

Color Nonuniformity in Projection-Based Displays: Analysis and Solutions

Aditi Majumder, Rick Stevens

Abstract—

Large-area displays made up of several projectors show significant variation in color. In this paper, we *identify* different projector parameters that cause the color variation and *study* their effects on the luminance and chrominance characteristics of the display. This work leads to the realization that luminance varies significantly within and across projectors while chrominance variation is relatively small, especially across projectors of same model.

To address this situation, we present a method to achieve luminance matching across all pixels of a multiprojector display that results in photometrically uniform displays. We use a camera as measurement device for this purpose. Our method comprises a one-time calibration step that generates a *per channel per projector luminance attenuation map (LAM)*, which is then used to correct any image projected on the display at interactive rates on commodity graphics hardware. To the best of our knowledge, this is the first effort to match luminance across all the pixels of a multiprojector display.

Keywords— Projection-Based Displays, Tiled Displays, Color Calibration

I. INTRODUCTION

Large-area, multiprojector displays offer high resolution, large field of view, and a compelling sense of presence. Thus, they are extremely useful for visualizing large scientific models and in immersive virtual environments used for 3D teleconferencing and entertainment purposes. Several such displays exist at Princeton, the University of North Carolina at Chapel Hill, Stanford, the Fraunhofer Institute (Germany), and different U.S. national laboratories. Recent efforts are directed toward building large displays comprising 40-50 projectors (Sandia National Laboratories and the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign).

The color of these displays shows significant spatial variation, which can be very distracting, thus breaking the illusion of having a single display. This problem is unique to such displays and can be caused by device-dependent reasons like *intraprojector color variation* (color variation within a single projector) and *interprojector color variation* (color variation across different projectors) or by other device-independent reasons such as non-Lambertian curved display surface and interreflections [1], [2], [3]. Further, deliberate alignment of projectors in an overlapping fashion introduces significant color variation.

Some existing solutions try to reduce the higher brightness in the overlap regions by blending techniques [4], [5] implemented either in software or optically. But since this

Department of Computer Science, University of North Carolina at Chapel Hill. Email: majumder@cs.unc.edu .

Department of Computer Science, University of Chicago/ Mathematics and Computer Science Division, Argonne National Laboratory. E-mail: stevens@mcs.anl.gov .

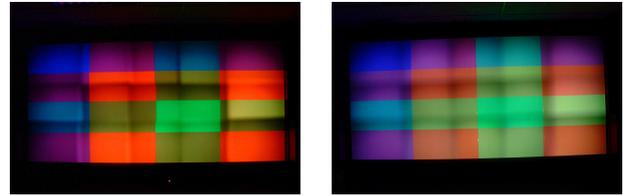


Fig. 1. Left: A display of 5×3 array of 15 projectors where the overlap regions are blended by using a physical shadow mask on the light path of the projector. Right: The same display with the overlap region blended by a linear ramp in software. For this, it is necessary to have the knowledge of the exact location of the overlap region. Note that the overlapping regions are distinctly noticeable in both cases.

approach does not account for either intra- or interprojector variations, the seams between projectors are still visible, and one can easily notice the boundaries of the projectors that make up the display, as shown in Figure 1. The solution presented in [1] matches the luminance across multiple projectors but does not account for the variation within a single projector's field of view and hence fails to generate photometrically uniform displays. The comments in recent work [6], [7], [8], [9], [10], [11], [12], [13] and our experience have led us to believe that this problem is non-trivial and needs to be analyzed in a structured manner.

A. Main Contributions

1. In this paper we first *identify* the different device-dependent parameters of a projector that can cause color variation in a multiprojector display such as position, zoom, axis of alignment, lamp age and projector controls of brightness, contrast, and white balance. Next, we *analyze the effects of the changes* in these parameters on the color variation and provide insights into the possible reasons for these variations. There has been some work on characterizing specifically the color gamuts of both LCD and DLP projectors [2], [3]. Our study in this paper complements this work.
2. From this analysis we make the *key observation* that the most significant cause of the spatial variation in color of a multiprojector display is the variation in luminance. Most tiled displays are made of multiple projectors that vary little in chrominance. Further, humans are more sensitive to luminance variation than to chrominance variation [14]. Hence, we may be able to achieve acceptable photometric uniformity by correcting for luminance variation alone.
3. Next, we present a method to do a per channel per pixel luminance matching. In the one-time calibration step, we first use a camera to measure the per channel luminance response of a multiprojector display and find the pixel with

the most “limited” luminance response. Then, for each projector, we generate a per channel *luminance attenuation map (LAM)* that assigns a weight to every pixel of the projector to scale the luminance response of that pixel to match the most limited response. This LAM is used to correct any image projected by the projector.

4. To the best of our knowledge, the algorithm presented here is the first effort to solve for all different kinds of luminance variations (within a single projector, across different projectors, and in overlap regions) in an *automated, unified manner* that is completely transparent to the user.

5. To the best of our knowledge, this is the first effort to achieve photometric uniformity in multiprojector displays by using a commodity off-the-shelf product like an inexpensive digital camera. Previous work in this direction [1] uses a high-precision, expensive point measurement instrument such as a radiometer, which is impractical for high-resolution measurements.

6. The image correction using LAM can be implemented at interactive rates on commodity graphics hardware to correct any image projected on the display.

This paper is organized in two parts. The first part (Section II) presents a detailed analysis of the nature of color variation. We study the *intra-* and *interprojector* color variations. Next, we identify and study the parameters that causes the color variation. Aided by the observations made in this part, in the second part of the paper (Section III) we present an algorithm to match the luminance response of every pixel of a multiprojector display followed by the implementation details, results, and issues. Finally, we conclude with future work in Section IV.

II. ANALYSIS OF THE COLOR VARIATION PROPERTIES

In this part of the paper, we present the different experiments performed to study the nature of the color variation across a multiprojector display in detail. At the end of this section, we infer some important properties of this color variation from the results of these experiments.

A. Background

Color can be specified by three parameters (Y, x, y) . Y is the luminance (the amount of achromatic light present in a color) and (x, y) are the chromaticity coordinates that define chrominance.

When two colors $c_1 = (Y_1, x_1, y_1)$ and $c_2 = (Y_2, x_2, y_2)$ in proportions p_1 and p_2 are combined additively (as the primaries are combined in display systems), such that $p_1 + p_2 = 1$, the resulting color produced is $c_3 = (Y_3, x_3, y_3)$, where

$$Y_3 = Y_1 + Y_2; \quad x_3 = p_1 x_1 + p_2 x_2; \quad y_3 = p_1 y_1 + p_2 y_2. \quad (1)$$

The three colors (say R,G,B) used to create a display are called primaries, and the input paths for these primaries are called channels. The input for each primary has a range from 0.0 to 1.0. Let the colors projected by the display for the input of 1.0 at each channel (and 0.0 in other two channels) be (Y_R, x_R, y_R) , (Y_G, x_G, y_G) , and (Y_B, x_B, y_B) ,

respectively. The triangle formed by the chromaticity coordinates of these primaries is called the *color gamut* of the display. Readers are referred to [15] for additional in-depth treatment on colorimetry.

Ideally, it is desirable to have a display where, given the properties of the primaries, one can predict, using simple formulae, the properties of any color produced by the combination of the primaries. This becomes easy if the display satisfies the following properties.

1. *Channel Independence*: This assumes that the light projected from one channel is independent of the other two. Thus, this indicates that no leakage light from other channels interferes with the light projected from a channel.
2. *Channel Constancy*: This assumes that only luminance changes with changing channel inputs. For input $0.0 \leq r \leq 1.0$, the chromaticity coordinates (x_r, y_r) of r are constant at (x_R, y_R) , and only the luminance Y_r changes.
3. *Spatial Homogeneity*: The response of all the pixels of the display is identical for any input.
4. *Temporal Stability*: The response for any input at any pixel of the display does not change with time.

The property of optical superposition states that light falling at the same physical location from different sources adds up. The properties of channel constancy, independence and superposition along with the assumption that with an input of $(0, 0, 0)$ the display outputs zero light, indicate that the color projected at a pixel is a linear combination of the color projected by the maximum values of the red, green, and blue channels alone when the values of the other two channels are set to zero. Hence, for any input $c = (r, g, b)$, $0.0 \leq r, g, b \leq 1.0$, the luminance Y_c is given by $Y_r + Y_g + Y_b$, and the chromaticity coordinate is given by the barycentric coordinates

$$x_c = r x_R + g x_G + b x_B; \quad y_c = r y_R + g y_G + b y_B. \quad (2)$$

This is referred to as the *linear combination property*.

Given the linear combination, the spatial homogeneity, and the temporal stability property, we can predict the color at any pixel at any input from the response of the primaries at any one pixel of the display. Most traditional display devices, such as CRT monitors, satisfy these properties to a reasonable accuracy, or the deviation from this ideal behavior is simple enough to be modeled by simple linear mathematical functions [16]. However, as we will see in the following sections, a projector is not such an ideal device.

Before we delve deep into the study of the property of the color variation, we present the details of the measurement process, instruments, and other factors that may have an effect on our analysis.

B. Measurement

As test projectors for our experiments, we used multiple Sharp, NEC, Nview, Proxima, and Epson projectors (both LCD and DLP) and both front- and back-projection systems. The graphs and charts we present are only samples of the different similar results we have achieved.

B.1 Measuring Devices

The goal of the process of measurement is to find the luminance and the chrominance properties at different points of the tiled display accurately. There are two options for the optical sensors that one might use for this purpose.

A *spectroradiometer* is an expensive precision-instrument that can measure any color projected by the projector accurately, enabling us to work in a laboratory-calibrated device-independent color space. But, it can measure only one point at a time at a very slow rate of about 1-20 seconds per measurement. Further, it is not possible to measure the response of every pixel separately at a reasonable geometric accuracy. Thus, it is unsuitable for acquiring high resolution data.

A *camera*, on the other hand, is relatively inexpensive and is suitable for acquiring high-resolution data in a very short time. There are many existing algorithms to find the geometric correspondence between the camera coordinates and the projector coordinates assuring geometric accuracy of the data acquired. But, the limitations lie in its relative photometric inaccuracy.

With these options in hand, we use both types of sensors, but for different purposes. For point measurements, we used a precision spectroradiometer (Photo Research PR 715). But, for finding the spatial color/luminance variation across a projector, where we need to measure the color at potentially every pixel of a projector, we use a high-resolution digital camera. To reduce the photometric inaccuracies of the camera by a reasonable amount, we use the following methods.

Using the Camera as a Reliable Measuring Device: The non-linearity of the camera is recovered by using the algorithm presented in [17]. From this we generate a color look-up-table (LUT) that linearizes the camera response. Every picture from the camera is linearized by using this color LUT.

It is important that the camera not introduce additional spatial variation beyond that is already present on the display wall. Hence, the camera must produce flat fields when it is measuring a flat color. It is mentioned in [17] that most cameras satisfy this property at lower aperture settings, especially below $F8$. Our camera showed a standard deviation of 2-3% for flat field images. These flat field images were generated by taking pictures of nearly diffused planar surfaces illuminated by a studio light with a diffusion filter mounted on it.

To ensure that a camera image is not under- or overexposed, we run simple under- or oversaturation tests. The differing exposures are accounted for by appropriate scaling factors [17].

Finally, to assure geometric accuracy of the measurements, we use a geometric calibration method [18] to find accurate camera-to-projector correspondence.

We cannot measure all the colors projected by the projector if the color gamuts of the camera does not contain the gamut of the projector, thus restricting us to work in a device-dependent color gamut. However, we do not use the

camera for any chrominance measurements, but only luminance measurements. So, this does not pose a problem.

B.2 Screen Material and View Dependency

For experiments on front-projection systems we use a (close to) Lambertian screen that does not amplify the color variations. But the Jenmar screen we use for our back projection system is not Lambertian. Thus, the measuring devices are sensitive to viewing angles. We orient the spectroradiometer perpendicular to the point that is being measured. But for the camera, the view dependency cannot be eliminated. However, we use the camera in two cases. First, we use it only for qualitative analysis of the nature of the luminance variation in a single projector. Since we are not trying to generate an accurate quantitative model of the variation, the view dependency is not critical. Second, we use it in our algorithm presented in Section III, which corrects the luminance variation accurately for one viewer position anyway.

B.3 Ambient Light

We try to reduce ambient light seen by the sensors as much as possible by taking the readings in a dark room turning off all lights. When taking the measurement of a projector, we turn off all adjacent projectors. Further, we use black material to cover up the white walls of the room to avoid interreflected light.

In the next few sections we study the intra- and inter-projector color properties from the measurements taken by using a spectroradiometer or a camera.

C. Intraprojector Variations

First, we study the intraprojector variations and show that the projectors do not follow the desirable properties mentioned in Section II-A. A few of these results are also confirmed in [19], [2], [3].

One important consequence of a display to satisfy *channel independence* and *channel constancy* is that the response for black (input of $(0, 0, 0)$) should have zero light. In projectors, however, some leakage light is projected even for black. This is called the *black offset*. Hence the chromaticity for any channel at zero is the chromaticity of this achromatic black. As the inputs increase, the chromaticity reaches a constant value, as it should for a device following channel constancy. This is demonstrated in Figure 2. The contours in Figure 2 show how the gamut starts out as a single point for 0 in all three channels and then attains the final red triangle at the highest input value. However, if this black offset is modeled by a linear offset term subtracted from the response of all inputs, the chromaticity coordinate curves for the three channels (shown in right picture of Figure 2) will be constant [2], [3].

From the various measurements we had from the spectroradiometer, we found that this black offset can be up to 2% of the maximum luminance projected per channel. However, this can deviate considerably across different projectors. We studied the black offset of fifteen same-model

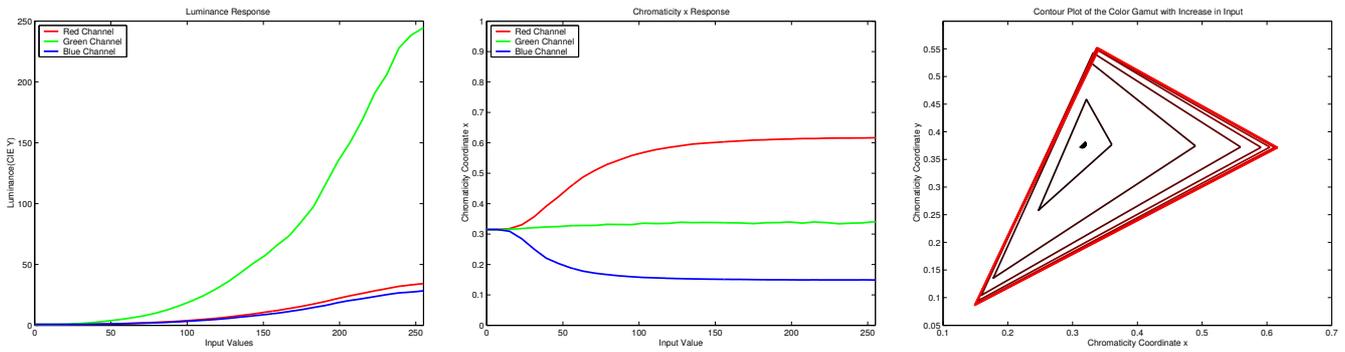


Fig. 2. Left: Luminance response of the three channels. Middle: Chromaticity x for the three channels. The shape of the curves for chromaticity y are similar. Right: Gamut contour as the input changes from 0 to 255 at intervals of 32.

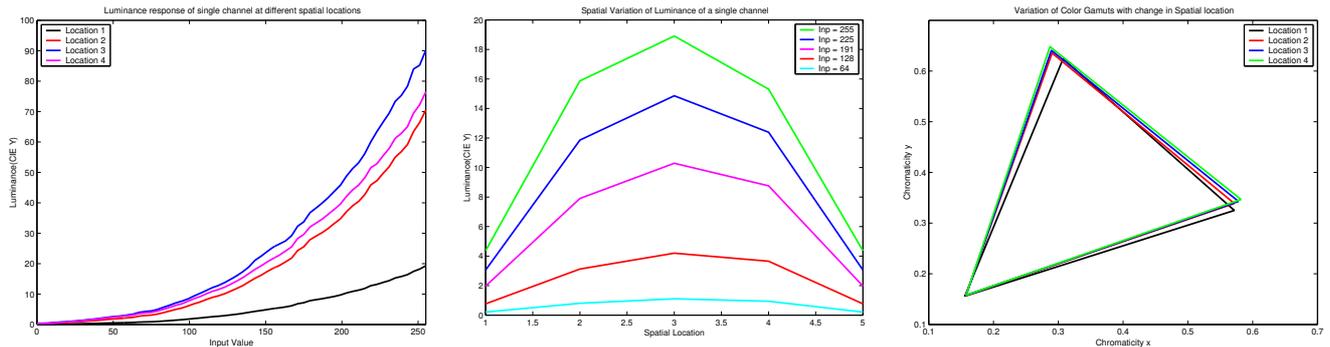


Fig. 3. Left: Luminance response of the red channel plotted against input at four different spatial locations. Middle: Luminance variation of different inputs of the red channel plotted against spatial location. The responses are similar for other channels. Right: Color gamut at four different spatial locations of the same projector.

projectors with similar control settings. They showed a relative standard deviation of about 25%.

If the black offset is accounted for by the linear offset term, almost all projectors exhibit the linear combination property. However, some DLP projectors do not exhibit the linear combination property for the grays. We found that instead of adding up the contributions from the red, green, and blue channels, these projectors use a clear filter while projecting the grays. Hence, luminance of the grays is much higher than the sum of the luminances of the constituting red, green, and blue. This is modeled in [2] by an additive gamut with an extrusion at the white point.

Projectors are not spatially homogeneous either. Accurate luminance and chrominance readings were taken at five equally spaced locations on the projector diagonal using the spectroradiometer. We named these locations from 1 to 5 starting at the top left corner. The luminance reaches a peak at the center (location 3) as seen in Figure 3. The luminance falls off at the fringes by a factor that may be as high as 80% of the peak luminance for rear-projection systems, and about 50% for front-projection systems. This considerable fall-off in luminance indicates that having wide overlaps between projectors in a multiprojector display can help us to get a better overall dynamic range.

Further, note that only the luminance changes spatially, while the color gamut remains almost identical, as shown in Figure 3. The gamut is measured from the chromaticity

coordinates of the primaries at their highest intensities. We measured the input chromaticity response at different spatial locations and found that the gamut does not vary spatially for the whole range of inputs.

Given these observations from the spectroradiometer measurements, we used a camera to measure the intra-projector spatial luminance variation at a much higher resolution (Figure 11). The readings are taken by carefully aligning the camera perpendicular to the screen. The luminance response shows a peak somewhere near the center, falling off radially toward the fringes asymmetrically. The pattern of this fall-off varies from projector to projector.

The above observations can be explained easily. The chrominance depends on the physical red, green, and blue filters of the projectors that do not change spatially within a single projector. Hence, the chrominance is spatially constant. The luminance fall-off is due to the distance attenuation of light, further amplified by the non-Lambertian nature of the display. The asymmetry in the fall-off pattern gets pronounced with off-axis projection, as we will see in the following sections. This indicates that the orientation of the projector is responsible for this asymmetry.

Finally, we find that the projectors are not temporally stable. The lamp in the projector ages with time and changes the color properties of the projector. Figure 9 shows a significant difference in luminance even within a short amount of time while the chrominance remains almost same. The luminance variation is also due to the

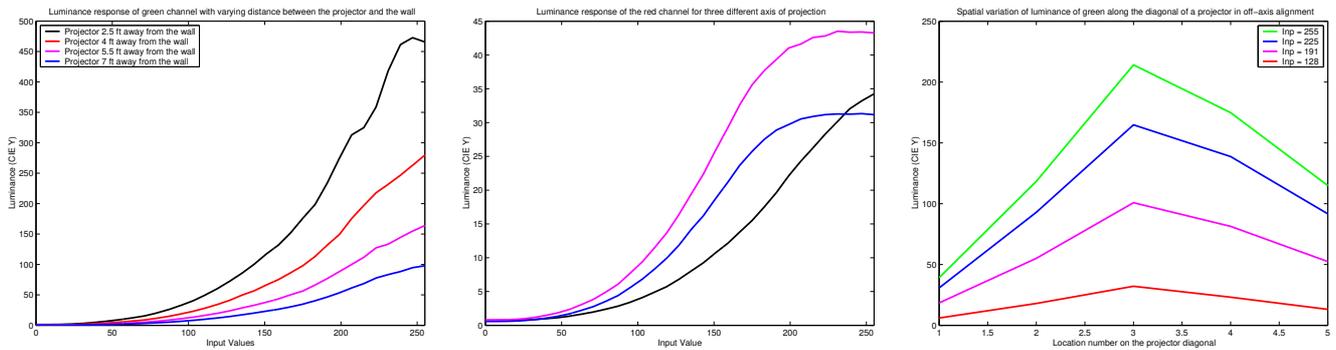


Fig. 4. Left: Luminance response of the green channel as the distance from the wall is varied along the axis of projection. Middle: Luminance Response of the red channel with varying axis of projection. Right: Luminance response of different inputs in the green channel plotted against the spatial location along the projector diagonal for oblique axis of projection.

unpredictable temporal change in the position of the arc in the lamp. Further, color characteristics also drift a little after extensive use of about 800-900 hours.

D. Projector Parameters That Change Color Properties

In this section, we identify the different projector parameters that can change the color properties of a projector, study the effects of varying these parameters on the color properties of a large area multiprojector display, and provide insights for the possible reasons behind such effects.

D.1 Position

Position is defined by the distance of the projector from the screen along the axis of projection and the alignment of the axis of projection with the planar display surface. We study the color properties with two sets of experiments. In one we keep the orientation constant while changing the distance from the screen, and in the other we keep the distance constant while changing the orientation.

1. *Distance to the Screen:* Figure 4 shows the luminance response as we move the projector at different positions along its axis of projection. The chrominance remains constant. Further, the shape of the spatial variance of the luminance also remains the same, as shown in Figures 3 and 11. By moving the projector away from the screen, the projection area increases. Hence the amount of light falling per unit area changes, but the nature of the fall-off does not change, as reflected in the observation.

2. *Orientation:* In this set of experiments, we kept the projector at the same distance from the screen while we rotated it about the x , y , and z direction to have an off-axis projection at four orientations from orthogonal to angled axis of projection of 60 degrees. In this case also we found that the chrominance remains constant while the luminance response changes. Figure 4 shows the results. The nature of the spatial variation is no longer symmetric as in the case of orthogonal position (Figure 3). Near the longer boundary of the key-stoned projection, which is physically farther away from the projector, there is a higher drop in luminance. As the orientation becomes more oblique, the luminance attenuation at the projector boundary farther away from the screen increases, resulting in asymmetric

fall-off. This is due to two reasons. First, the light from each pixel gets distributed over a larger area. Second, the angled surface receives less incident light because of the cosine effect. The results for vertical direction do not show a symmetry even for orthogonal projection because of the offset projection.

In both the above cases, since moving the projector around does not change the internal filters of the projector, the chrominance remains constant, as expected.

D.2 Controls

The projectors offer us various controls such as zoom, brightness, contrast, and white balance. Knowing how these controls affect the luminance and chrominance properties of the projector can help us decide the desirable settings for the projector controls that reduce variation within and across projectors. Thus, we can avail ourselves of the best possible dynamic range and color resolution offered by the device.

1. *Zoom:* We test the projector at four different zoom settings. Both luminance and chrominance remain constant with the change in zoom settings of the projector. With the change in zoom, the amount of light for each pixel gets distributed over a different area. For a focused projector it is distributed over a small area, while for a unfocused projector it is distributed over a larger area. However, the total area of projection remains the same, and the total amount of light falling in that area remains the same. Hence the light per unit area remains unchanged, while the percentage of light that each unit area receives from the different pixels changes.

2. *Brightness:* Luminance and chrominance response is measured by putting the brightness control at 5 different positions. Poynton [20] mentions that usually the brightness control in displays change the black offset. In projectors, however, this control affects both the gain and black offset of the luminance response of *all the three channels similarly and simultaneously*. The results are illustrated in Figure 5. As the brightness is increased, both the black offset and the gain of the luminance increase. If the brightness is too low, however, the luminance response gets clipped at the lower input range. In these settings, since the lumi-

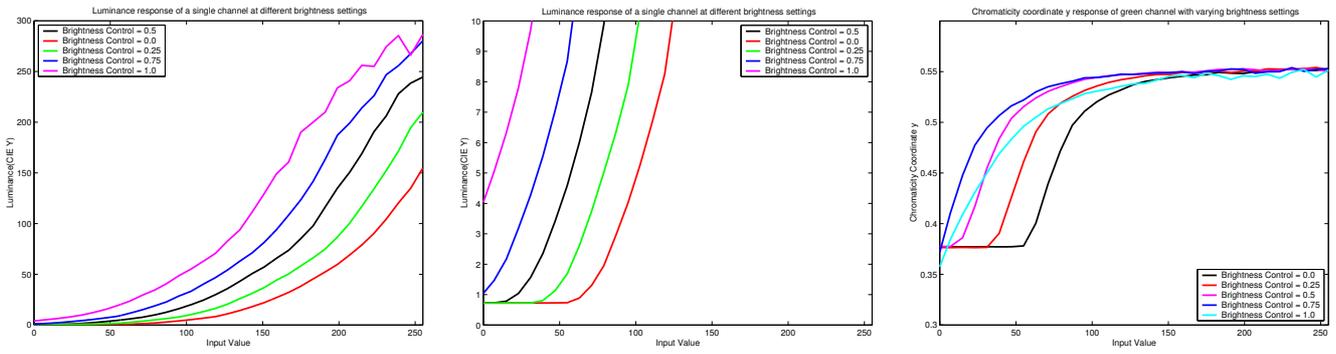


Fig. 5. Left: Luminance response of the green channel with varying brightness settings. Middle: Luminance response of the green channel with varying brightness settings zoomed near the lower input range to show the change in the black offset. Right: Chrominance response of the green channel with varying brightness settings.

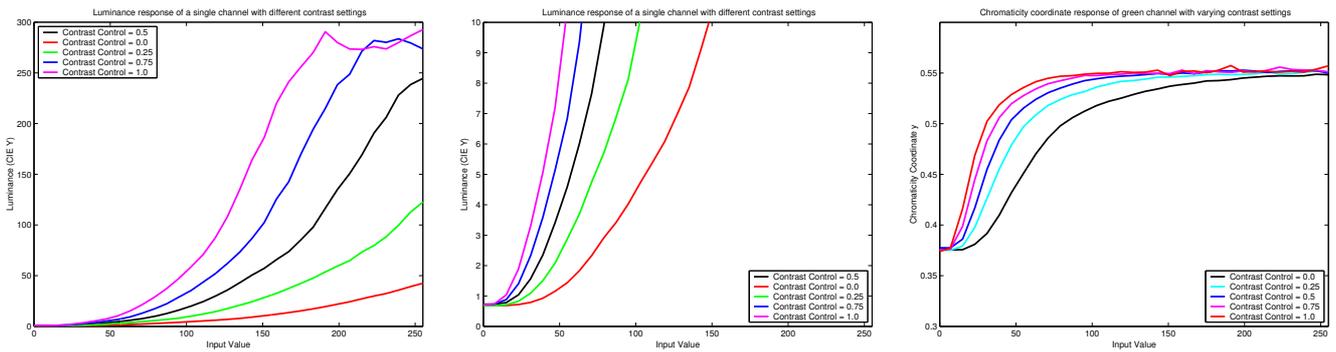


Fig. 6. Left: Luminance response of the green channel with varying contrast settings. Middle: Luminance response of the green channel with varying contrast settings zoomed near the lower luminance region to show that there is no change in the black offset. Right: Chrominance response of the green channel with varying contrast settings.

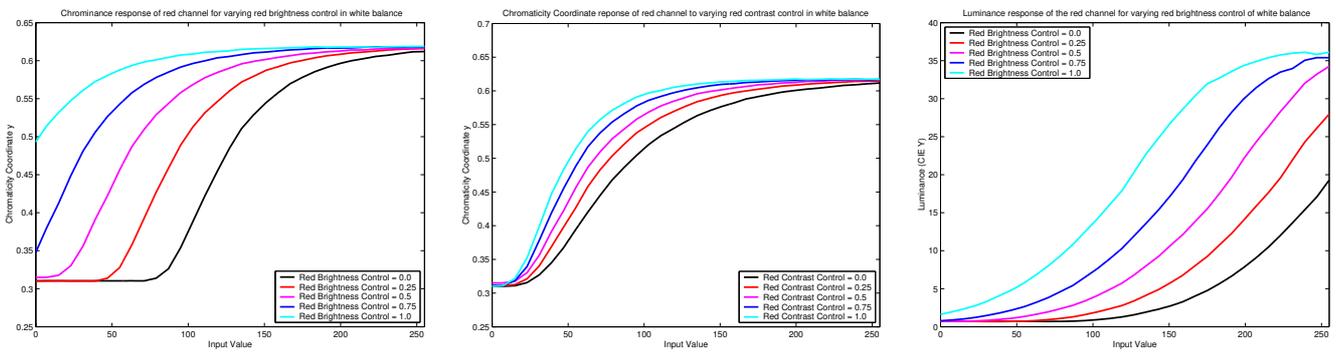


Fig. 7. Left: Chrominance response of the green channel with varying green brightness settings for white balance. Middle: Chrominance response of the red channel with varying red contrast settings for white balancing. Right: Luminance response of the red channel with varying red brightness settings in white balance.

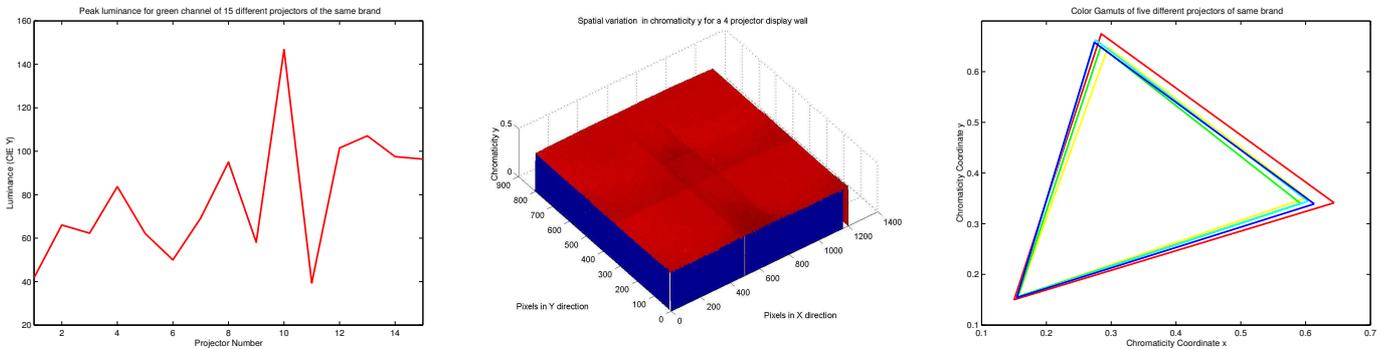


Fig. 8. Left: Peak luminance of green channel for fifteen different projectors of the same model with same control settings. Middle: Chrominance response of a display wall made of four overlapping projectors of same model. Right: Color gamut of 5 different projectors of the same model. Notice the large variation in luminance and small variation in chrominance.

TABLE I
CHROMATICITY COORDINATES OF THE PRIMARIES OF DIFFERENT
MODELS OF PROJECTORS

Projector Brand	Red		Green		Blue	
	x	y	x	y	x	y
Sharp XG-E3000U	0.62	0.32	0.33	0.62	0.14	0.07
NEC MT-1035	0.55	0.31	0.35	0.57	0.15	0.09
nView D700Z	0.54	0.34	0.28	0.58	0.16	0.07
Epson 715c	0.64	0.35	0.30	0.67	0.15	0.05
Proxima DX1	0.62	0.37	0.33	0.55	0.15	0.07
<i>Max Distance</i>	0.085		0.086		0.028	

nance remains at the same level for many lower inputs, the chromaticity coordinates also remain constant. At very high brightness settings, we observed some nonmonotonicity in the luminance response for the higher input range. As a consequence, the chromaticity coordinates also show some nonmonotonicity at the higher brightness settings. Thus, it is ideal to have the brightness control set so that there is no clipping in the lower input range or nonmonotonicity at higher input ranges. For example, in these illustrations, the ideal setting is between 0.5 and 0.75.

3. *Contrast*: We perform similar experiments for the contrast control. This also affects *all the three channels similarly and simultaneously*. The results are illustrated in Figure 6. Poynton [20] mentions that usually the contrast control changes the gain of the luminance curve. We found the same with the projectors. As the gain increases, the luminance difference becomes significant enough at lower input ranges to push the chromaticity away from the gray chromaticity values toward the chromaticity coordinates of the respective primaries. However, the luminance response starts to show severe nonmonotonicity at higher contrast settings, thus reducing the input range of monotonic behavior. So, the contrast setting should be in the monotonic range to maximally use the available color resolution.

4. *White Balance*: The white balance usually has a brightness and contrast control for each of the three channels separately. We put these in five different settings for our readings. The luminance and the chrominance response change exactly the same way as for the independent brightness and contrast controls, but the change affects *only one channel at a time* instead of affecting all of them similarly. Thus, it controls the proportion of the contribution from each channel to a color that in turn changes the white balance (Figure 7).

E. Interprojector Color Variations

In this section we study how these properties of a projector vary across different projectors.

Figure 8 shows the luminance and color gamut response for the maximum intensity of a single channel for different projectors of *same model* having exactly the same values for all the parameters defined in Section II-D. There is nearly 66% variation in the luminance, while the variation in color gamut is relatively much smaller. Figure 8 also shows the high-resolution chrominance response of a dis-

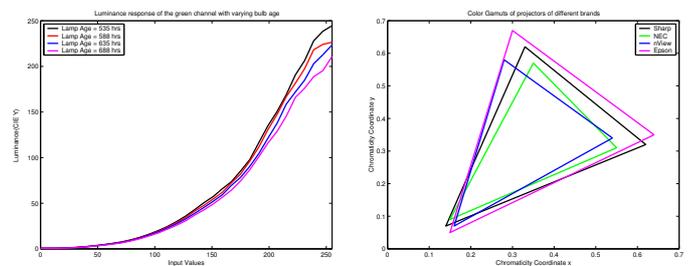


Fig. 9. Left: Luminance response of the green channel at four different bulb ages. Right: Color gamut of projectors of different models.

play wall made of four overlapping projectors of the same model, projecting the same input at all pixels. Projectors of the same model usually use the same brand bulb (which have similar white points) and similar filters, justifying similarity in the color gamut. However, this is not true for the grays of the DLP projectors that use the clear filter, where the chrominance of grays differ significantly across different projectors because of large variation in this clear filter.

But, the color gamut across projectors of different models vary much more, as shown in Table II-D.2 and Figure 9. However, this is relatively much smaller when compared with the luminance variation.

F. Inference

The key observations from experiments and analysis of Sections II-C, II-D, and II-E can be summarized as follows.

1. Within a single projector's field of view, only luminance varies, while chrominance remains almost constant.
2. Chrominance varies negligibly across different projectors of same model, but luminance variation is significant.
3. The variation in chrominance across projectors of different models is very small when compared with the variation in luminance.
4. With the change in various projector parameters like brightness, contrast, zoom, distance, and orientation, only luminance changes, while chrominance remains constant.

III. LUMINANCE-MATCHING ALGORITHM

In almost all cases, projection-based displays are made of projectors of the same model, the chrominance variation across which is very small (Figure 8). Further, humans are much less sensitive to chrominance variation than to luminance variation [21], [22], [23]. Our analysis showed that the chrominance variation across same model projectors is within the perceptual threshold presented in [24]. Aided with this observation, we present a method that matches the luminance response at every pixel of a multiprojector display and achieves photometric uniformity.

A. Algorithm Overview

We describe the algorithm for a single channel. All three channels are treated similarly and independently. The method comprises two steps. The first step is a one-time *calibration step* where a per projector luminance attenuation map (LAM) is generated. In the *image correction step*, this LAM is used to correct any image to be displayed.

A.1 Calibration Step

The calibration step consists of three stages.

1. *Measuring the Luminance Response*: The *luminance response* of any pixel is defined as the variation of luminance with input at that pixel. We measure the luminance response of every pixel of the display with a camera.

2. *Finding the Common Achievable Response*: We find the common response that every pixel of the display is capable of achieving. The goal is to achieve this *common achievable response* at every pixel.

3. *Generating the Luminance Attenuation Map*: We find a luminance attenuation function that transforms the measured luminance response at every pixel to the common achievable response.

For calibration, we assume that every display pixel has a linear luminance response. By linear response we mean that the luminance of black is zero, the maximum luminance occurs for the maximum input, and the luminance response for every other input is a linear interpolation between these two values. This simplifies each of the above three stages. First, in the luminance measurement stage, instead of measuring the luminance response of every input, we can now measure the luminance of only the maximum input. Second, the common achievable response can now be defined as the linear response with minimum luminance range. Third, the luminance attenuation function is just a scaling function that is encoded in the luminance attenuation map. Section III-B shows how we satisfy this assumption in the actual implementation.

1. *Measuring the Luminance Response*: Let us assume that the display D of resolution $W_d \times H_d$ is made up of n projectors each of resolution $W_p \times H_p$. Let us refer to the projectors as $P_i, 0 \leq i < n$. We use a static camera C of resolution $W_c \times H_c$ to measure the luminance of D . Let us denote the luminance response for the *maximum input* of the channel at a display location (x_d, y_d) as $L_d(x_d, y_d)$. The light at (x_d, y_d) can come from one or more projectors. If it comes from more than one projector, then (x_d, y_d) is in the region of the display where multiple projectors overlap. We want to find $L_d(x_d, y_d)$ for all pixels (x_d, y_d) .

Geometric Calibration: First, we perform a geometric calibration that defines the geometric relationships between the projector pixels (x_{P_i}, y_{P_i}) , the camera pixels (x_c, y_c) , and the display pixels (x_d, y_d) . This geometric calibration uses the static camera to take pictures of some known static patterns projected on the display. By processing these pictures, the geometric calibration procedure defines two warps: $T_{P_i \rightarrow C}(x_{P_i}, y_{P_i})$, which maps a pixel (x_{P_i}, y_{P_i}) of projector P_i to the camera pixel (x_c, y_c) , and $T_{C \rightarrow D}(x_c, y_c)$, which maps a camera pixel (x_c, y_c) to a display pixel (x_d, y_d) . The concatenation of these two warps defines $T_{P_i \rightarrow D}(x_{P_i}, y_{P_i})$, which maps a projector pixel (x_{P_i}, y_{P_i}) directly to display pixel (x_d, y_d) . These three warps give us the geometric information we need to find $L_d(x_d, y_d)$.

Data Capture for Luminance Correction: Keeping the camera in the same position, we take the image of each projector P_i projecting the maximum input for the channel. From these images we extract the luminance image for P_i

in the camera coordinate space, denoted by I_i .

Generation of the Luminance Surface: Next we generate the luminance surface $L_{P_i}(x_{P_i}, y_{P_i})$ for every projector P_i . For this, we first transform every projector pixel (x_{P_i}, y_{P_i}) by $T_{P_i \rightarrow C}$ into the camera coordinate space and read the luminance at that transformed pixel from I_i . Hence

$$L_{P_i}(x_{P_i}, y_{P_i}) = I_i(T_{P_i \rightarrow C}(x_{P_i}, y_{P_i})). \quad (3)$$

Once we have the luminance surface L_{P_i} for every projector P_i , we find the contribution of every projector at (x_d, y_d) by the inverse warp of $T_{P_i \rightarrow D}$ denoted by $T_{D \rightarrow P_i}(x_d, y_d)$ and add them up.

$$L_d(x_d, y_d) = \sum_{i=1}^n L_{P_i}(T_{D \rightarrow P_i}(x_d, y_d)) \quad (4)$$

2. *Finding the Common Achievable Response*: The common achievable response is defined as a linear response for which the luminance response for the maximum input is minimum of all $L_d(x_d, y_d)$. This minimum luminance is denoted by L_{min} . Conceptually, this is equivalent to finding a common response that every pixel is capable of achieving. Figure 12 illustrates this.

3. *Generating the Luminance Attenuation Map (LAM)*: The LAM, denoted by $A_d(x_d, y_d)$, is first generated in the display coordinate space and is given by

$$A_d(x_d, y_d) = \frac{L_{min}}{L_d(x_d, y_d)}. \quad (5)$$

This signifies the pixelwise scale factor (less than 1.0) by which L_d should be attenuated to achieve luminance matching.

The next step is to generate the per projector luminance attenuation maps $A_{P_i}(x_{P_i}, y_{P_i})$ from A_d . Since we know the warp $T_{P_i \rightarrow D}$, this is achieved by

$$A_{P_i}(x_{P_i}, y_{P_i}) = A_d(T_{P_i \rightarrow D}(x_{P_i}, y_{P_i})). \quad (6)$$

A.2 Image Correction Step

Once this per projector LAM is generated, it is used to attenuate any image. When an image $M(x_d, y_d)$ of resolution $W_d \times H_d$ is projected on the display wall, the warp $T_{P_i \rightarrow D}$ is used to generate $M_{P_i}(x_{P_i}, y_{P_i})$, which is the part of M that projector P_i should project.

$$M_{P_i}(x_{P_i}, y_{P_i}) = M(T_{P_i \rightarrow D}(x_{P_i}, y_{P_i})) \quad (7)$$

Finally, M_{P_i} is multiplied by A_{P_i} to create the final image for projector P_i , denoted by F_{P_i} .

$$F_{P_i}(x_{P_i}, y_{P_i}) = M_{P_i}(x_{P_i}, y_{P_i}) \times A_{P_i}(x_{P_i}, y_{P_i}) \quad (8)$$

B. Implementation

In this section we describe the implementation of our algorithm. We first implemented it on a wall of resolution 1200×800 made up of a 2×2 array of four projectors. Then, we extended this to a wall of resolution 4500×2000 made of a 5×3 array of fifteen projectors.

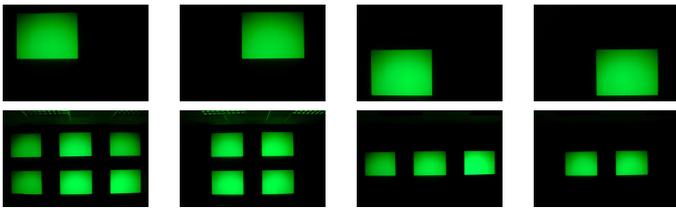


Fig. 10. The pictures taken to generate the luminance attenuation for green channel. Top: For a display made of 2×2 array of 4 projectors. Bottom: For a display made of 5×3 array of 15 projectors.

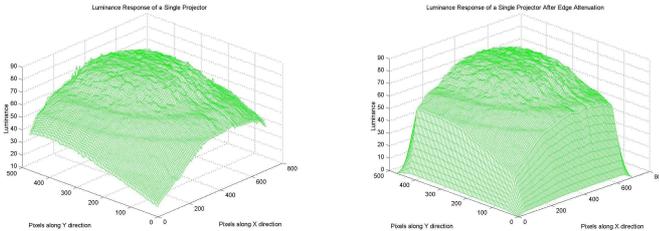


Fig. 11. Left: The luminance surface generated for one projector. Right: The same luminance surface after edge attenuation.

B.1 Luminance Response Measurement

1. Geometric Calibration

We need an accurate geometric calibration algorithm for our photometric calibration. Several geometric calibration algorithms have been designed in the past [10], [25], [11]. Any geometric calibration algorithm that can define accurately the two warps, $T_{P_i \rightarrow C}$ and $T_{C \rightarrow D}$, can be used for our method. For our implementation, we adopt the technique of [18] and use two *cubic nonlinear* warps to define $T_{P_i \rightarrow C}$ and $T_{C \rightarrow D}$. These nonlinear warps include the radial distortion correction for both the camera and the projectors and can be implemented in real time on a traditional graphics pipeline using texture mapping.

2. Data Capture for Luminance Correction

As mentioned in the preceding section, we need to capture images for every projector P_i when it is projecting the maximum luminance for each channel. During this time we turn off all the projectors that overlap with P_i to capture the luminance contribution solely from P_i accurately. To capture the data for all projectors in the display, we need to take a total of four pictures per channel. In each picture alternate projectors are turned on so that none of them overlap with each other. The pictures taken for the two different wall configurations are shown in Figure 10. For our implementation we use a Fujifilm MX-2900 camera. To satisfy the linearity assumption made in the preceding section, we linearize every camera image as explained in Section II.

3. Generation of the Luminance Surface

Generating the Luminance Surface in Camera Coordinate Space: First, we generate the luminance surface in the camera coordinate space corresponding to linearized images generated in the preceding section. For this we use the standard linear transformation usually used to convert RGB colors to YUV space ($Y = 0.299R + 0.587G + 0.114B$).

Generating the Per Projector Luminance Surface: In this step, we generate L_{P_i} for each projector P_i (Figure 11). For every pixel of the projector we find the corresponding camera coordinate using $T_{P_i \rightarrow C}$ and then interpolate the corresponding luminance from the luminance of the four nearest neighbors in the camera coordinate space.

Edge Attenuation: In most projection-based displays, adjacent projectors are overlapped to avoid rigid geometric alignment. However, the transition of luminance from the nonoverlap to the overlap region is very sharp. Theoretically, to reconstruct this edge between the overlap and nonoverlap regions, we would need a camera resolution at least twice the display resolution. Given the resolution of today's display walls, this is a severe restriction.

Instead, we smooth out this sharp transition by attenuating a few pixels at the edge of each projector. This increases the tolerance of our method to inaccuracies in reconstructing this sharp edge. We do this attenuation completely in software. After generating the luminance image for each projector, we attenuate the 40-50 pixels at the edge of the projector using a linear function as shown in Figure 11. We do not need information about the exact geometric location of the overlap regions or the geometric correspondences between projectors in this region but just an approximate idea about the width of the overlap. Further, the width of the attenuation can be changed as long as it is less than the width of the overlap region. Or, a different function can be used (e.g. a cosine ramp).

Adding Them Up: The next step is to add up the luminance surface of each projector in the display coordinate space to generate L_d . For every projector pixel, we use $T_{P_i \rightarrow D}$ to find the corresponding display coordinate and then add the contribution of the luminance to the nearest four display pixels in a bilinear fashion. Figure 12 shows the luminance surface for the 2×2 array of four projectors and the 5×3 array of fifteen projectors thus generated.

B.2 Luminance Attenuation Map Generation

We define the common achievable response as the minimum of L_d designated by L_{min} . Then we generate the luminance attenuation map A_d , in the display coordinate space by dividing L_{min} by L_d (Figure 13). Note that the LAM is dimmer in the overlap regions and near the center of each projector to compensate for the brighter regions of the display luminance surface. Further, because of the large luminance fall-off at the edges of the boundary projectors where there is no overlap, the reduction in dynamic range can be drastic. So, we ignore about 200 pixels in the boundary of the display coordinate system while generating the LAM.

To generate the per projector attenuation map A_{P_i} , for every projector pixel we use $T_{P_i \rightarrow D}$ to convert it to display coordinate space and then interpolate the value of A_d from the nearest four neighbors in bilinear fashion.

Finally, we put in the edge attenuation in the luminance attenuation map for each projector by attenuating the same number of edge pixels in the same way in software as was done while generating the per projector luminance surface.

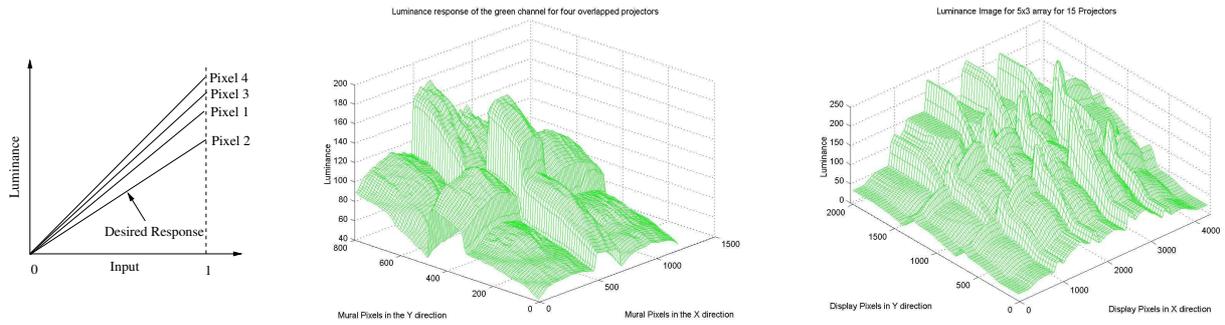


Fig. 12. Left: The common achievable response with four sample pixel response. The response with the least range is the common achievable response. Middle: The display luminance surface for 2×2 array of four projectors. Right: The display luminance surface for 5×3 array of 15 projectors.

Figure 14 shows an example of the per projector LAM. The fifteen-projector wall had larger luminance variation. Hence the attenuation in the fifteen-projector display is higher than that in the four-projector display.

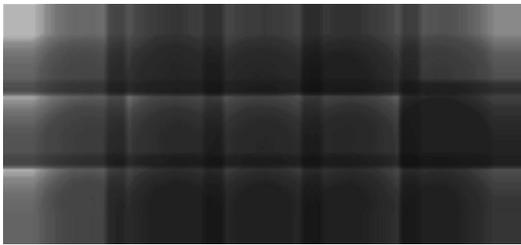


Fig. 13. LAM for a display made of the 5×3 array of 15 projectors.

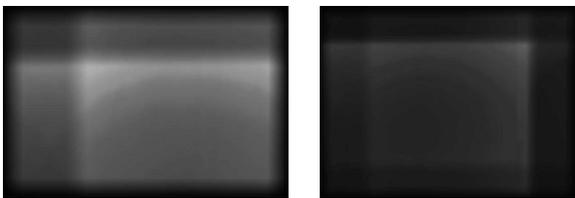


Fig. 14. Left: LAM for a single projector in the four-projector display. Right: LAM for a single projector in the fifteen-projector display.

B.3 Image Correction

The image correction is done in two steps.

1. *Image Attenuation*: The per projector LAM thus generated in the calibration step is first multiplied with the image projected by each projector.
2. *Linearization of Projectors*: Since we have assumed linear response for the projectors in the calibration step, we linearize the projectors. This is done by per projector pre-generated look-up-tables (LUT). It is shown in [19] that the projector nonlinearity response does not vary spatially. So, we measure the per channel nonlinearity in luminance response only at the center of every projector with a photometer. Then, we find a LUT that linearizes this luminance response and use it for all pixels of the projector.

Real-Time Implementation: The LAM can be implemented in real time by using commodity graphics hard-

ware. We use alpha blending to multiply the LAM with the image and 3D color LUT of pixel shaders for the linearization.

C. Results

In this section we present and discuss our results. Figure 15 shows the results on the four-projector display. These images are taken by a digital camera using the same exposure so that they can be compared. The worst test patterns for this algorithm are images with flat test colors. Figure 16 shows our algorithm on one such image. A faint vertical line that can be seen in the images is not the projector boundaries but is the physical crack between the vertical planks that make our display screen.

Figure 16 shows results of the fifteen-projector display. These show two types of artifacts. Some contours are visible, and some of the projector edges are faintly visible. These artifacts are due to insufficient sampling or limited camera dynamic range and will be explained in detail in the next section. The bright spots in the center are due to light leaking through the cracks between the planks making up the display. Because of larger variation in luminance, the attenuation is larger for the fifteen-projector display. Hence, the images of the corrected display are taken at a higher exposure than the images of the uncorrected display.

D. Issues

In this section, we address some implementation issues.

1. *Accuracy of Geometric Calibration*: Even with geometric accuracy about 0.2 pixels (pixels about 2.5 mm), a misalignment of even a couple of pixels in the reconstructed luminance response causes perceived discontinuities without the edge attenuation. The edge attenuation helps us to tolerate greater errors of about 5-6 pixels.
2. *Sampling Density*: The minimum sampling density required to reconstruct the display luminance surface accurately varies across displays. To get an approximate idea, we reconstructed the luminance response of a four-projector region of the display sampled at the ideal sampling density. The frequency content of the luminance of this region is representative of that of a larger display that is potentially made of several such four-projector configu-



Fig. 15. The top row shows the image before correction, and the bottom row shows the image after luminance matching for 2×2 array of four projectors.

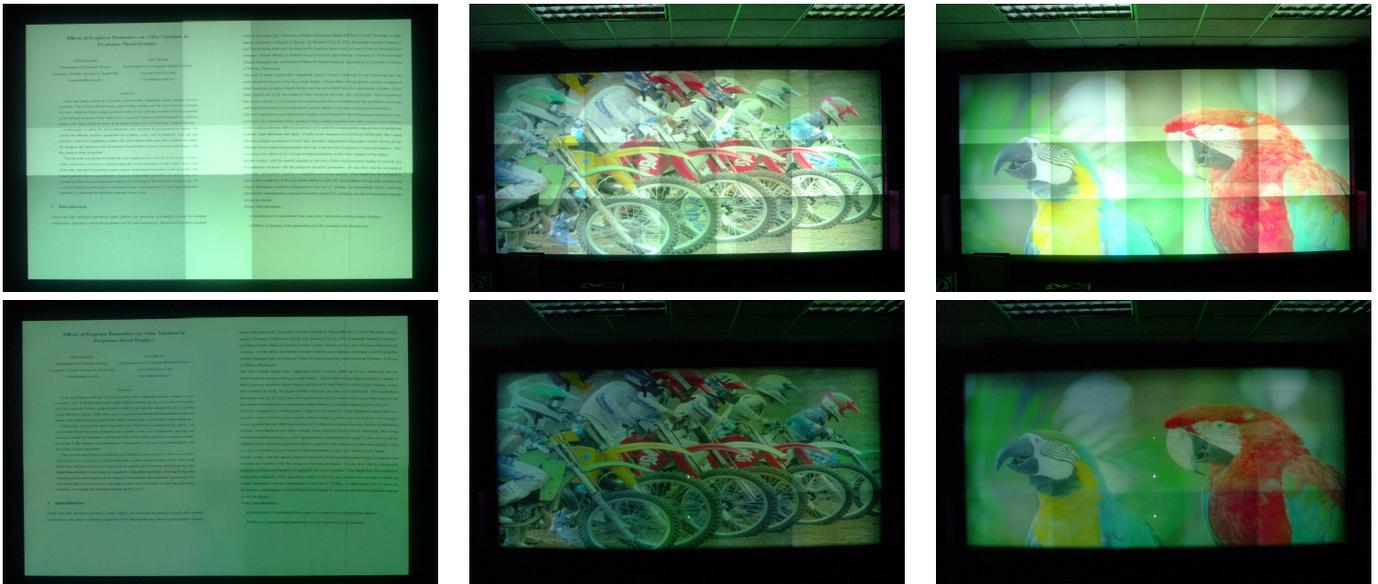


Fig. 16. The top row shows the image before correction, and the bottom row shows the image after luminance matching. Left column: 2×2 array of four projectors. Right two columns: 5×3 array of fifteen projectors

rations. Fourier analysis of this luminance image after edge attenuation shows that the sampling resolution should be about one-fifth of the display resolution.

3. *Dynamic Range of the Calibration Images:* Using the same exposure for all projectors can cause over- or under-saturation leading to contouring artifacts. So, the variation in brightness across projectors is handled by adjusting the camera exposure appropriately.

4. *Black Offset:* Since we assume black offset to be zero, for near-black images faint seams are visible. From our experience, we find that the black offset has less effect on images with high-frequency content.

5. *White Balance:* Since we use a per channel LAM leading to different attenuation for each channel, the grays are

often not retained as grays by the correction. This leads to faint color blotches. So, we use the LAM generated for the green channel for all other channels. The nature of luminance variation being similar across the three channels, the small inaccuracy thus introduced does not show any visible artifacts.

6. *View Dependency:* Though our correction is accurate from one location only, we found our results to be acceptable for a wide angle of about 120 degrees.

IV. CONCLUSION

This paper presents a comprehensive analysis of the nature of the color variation in projection-based displays. This leads to the realization that luminance matching can

achieve acceptable photometric uniformity in such displays. To demonstrate this, we present an algorithm that matches the luminance across all the pixels of the display. This method is practical and automatic and addresses all different types of luminance variation in a unified manner. To the best of our knowledge, this is the first system that achieves photometric uniformity across all display pixels.

However, we believe that this is just the first step toward achieving seamless displays. Following are some of the areas we are working on currently.

1. Matching the response of all pixels to the worst possible response reduces the dynamic range of the display dramatically. We are currently developing algorithms that remove seams while maintaining high dynamic range.
2. We are designing photometric and perceptual metrics that quantify the different display properties that are improved or degraded by different algorithms.
3. Our method is limited by camera resolution and can lead to sampling artifacts for walls of extremely high resolution. So, we are designing scalable algorithms that correct parts of the wall at a time at high resolution and then stitch them together.

ACKNOWLEDGMENTS

We acknowledge Kodak for providing us with high-resolution pictures to test our algorithm. This work was supported in part by the U.S. Department of Energy (DOE), under Contract W-31-109-Eng-38.

REFERENCES

- [1] Aditi Majumder, Zue He, Herman Towles, and Greg Welch, "Achieving color uniformity across multi-projector displays," *Proceedings of IEEE Visualization*, October 2000.
- [2] Maureen C. Stone, "Color balancing experimental projection displays," *9th IS&T/SID Color Imaging Conference*, April 2001.
- [3] Maureen C. Stone, "Color and brightness appearance issues for tiled displays," *IEEE Computer Graphics and Applications*, September/October 2001.
- [4] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stessin, and H. Fuchs, "The office of the future: A unified approach to image based modeling and spatially immersive display," in *Proceedings of ACM Siggraph*, 1998, pp. 168–176.
- [5] Samanta Rudro, Jiannan Zheng, Thomas Funkhouse, Kai Li, and Jaswinder Pal Singh, "Load balancing for multi-projector rendering systems," in *SIGGRAPH/Eurographics Workshop on Graphics Hardware*, August 1999.
- [6] Greg Humphreys, Matthew Eldridge, Ian Buck, Gordon Stoll, Matthew Everett, , and Pat Hanrahan, "Wiregl: A scalable graphics system for clusters," *Proceedings of ACM SIGGRAPH*, 2001.
- [7] Ian Buck, Greg Humphreys, and Pat Hanrahan, "Tracking graphics state for networked rendering," *Proceedings of Eurographics/SIGGRAPH Workshop on Graphics Hardware*, 2000.
- [8] Greg Humphreys, Ian Buck, Matthew Eldridge, and Pat Hanrahan, "Distributed rendering for scalable displays," *Proceedings of IEEE Supercomputing*, 2000.
- [9] G. Humphreys and P. Hanrahan, "A distributed graphics system for large tiled displays," in *Proceedings of IEEE Visualization*, 1999.
- [10] R. Raskar, M.S. Brown, R. Yang, W. Chen, H. Towles, B. Seales, and H. Fuchs, "Multi projector displays using camera based registration," *Proceedings of IEEE Visualization*, 1999.
- [11] Ruigang Yang, David Gotz, Justin Hensley, Herman Towles, and Michael S. Brown, "Pixelflex: A reconfigurable multi-projector display system," *Proceedings of IEEE Visualization*, 2001.
- [12] Greg Humphreys, Mike Houston, Ren Ng, Randall Frank, Sean Ahem, Peter Kirchner, and James Klosowski, "Chromium: A stream processing framework for interactive rendering on clusters," *ACM Transactions of Graphics*, 2002.
- [13] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. Defanti, "Surround-screen projection-based virtual reality : The design and implementation of the CAVE," in *Proceedings of ACM Siggraph*, 1993.
- [14] R. A. Chorley and J. Laylock, "Human factor consideration for the interface between electro-optical display and the human visual system," in *Displays*, 1981, vol. 4.
- [15] Edward J. Giorgianni and Thomas E. Madden, *Digital Color Management: Encoding Solutions*, Addison Wesley, 1998.
- [16] D. H. Brainard, "Calibration of a computer controlled color monitor," *Color Research and Applications*, vol. 14, no. 1, pp. 23–34, Feb 1989.
- [17] Paul E. Debevec and Jitendra Malik, "Recovering high dynamic range radiance maps from photographs," *Proceedings of ACM Siggraph*, pp. 369–378, 1997.
- [18] Mark Hereld, Ivan R. Judson, and Rick Stevens, "Dottytoto: A measurement engine for aligning multi-projector display systems," *Argonne National Laboratory preprint ANL/MCS-P958-0502*, May, 2002.
- [19] Aditi Majumder, "Properties of color variation across multi-projector displays," *Proceedings of SID Eurodisplay*, 2002.
- [20] Charles Poynton, *A Technical Introduction to Digital Video*, John Wiley and Sons, 1996.
- [21] Charles J. Lloyd, "Quantifying edge blend quality: Correlation with observed judgements," *Proceedings of Image Conference*, 2002.
- [22] Scott Daly, "The visual differences predictor: An algorithm for the assessment of image fidelity," *Digital Images and Human Vision*, editor: A. B. Watson, MIT Press, pp. 179–206, 1993.
- [23] Russell L. De Valois and Karen K. De Valois, *Spatial Vision*, Oxford University Press, 1988.
- [24] D. L. MacAdam, "Visual sensitivity to color differences in daylight," *Journal of Optical Society of America*, pp. 247–274, 1942.
- [25] Ramesh Raskar, "Immersive planar displays using roughly aligned projectors," in *Proceedings of IEEE Virtual Reality*, 2000.



and human computer



Aditi Majumder is a Ph.D. candidate at the Department of Computer Science, University of North Carolina at Chapel Hill, and a research associate at the Mathematics and Computer Science Division of Argonne National Laboratory. She received her B.E. in Computer Science and Engineering from Jadavpur University, Calcutta, India, in 1996 and M.S. from the Department of Computer Science, University of North Carolina at Chapel Hill, in 1998. Her research is focused on color seamlessness

interaction in tiled multiprojector displays.

Rick Stevens is director of the Mathematics and Computer Science Division at Argonne National Laboratory (ANL) and professor at the Department of Computer Science at University of Chicago (UC). He is also the director of Argonne's advanced computing initiative targeting the development of petaflops computing systems. He is director of the ANL/UC Computation Institute, a multidisciplinary institute aimed at connecting computing to all areas of inquiry at the university and the laboratory. Recently he has been appointed the project director for the National Science Foundation (NSF) supported TeraGrid project to build the US's most comprehensive open scientific computing infrastructure (linking ANL/UC, National Center for Supercomputing Applications (NCSA), San Diego Supercomputer Center (SDSC), Caltech, and the Pittsburg Supercomputing Center). He also heads the Argonne/Chicago Futures Laboratory a research group he started in 1994 to investigate problems in large-scale scientific visualization and advanced collaboration environments. His group in the Futures Lab has developed the widely deployed Access Grid collaboration system. Prof. Stevens is also interested in the development of innovative tools and techniques that enable computational scientists to solve important large-scale problems effectively on advanced scientific computers. Specifically, his research focuses on three principal areas: advanced collaboration and visualization environments, high-performance computer architectures (including Grids), and computational problems in the life sciences, most recently the computational problems arising in systems biology.