sections have focused on image quality benefits of dual modulation display systems. Not to be overlooked is the substantial energy efficiency gain that such devices can deliver. Energy consumption is a key differentiating feature in a world where the ecological impact of consumer electronic devices has become a major consideration.

Emerging global dimming displays with the ability to modulate the entire backlight consume less power. Studies have shown that global backlight dimming can typically, depending on image content; reduce backlight power consumption by up to 50% (Chen, Sung, Ha, & Park, 2007). Overall, global dimming technology can provide substantial power savings for scenes that are uniformly dark but struggles to maintain this benefit when a small portion of the image is bright.

Advanced dual modulation displays or projectors use a variable intensity light source array to create the final image. This means that, unlike a conventional static backlight or global dimming display, the dual modulation system only generates light in proportion to the required image luminance. The energy consumption of the device is therefore proportional to the intensity of the image content. Most content has an average luminance level that is much lower than the peak level so that the average power consumption of advanced dual modulation systems is very low. Figure 51 shows the per-frame power consumption of the backlight of a 37” advanced dual modulation display for a compilation of industry test images and representative motion sequences (e.g. movie trailer, TV sequences, etc). The average power consumption of this device for this sequence is only 27% of that of a constant backlight design. Since the backlight of a display consumes most of the system power, this leads to a substantial improvement of energy efficiency of the overall system.

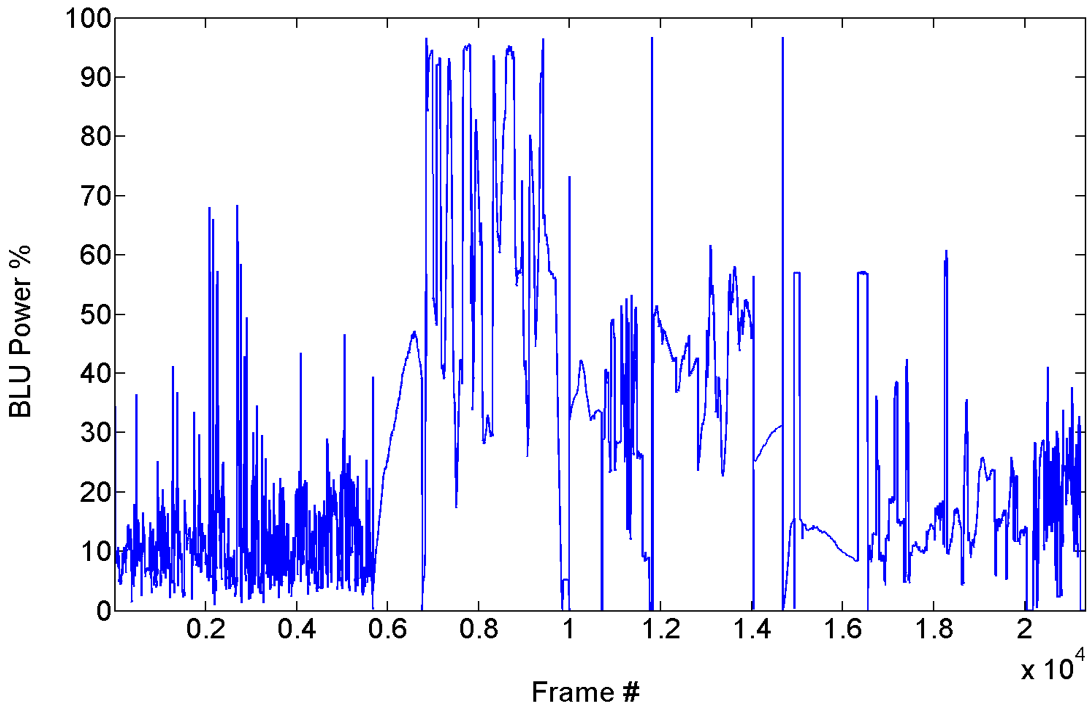

Figure 51: Backlight power level over 21,000 representative image frames for a 37" advanced dual modulation display with 1380 LED[29].

This energy saving is highly content dependent. Figure 52 presents histograms of average power levels of 1,000 representative frames for different types of image content. NTSC TV content shows an average power level in the range of 25-30% (see also Figure 51 for a longer sequence of TV frames averaging at 27%). Representative computer gaming content is generally at even lower intensity levels (~20%) due to the darker nature of gaming imagery. The highest average power consumption at 40-50% can be found in desktop applications due to the predominance of white backgrounds (e.g. text

[29] Data collected by Dr. Peter Longhurst of Dolby Canada Corp., used with permission.

editors). Even at 50% average backlight level the overall power saving is still significant compared to conventional displays. Power saving of this kind is inherently related to the peak luminance of the display. In general, at higher system luminance the ratio of relative mean image intensity to peak luminance is lower and thus the relative power saving higher. The same characteristic that enables the power saving also improves the lifetime and uniformity of dual modulation systems. Lower average power levels imply a lower thermal load for the light sources and consequently both longer lifetime and higher luminous efficiency.

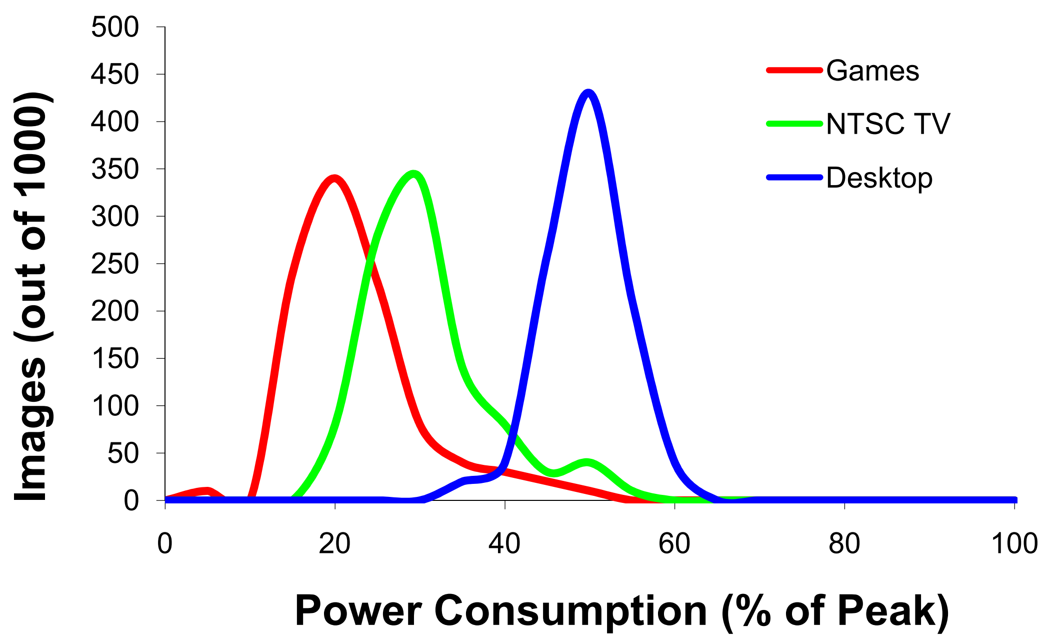

Figure 52: Histogram of advanced dual modulation display power consumption. Data represents 1,000 representative Video Game, NTSC TV and Windows Desktop images. The display has 1380 LED behind a 37” LCD and offers a peak luminance of 3,000cd/m^2.

8.4 Viewer Experience

Of course the ultimate benefit of a true HDR system is the exceptional viewing experience. All other technical specifications and benefits described so far combine to create a high quality image that is only rivalled in appearance by the scene that the image represents. It is hard to quantify or describe the effect of seeing a digital image that is comparable to the real world. Decades of viewing conventional displays have conditioned us to accept that digital reproductions are only pale abstractions of the real experience. We have effectively trained ourselves to live with the limitations of displays. In this environment the first encounter of a true HDR display prompts many to regard it as a window to a real scene rather than a digital display.

This paradigm shift in viewing experience makes it challenging to evaluate HDR displays against conventional displays using user studies. Preference studies between conventional and HDR displays lead to very binary results. A study conducted with 12 participants (9 male, 3 female, aged 21 to 39 with normal or corrected to normal vision) comparing pairs of 10 low and high dynamic range images resulted in all participants strongly preferring the HDR version for all images.

Chapter 4 provides the results of user studies where image processing is used to gradually scale the increase in dynamic range. This allows a single HDR display to map out the range between conventional and HDR displays. The results of this study show a strong preference for higher luminance and higher contrast images.

Undesirable HDR images can of course be created by inappropriate image adjustments. All comparisons are heavily influenced by the choice of image processing for both the conventional and HDR display. Image processing is necessary because conventional displays cannot show HDR content. The two choices therefore are to enforce either high or low dynamic range content for both devices. Allowing each device its own type of content stacks the deck unfairly in favour of HDR displays since HDR content has more image information in the first place.

If the content is high dynamic range then the conventional display requires a choice of tone mapping algorithm to create an acceptable input image. Ledda et al. compared six common tone mapping operators against each other and against an unmapped HDR display as a reference (Ledda, Chalmers, Troscianko, & Seetzen, August 2005). In addition to showing the advantages of different operators compared to others, the study indirectly shows the preferences for high dynamic range. None of the operators reaches the reference quality rating of the HDR display in any of the visual quality tests.

If the content is low dynamic range then the HDR display needs to boost the content range algorithmically. In essence this is a reverse tone mapping algorithm which has a strong influence on image quality. The colour appearance of HDR displays is often different than in conventional displays (Oğuz Akyüz & Reinhard, 2006) and thus care needs to be taken to preserve a comparable look of the image for direct quality comparisons. Rempel et al. (August 2007) propose a sophisticated algorithm to increase dynamic range and demonstrate that such a solution yields an improved viewing experience. Yoshida et al. (2006) conducted user studies on HDR displays to develop a general mapping mechanism between different dynamic ranges. Surprisingly, HDR displays are even preferred if a completely simplistic linear scaling factor is used (Oğuz Akyüz, Fleming, Riecke, Reinhard, & Bülthoff, 2007).

The image processing can be avoided by assuming that reality is a desirable benchmark for images. This is quite reasonable since we have evolved in this environment and presumably tuned our preference measures to it. Ledda et al. compared a basic dual modulation display against real scenes and found that the HDR display is a compelling representation solution for real scenes (Ledda, Chalmers, & Seetzen, October 2004).

Not captured by any of these evaluations or comparisons is the immersive experience of a large scale HDR image. The three dual modulation projection systems in chapter 7 can illuminate very large screens that fill the complete field of view. With proper calibration such a screen can be made to exactly represent real scene luminance and colours to create an eerily realistic viewing experience. The only improvement would be to enclose the viewer with a surrounding high dynamic range environment for complete immersion (Ghosh, Trentacoste, Seetzen, & Heidrich, 2005).

9 Conclusions

High dynamic range imaging is a gateway to a new viewer experience. Not only does it remove one of the last fundamental shortcomings of two dimensional imaging devices, it blurs the line between real scenes and display imagery in terms of luminance, contrast and colour. Recent advances in the areas of image capture, processing and transmission have made it possible to deliver HDR content to the display system. The techniques described in this dissertation enable the display of this content and effectively close the loop on the HDR chain.

The enabling condition for this HDR chain is our ability to perceive a wide range of luminance. Our visual system has evolved in an environment with a variety of illumination environments. Using the mechanisms outlined in the beginning of this dissertation we are able to resolve up to five orders of magnitude of luminance values at any point in time and further expand this impressive range through temporal adaptation. Moreover, not only do we have the capability to psychophysically process such a high dynamic range of luminance, we in fact like seeing it. The viewer preference studies described herein show a strong preference for higher contrast and luminance than conventional displays can deliver. This creates the demand for HDR experiences.

The technology concepts described in this dissertation provide the solutions to address this demand. The two front-view display solutions enable a range of applications. The basic dual modulation concept provides a scientific HDR visualisation tool with very low development cost and very high accuracy. This is augmented by the advanced dual modulation concept for mass consumer applications where a slimmer package and reliable solid state design are instrumental for success.

The three projection concepts further extended the dual modulation concept into the large screen application space. The basic projection concept supports mid sized applications and professional visualisation. It would also be easily adaptable to more commercial applications such as home theatre. The distributed projection technique extends the benefits of HDR imagery to very large venues that can modify the existing imaging architecture including the screen. Finally, the active screen design offers a highly efficient path to high dynamic range in a compatible approach to conventional cinema installations.

Using the design, algorithm and system architectures described herein it is possible to implement each of these concepts into effective devices. All five solutions deliver exceptionally high dynamic range of luminance as well as a number of additional advantages. These advantages range from efficiency improvements such as energy saving to image quality enhancement in the form of colour and amplitude resolution improvements. In short, they enable a display experience that is beyond compare and outperform all conventional display solutions available today. Such is the attraction of these concepts that nearly every major television manufacturer has commercially introduced advanced dual-modulation displays using the techniques described in this dissertation.

BIBLIOGRAPHY

References indicated with an asterisk (*) have been incorporated directly into this dissertation.

American Society for Testing and Materials (ASTM). (1987). *ASTM Standards for Color and Appearance Measurement.* Philadelphia, PA: American Society for Testing and Materials.

Ashikhmin, M. (2002). A tone mapping algorithm for high contrast images. *In Proc. of Eurographics Workshop on Rendering 2002*, (pp. 145–156).

Barron, C., & Cotta Vaz, M. (2002). *The Invisible Art: The Legends of Movie Matte Painting.* Chronicle Books.

Barten, P. (1992). Physical model for the contrast sensitivity of the human eye. *In Proc. SPIE, vol. 1666*, (pp. 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.*

Becker, M. E. (2006). Display reflectance: Basics, measurement, and rating. *Journal of the Society for Information Display 14* , 1003-1017.

Berns, R. S., J, M. R., & Gorzynski, M. E. (1993). CRT colorimetry, Part I: Theory and practice. *Color Research and Application 18* , 299-314.

Bimber, O., & Iwai, D. (2008). Superimposing dynamic range. *ACM Trans. Graph. (Proc. Siggraph Asia) Vol. 27 #5* , 1-8.

Boff, K. R., & Lincoln, J. E. (1988). Light Adaptation. In *Engineering Data Compendium: Volume I - Human Perception and Performance* (pp. 137-163). Wright-Patternson AFB, Ohio: Harry G. Armstrong Aerospace Medical Research Laboratory.

Bordes, N., Bleha, W. P., & Pailthorpe, B. (2003). Design and characterization of an ultraresolution seamlessly tiled display for data visualization. *Proceedings of SPIE.*

Chen, H., Sung, J., Ha, T., & Park, Y. (2007). Locally Pixel Compensated Backlight Dimming for Improving Static Contrast on LED Backlight LCDs. *SID Symposium Digest*, (pp. 1339-1342).

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*, (pp. 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*, (pp. 186–196).

Commission Internationale de L'Eclairage (CIE). (1979). *Absolute Methods For Reflection Measurements.* Technical Report CIA 44-1979.

Commission Internationale de L'Eclairage (CIE). (1987). *International Lighting Vocabolary.* Technical Report CIE 17.4-1097.

Commission Internationale de L'Eclairage (CIE). (1977). *Radiometric And Photometric Characteristics of Materials And Their Measurement.* Technical Report CIE 38-1977.

Commission Internationale de L'Eclairage (CIE). (1996). *The Relationship Between Digital and Colorimetric Data For Computer-Controlled CRT Displays.* Technical Report CIE 122-1996.

Damberg, G., Seetzen, H., Whitehead, L. A., & Ward, G. (2007). High Dynamic Range Projection Systems. *Proceedings of the 2006 Society for Information Display Annual Symposium.*

Debevec, P., & And Malik, J. (1997). Recovering high dynamic range radiance maps from photographs. *In Proc. of ACM SIGGRAPH '97*, (pp. 369–378).

Devlin, K., Chalmers, A., Wilkie, A., & Purgathofer, W. (2002). STAR Report on Tone Reproduction and Physically Based Spectral Rendering. *Eurographics 2002.*

DICOM. (2003). *Part 14: Grayscale Standard Display Function.* Rosslyn, Virginia: National Electrical Manufacturer's Assoc. Technical Report PS 3.14-2003.

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*, (pp. 257–266).

Fairchild, M. D. (2005). *Color Appearance Models, Second Edition.* Chichester, UK: Wiley-IS&T Series in Imaging Science and Technology.

Fairchild, M. D., & Wyble, D. R. (2007). Mean observer metamerism and the selection of display primaries. *IS&T/SID 15th Color Imaging Conference*, (pp. 151-156). Albuquerque.

Fattal, R., Lischinski, D., & Werman, M. (2002). Gradient domain high dynamic range compression. *ACM Trans. Graph. (special issue SIGGRAPH 2002) 21, 3*, (pp. 249–256).

Ferwerda, J. A., Pattanaik, S., Shirley, P., & Greenberg, D. P. (1996). A model of visual adaptation for realistic image synthesis. *Proceedings of ACM SIGGRAPH 1996*, (pp. 249-258).

Fisekovic, N., Nauta, T., Cornelissen, H., & Bruinink, J. (2001). Improved motion-picture quality of AM-LCDs using scanning backlight. *In IDW 2001 Proceedings*, (pp. 1637–1640).

Gardner, A., Tchou, C., Hawkins, T., & Debevec, P. (2003). Linear light source reflectometry. *ACM Trans. Graph. (special issue SIGGRAPH 2003) 22, 3*, (pp. 749–758).

Ghosh, A., Trentacoste, M., Seetzen, H., & Heidrich, H. (2005). Real Illumination from Virtual Environments. *ACM Transactions on Graphics, Special issue on Proc. of ACM SIGGRAPH 2004.*

Glassner, A. (1995). Principles of Digital Image Synthesis. San Francisco: Morgan Kaufman.

Heckaman, & R L Fairchild, M. D. (2007). Beyond the locus of pure spectral colors and the promise of high-dynamic-range display technology. *Information Display 23:7* , 22-27.

Heckaman, R. L., & Fairchild, M. D. (2006). Effect of DLP projector white channel on perceptual gamut. *Journal of the Society of Information Display 14* , 755-761.

Heckaman, R. L., & Fairchild, M. D. (2006). Expanding display color gamut beyond the spectrum locus. *Color Research and Application 31* , 475-482.

Hood, D. C., & Finkelstein, M. A. (1986). Sensitivity to light. In L. K. K. R. Boff, *Handbook of Perception and Human Performance: Volume I - Sensory Processes and Perception* (pp. 5/1-5/66). New York: John Wiley & Sons.

Hunt, R. W. (2005). *The Reproduction of Color, Sixth Edition.* Chichester, UK: Wiley-IS&T Series in Imaging Science and Technology.

International Color Consortium (ICC). (2004). *Image Technology Color Management – Architecture, Profile Format, and Data Structure.* Specification ICC.1:2004-10 / Profile version 4.2.0.0.

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

Kang, S., Uyttendaele, M., Winder, S., & Szeliski, R. (2003). High dynamic range video. *ACM Trans. Graph. (special issue SIGGRAPH 2003) 22, 3*, (pp. 319–325).

Kelley, E. F., & Jones, G. R. (1997). Utilizing the bi-directional reflection distribution function to predict reflections from FPDs. *SID Symposium Digest of Technical Papers, 28*, (pp. 831-834).

Kelly, E. F. (2006). *Display Measurements for Flat Panel Displays.* SID Application Tutorial Notes, A-6/1-A-6/73.

Kennel, G. (2006). *Color and Mastering for Digital Cinema.* New York: Focal Press.

Krantz, J. H., Silverstein, L. D., & Yeh, Y. Y. (1992). Visibility of transmissive liquid crystal displays under dynamic lighting conditions. *Human Factors, 34* , 615-632.

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

Larson, G. W., Rushmeier, H., & 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.

Ledda, P., Chalmers, A., & Seetzen, H. (October 2004). High Dynamic Range Displays: a Validation Against Reality. *International Conference on Systems, Man and Cybernetics.*

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

Li, Y., Sharan, L., & Adelson, E. (2005). Compressing and Companding High Dynamic Range Images with Subband Architectures. *Proc. of SIGGRAPH '05, Special issue of ACM Transactions on Graphics* .

Lubin, J., & Pica, A. P. (1991). A non-uniform quantizer matched to human visual performance. *Society for Information Display Technical Digest* , 619-622.

Mann, S., & 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.

Mantiuk, R., Efremov, A., Myszkowski, K., & Seidel, H.-P. (2006). Backward compatible high dynamic range MPEG video compression. *ACM Trans. Graph. 25(3)* , 713-723.

Mantiuk, R., Krawczyk, G., Myszkowski, K., & Seidel, H.-P. (2004). Perception-motivated High Dynamic Range Video Encoding. *Proc. of SIGGRAPH '04 Special issue of ACM Transactions on Graphics* .

Mantiuk, R., Myszkowski, K., & Seidel, H.-P. (2005). A Perceptual Framework for Contrast Processing of High Dynamic Range Images. *Proc. of Second Symposium on Applied Perception in Graphics and Visualization*, (pp. 87-94).

Mark, W. R., Glanville, R. S., Akeley, K., & 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*, (pp. 896–907).

Merrifield, R. M., & Silverstein, L. D. (1988). An automatic brightness and contrast control system for vehicular displays. *Society for Information Display Technical Digest* , 178-181.

Mitsunaga, T., & Nayar, S. K. (1999). Radiometric self calibration. *In Proc. of IEEE CVPR*, (pp. 472–479).

Moon, P., & Spencer, D. (1945). The visual effect of non-uniform surrounds. *Journal of the Optical Society of America 35, 3* , 233–248.

Mossman, M. A., Kotlicki, A., Whitehead, L. A., Biernath, R. W., & Rao, S. P. (2001). New Reflective Color Display Technique Based on Total Internal Reflection and Subtractive Color Filtering. *Society for Information Display Annual Symposium.*

Muka, E., & Reiker, G. (2002). Reconsidering bit depth for radiological images – is eight enough? *In Proc. SPIE, vol. 4686*, (pp. 177–188).

Oğuz Akyüz, A., & Reinhard, E. (2006). Color Appearance in High Dynamic Range Imaging. *Journal of Electronic Imaging* .

Oğuz Akyüz, A., Fleming, R., Riecke, B., Reinhard, E., & Bülthoff, H. (2007). Do HDR displays support LDR content? A psychophysical evaluation. *ACM Transactions on Graphics 26 (3)* , Article 38.

Ohtsuki, H., Nakanishi, K., Mori, A., Sakai, S., Yachi, S., & 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*, (pp. 1154–1157).

Olsson, S., Stroppiana, M., & Baïna, J. (1997). Objective methods for assessment of video quality: state of the art . *IEEE Trans. on Broadcasting, vol. 43, no. 4* , 487-95.

Pattanaik, S. N., Ferwerda, J. A., Fairchild, M. D., & And Greenberg, D. (1998). A multiscale model of adaptation and spatial vision for realistic image display. *Proceedings of the 25th Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH '98.*

Pinson, M. H., & Wolf, S. (n.d.). Comparing subjective vieo quality testing methodologies. *SPIE 2003 Volume 5150* , 573.

Poynton, C. (2003). *Digital Video and HDTV: Algorithms and Interfaces.* San Francisco: Morgan Kaufmann Publishers.

Poynton, C. (2007). Wide-gamut displays. *Information Display 23:7* , 10-15.

Pratt, W. K., Sawkar, S. S., & O'Reilly, K. (January 1998). Automatic blemish detection in liquid crystal flat panel displays. *IS&T/SPIE symposium on Electronic Imaging: Science and Technology.*

Reinhard, E., Stark, M., Shirley, P., & Ferwerda, J. (2002). Photographic tone reproduction for digital images. *ACM Trans. Graph. (special issue SIGGRAPH 2002) 21, 3*, (pp. 267–276).

Reinhard, E., Ward, G., Pattanaik, S., & Debevec, P. (2005). *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting.* San Francisco: Morgan Kaufmann Publishers.

Rempel, A., Seetzen, H., Young, D., Trentacoste, M., Heidrich, W., Ward, G., et al. (August 2007). Dynamic Range Expansion for Legacy Video and images. *ACM Transaction on Graphics, special issue on Proc. of ACM SIGGRAPH 2007* .

Robertson, M., & Borman, S. S. (1999). Dynamic range improvements through multiple exposures. *In Proc. of International Conference on Image Processing (ICIP) ’99*, (pp. 159–163).

Robinson, M. G., Chen, J., & Sharp, G. D. (2005). *Polarization Engineering for LCD Projection.* John Wiley & Sons Ltd.

Schlick, C. (1994). Quantization techniques for visualization of high dynamic range pictures. *In Proc. of Eurographics Workshop on Rendering ’94*, (pp. 7–20).

Seetzen, H., Heidrich, W., Stuerzlinger, W., Ward, G., Whitehead, L. A., Trentacoste, M., et al. (August 2004). High Dynamic Range Display Systems. *ACM Transactions on Graphics, special issue on Proc. of ACM SIGGRAPH 2004* .

Seetzen, H., Li, H., Ye, L., Heidrich, W., Whitehead, L. A., & Ward, G. (June 2006). Observation of Brightness, Contrast and Amplitude Resolution of Displays. *Proceedings of the 2006 Society for Information Display Annual Symposium.*

Seetzen, H., Makki, S., Ip, H., Wan, T., Kwong, V., Ward, G., et al. (February 2007). Self-calibrating wide color gamut high dynamic range display. *Proceedings of SPIE – Volume 6492.*

Seetzen, H., Wan, T., & Longhurst, P. (2008). Advanced High Dynamic Range Range Displays. *International Display Workshop.* Niigata.

Seetzen, H., Ward, G., & Whitehead, L. A. (December 2006). Dual-Modulation High Dynamic Range Display Technology. *Proceedings of the 2006 International Display Workshop.*

Seetzen, H., Ward, G., & Whitehead, L. A. (November 2005). High Dynamic Range Imaging Pipeline. *Proceedings of the 2005 Americas Display Engineering and Application Conference.*

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

Silverstein, L. D. (2004). *Color in Electronic Displays.* Society for Information Display Seminar Lecture Notes, Volume 1, M-13/1-M13/63.

Silverstein, L. D. (2003). Display visibility in dynamic lighting environments: Impact on the design of portable and vehicular displays. *Proceedings of the International Display Manufacturing Conference*, (pp. 15-18).

Silverstein, L. D., & Merrifield, R. M. (1985). *The Development and Evaluation of Color Display Systems for Airborne Applications: Phase I - Fundamental Visual, Perceptual, and Display System ConsiderationsTechnical Report DOT/FAA/PM-85-19.*

Spaulding, K., Woolfe, G. J., & Joshi, R. L. (n.d.). Extending the Color Gamut and Dynamic Range of an sRGB Image using a Residual Image. *Color Res. Appl. 28* .

Sproson, W. N. (1983). *Colour Science in Television and Display Systems.* Bristol, England: Adam Hilger, Ltd.

Stockman, MacLeod, & Johnson. (1993). *Journal of the Optical Society of America A* , 2491-2521.

Trentacoste, M. (2006). *Photometric Image Processing for High Dynamic Range Displays.* University of British Columbia, MSc Thesis.

Trentacoste, M., Heidrich, W., Whitehead, L. A., Seetzen, H., & Ward, G. (2007). Photometric Image Processing for High Dynamic Range Displays. *Journal of Visual Communication and Image Representation - Special Issue on HDRI* .

Tumblin, J., & Rushmeier, H. (1993). Tone reproduction for realistic images. *IEEE Computer Graphics and Applications 13, 6*, (pp. 42–48).

Tumblin, J., & Turk, G. (1999). LCIS: A boundary hierarchy for detail-preserving contrast reduction. *In Proc. of ACM SIGGRAPH '99*, (pp. 83–90).

Video Electronics Standards Association (VESA). (2001). *Flat Panel Display Measurements Standard (FPDM), Version 2.0.* Milpitas, CA.

Vos, J. (1984). Disability glare -a state of the art report. *CIE Journal 3, 2* , 39–53.

Wandell, B. A., & Silverstein, L. D. (2003). Digital color reproduction. In S. (. Shevell, *The Science of Color, 2nd Edition* (pp. 281-316). Washington, D. C: Optical Society of America.

Ward Larson, G. (November 1998). Overcoming Gamut and Dynamic Range Limitations in Digital Images. *IS&T/SID 6th Color Imaging Conference.*

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

Ward, G. (April 2002). A Wide Field, High Dynamic Range, Stereographic Viewer. *Proceedings of PICS 2002.*

Ward, G. (1991). Real Pixels. In J. (. Avro, *Graphics Gems II* (pp. 80-83). Academic Press.

Ward, G., & Simmons, M. (November 2005). JPEG-HDR: A Backwards-Compatible, High Dynamic Range Extension to JPEG. *Proceedings of the Thirteenth Color Imaging Conference.*

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

Wyszecki, G., & Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition.* New York: John Wiley & Sons, Inc.

Yoshida, A., Mantiuk, R., Myszkowski, K., & Seidel, H.-P. (2006). Analysis of Reproducing Real-World Appearance on Displays of Varying Dynamic Range. *Computer Graphics Forum (Proc. of Eurographics'06) 25 (3)*, (pp. 415-426).

Zou, W. Y., & Corriveau, P. J. *Methods for evaluation of digital television picture quality, Doc.G-2.1.6/28.* Boulder, Colorado: IEEE Broadcast Tech. Soc. – 4th meeting G-2.1.6, Audio Video Techniques Committee G-2.1, Compression and Processing Subcommittee.

PATENT REFERENCES

The concepts described in this document are in part the subject matter of the following US and PCT patents and applications. International counterparts to the listed patents and applications have been omitted for brevity. All patents are assigned to Dolby Licensing Corp.

US 6,891,672 L. A. Whitehead, G. J. Ward, W. Stuerzlinger, H. Seetzen, High Dynamic Range Display Devices, issued May 10th 2005

US 7,106,505 L. A. Whitehead, G. J. Ward, W. Stuerzlinger, H. Seetzen, High Dynamic Range Display Devices, issued September 12th 2006

US 7,172,297 L. A. Whitehead, G. J. Ward, W. Stuerzlinger, H. Seetzen, High Dynamic Range Display Devices, issued February 6th 2007

US7370979 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, W. Stuerzlinger, Calibration of Displays Having Spatially Variable Backlight, issued May 13th 2008

US7377652 L. A. Whitehead, H. Seetzen, W. Stuerzlinger, G. J. Ward, High Dynamic Range Displays Having Location Specific Modulation, issued May 27th 2008

US7403332 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, W. Stuerzlinger, High Dynamic Range Display Devices, issued July 22nd 2008

US7413307 L. A. Whitehead, H. Seetzen, W. Stuerzlinger, G. J. Ward, High Dynamic Range Colour Display, issued August 19th 2008

US7413309 L. A. Whitehead, H. Seetzen, W. Stuerzlinger, G. J. Ward, High Dynamic Range Colour Display, issued August 19th 2008

US7419267 L. A. Whitehead, H. Seetzen, W. Stuerzlinger, G. J. Ward, High Dynamic Range Displays with Overlapping Dual Modulation, issued September 2nd 2008

US12/250,435 L. A. Whitehead, H. Seetzen, W. Stuerzlinger, G. J. Ward, High Dynamic Range Displays and Control System therefore, filed October 13th 2008

PCT/CA2002/00255 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, High Dynamic Range Display Devices, filed February 27th 2002

US11/831,749 L. A. Whitehead, H. Seetzen, W. Stuerzlinger, G. J. Ward, High Dynamic Range Displays and Control System therefore, filed July 31st 2007

PCT/CA03/00350 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, W. Stuerzlinger, High Dynamic Range Display Devices, filed May 27th 2005

US11/608,771 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, W. Stuerzlinger, Motion-Blur Compensation in Backlit Displays, filed December 8th 2006

US11/831,851 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, W. Stuerzlinger, High Dynamic Range Display with Individual Color Backlights, filed July 31st 2007

US11/831,827 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, W. Stuerzlinger, High Dynamic Range Displays with Dual Modulators with Different Resolutions, filed July 31st 2007

US11/831,864 L. A. Whitehead, H. Seetzen, G. J. Ward, D. Graham, W. Stuerzlinger, High Dynamic Range Displays Having Light Estimating Controllers, filed July 31st 2007

PCT/CA2004/002198 L. A. Whitehead, H. Seetzen, G. J. Ward, Methods and Systems for Converting Images from Low Dynamic Range to High Dynamic Range, filed December 24th 2004

US10/599,955 L. A. Whitehead, H. Seetzen, G. J. Ward, Methods and Systems for Converting Images from Low Dynamic Range to High Dynamic Range, filed December 24th 2004

US11/831,521 L. A. Whitehead, H. Seetzen, G. J. Ward, Methods and Systems for Converting Images from Low Dynamic Range to High Dynamic Range, , filed July 31st 2007

PCT/CA04/002197 H. Seetzen, Methods for Efficient Computation of Image Frames for Dual Modulation Display Systems using Key Frames, filed December 24th 2004

US10/599,954 H. Seetzen, Methods for Efficient Computation of Image Frames for Dual Modulation Display Systems using Key Frames, filed December 24th 2004

US11/831,920 H. Seetzen, Driving Dual Modulation Display Systems Using Key Frames, filed July 31st 2007

PCT/CA05/000807 L.A. Whitehead, H. Seetzen, G. J. Ward, W. Heidrich, Rapid Image Rendering on Dual-Modulation Displays, filed May 27th 2005

US11/627,936 L.A. Whitehead, H. Seetzen, G. J. Ward, W. Heidrich, Rapid Image Rendering on Dual-Modulation Displays, filed May 27th 2005

US11/831,854 L.A. Whitehead, H. Seetzen, G. J. Ward, W. Heidrich, Rapid Image Rendering on Dual-Modulation Displays, filed May 27th 2005

PCT/CA2004/002200 H. Seetzen, Wide Color Gamut Displays, filed December 24th 2004

US11/722,707 H. Seetzen, Wide Color Gamut Displays, filed December 24th 2004

US11/831,922 H. Seetzen, Wide Color Gamut Displays, filed July 31st 2004

PCT/CA05/001975 H. Seetzen, Field Sequential Display of Color Images, filed December 23rd 2005

US11/722,706 H. Seetzen, Field Sequential Display of Color Images, filed December 23rd 2005

PCT/CA06/000542 H. Seetzen, L. A. Whitehead, G. J. Ward, Modular Electronic Displays, filed April 18th 2006

US11/612,410 G. J. Ward, H. Seetzen, Representing and Reconstructing High Dynamic Range Images, filed December 18th 2006

US11/831,630 G. J. Ward, H. Seetzen, Representing and Reconstructing High Dynamic Range Images, filed July 31st 2007

PCT/US2007/002788 W. Heidrich, M. Trentacoste, H. Seetzen, Multiple Modulator Displays and Related Methods, January 31st 2007

US12/021,214 L. Silverstein, L. A. Whitehead, G. J. J. Ward, H. Seetzen, V. Kwong, G. Thomas, Retroreflective optical collimation layer and display devices incorporating same, filed January 28th 2008

US12/021,206 G. J. J. Ward, G. Damberg, H. Seetzen, Serial modulation display having binary light modulation stage, filed January 29th 2008

US12/030,448 G. Damberg, H. Seetzen, Temporal Filtering of Video Signals, filed February 13th 2008

PCT/CA06/000559 L. A. Whitehead, M. A. Mossman, H. Seetzen, Brightness Enhancement in TIR-modulated Electrophoretic Reflective Image Displays

PCT/CA2008/001413 H. Seetzen, A. Rempel, W. Heidrich, G. Ward, L. A. Whitehead, Enhancing Dynamic Range of Images, filed July 30th 2008

US12/182,121 H. Seetzen, A. Rempel, W. Heidrich, G. Ward, L. A. Whitehead, Enhancing Dynamic Range of Images, filed July 29th 2008

US12/183,033 H. Seetzen, A. Rempel, W. Heidrich, G. Ward, L. A. Whitehead, Enhancing Dynamic Range of Images, filed July 30th 2008