LAM: Luminance Attenuation Map for Photometric Uniformity in Projection Based Displays

Aditi Majumder
Department of Computer Science,
University of North Carolina at Chapel Hill, USA.
majumder@cs.unc.edu

Rick Stevens
Department of Computer Science,
University of Chicago, USA /
Mathematics and Computer Science Division,
Argonne National Laboratories, USA.
stevens@mcs.anl.gov

ABSTRACT

Large-area multi-projector displays show significant spatial variation in color, both within a single projector's field of view and across different projectors. Recent research in this area has shown that the color variation is primarily due to luminance variation. Luminance varies within a single projector's field of view, across different brands of projectors and with the variation in projector parameters. Luminance variation is also introduced by overlap between adjacent projectors. On the other hand, chrominance remains constant throughout a projector's field of view and varies little with the change in projector parameters, especially for projectors of the same brand. Hence, *matching* luminance response of all the pixels of a multi-projector display should help us to achieve photometric uniformity.

In this paper, we present a method to do a per channel per pixel luminance matching. Our method consists of a one-time calibration procedure when a *luminance attenuation map (LAM)* is generated. This LAM is then used to correct any image to achieve photometric uniformity. In the one-time calibration step, we first use a camera to measure the per channel luminance response of a multi-projector display and find the pixel with the most "limited" luminance response. Then, for each projector, we generate a per channel LAM that assigns a weight to every pixel of the projector to scale the luminance response of that pixel to match with the most limited response. This LAM is then used to attenuate any image projected by the projector.

This method can be extended to do the image correction in real time on traditional graphics pipeline by using alpha blending and color look-up-tables. To the best of our knowledge, this is the first effort to match luminance across all the pixels of a multi-projector display. Our results show that luminance matching can indeed achieve photometric uniformity.

VRST'02, November 11–13, 2002, Hong Kong.

Categories and Subject Descriptors

I.3.2 [**Computer Graphics**]: Graphics Systems

General Terms

Algorithms

Keywords

Projection Based Displays, Tiled Displays, Color Calibration, Color Uniformity

1. INTRODUCTION

Large-area, high-resolution multi-projector displays have the potential to change the way we interact with our computing environments. The high resolution and large field of view make them extremely useful for visualizing large scientific models. The compelling sense of presence created by such displays makes them suitable for creating immersive virtual environments for 3D teleconferencing and entertainment purposes. Several such displays exist at Princeton, University of North Carolina at Chapel Hill, University of Minnesota, University of Illinois at Chicago, Stanford, MIT, Fraunhofer Institute (Germany), and U.S. national laboratories such as Lawrence Livermore, Argonne, and Sandia National Laboratories. Recent efforts are directed toward building large displays comprising $40 - 50$ projectors (Sandia National Laboratories and National Center for Supercomputing Applications at University of Illinois at Urbana-Champaign).

Geometric registration and photometric uniformity of the projected image are essential in multi-projector displays to provide the user with the illusion of a single display. Several algorithms perform the geometric registration [10, 11, 12, 15]. But, to the best of our knowledge, no single solution takes care of different photometric variations to achieve photometric uniformity in multi-projector displays. The comments in recent works [1, 4, 5, 6, 11] and our experience have led us to believe that this problem is not trivial and needs to be solved.

The color of multi-projector displays shows significant spatial variation, which can be distracting enough to be the sole cause of breaking the illusion of having a single display. Both *intra (within a single projector's field of view)* and *inter (across different projectors)* projector variation are the cause of such photometric non-uniformity. Further, adjacent projectors are overlapped to avoid rigid geometric alignment at the cost of introducing color variation. Some existing solutions try to reduce the higher brightness in the overlap regions by blending techniques [11] implemented either in software or in hardware. But since this does not account for either intra or inter projector 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 [7, 13, 14] 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.

Current research on analysis of the photometric variation [8, 9] shows that the complexity of this problem may be reduced by the fact that the color variation across a multi-projector display is primarily due to luminance variation. The luminance drops by almost 80% at fringes from the center of a single projector, but the chrominance remains constant. The chrominance of projectors of the same brand is so close that the difference can be ignored for all practical purposes. Even for projectors of different brands, luminance difference is much more significant than the chrominance difference. It has also been shown [8] that luminance changes significantly with the change in projector parameters like position, zoom, brightness, contrast and lamp-age. Further, luminance difference is also introduced by overlapping projectors. Thus, matching the luminance response of all the pixels of a multi-projector display may be sufficient to achieve the desired photometric uniformity.

In this paper we present a method that achieves such a luminance matching, which 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)* that can then be used to correct any image projected by the projector.

1.1 Main Contributions

Following are the main contributions of the paper.

1. To the best of our knowledge, this is the first effort that solves for both intra and inter-projector color variation and the variation introduced by the overlaps. All of these variations are handled by a single algorithm in an automated, unified manner that is completely transparent to the user. In the past, it was assumed that color variation exists only across different projectors but not within a single projector and hence methods were devised [7] to achieve photometric uniformity by accurately measuring the photometric response at only one location per projector. Further, because of the same assumption, a different algorithm was needed to handle the overlaps. As shown by the several studies [8, 9], this assumption is not true.

2. To the best of our knowledge, this is the first effort to achieve photometric uniformity in multi-projector displays that uses a commodity off-the-shelf product (an inexpensive digital camera) to measure the spatial luminance variation across the display. Previous work in this direction [7] used a expensive high-precision radiometer which measures one pixel at a time and takes several minutes. Using such a device is impractical for measuring the response of potentially millions of pixels on the display.

3. The use of a camera to accurately measure the luminance variation across the multi-projector display and the one-time calibration process to generate the LAM make this method easily scalable, practical, and general purpose. Further, this method has the potential to be used in traditional graphics pipelines to achieve this correction in real time.

In Section 2, we give overview of our algorithm. In Section 3, we discuss the implementation of each step of the algorithm in detail. Then we present the results in Section 4. In Section 5, we discuss several pertinent issues that can affect the quality of the results achieved by our algorithm. Finally, in Section 6 we discuss future work.

2. ALGORITHM OVERVIEW

In this section, the algorithm is described 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 is generated. In the second step, this LAM is used to correct any image content.

2.1 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 to 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.

If we assume a linear response for the projectors, then each of these three stages gets simplified. By linear response we mean that the luminance of black is zero, the maximum luminance occurs for the maximum input, and luminance response for every other input is a linear interpolation between these two values. First, the luminance measurement stage is simplified with this assumption because 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. Hence, we assume that every display pixel has a linear luminance response. In Section 3 we show how we satisfy this assumption in the actual implementation.

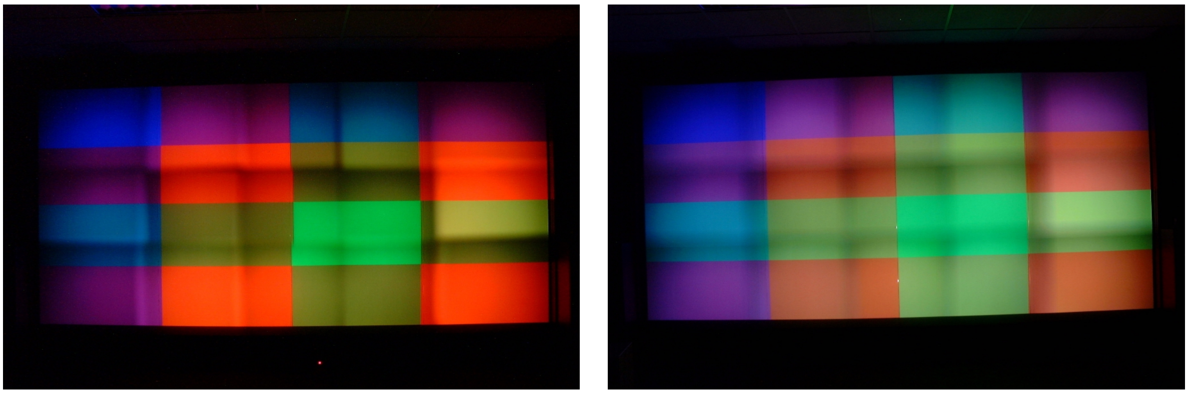

Figure 1: Left: A Display of a 5×3 array of 15 projectors where the overlap regions are blended by using a physical shadow mask on the light path of the projector. Right: Same display with the overlap region blended by a linear ramp in software. For this, one needs to know the exact location of the overlap region. Note that the projector boundaries are noticeable in both cases.

2.1.1 *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}), 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 : $T_{P_i \to C}(x_{P_i}, y_{P_i})$, which maps a pixel (x_{P_i}, y_{P_i}) of projector P_i to the camera pixel (x_c, y_c), and $T_{C \to D}(x_c, y_c)$, which maps a camera pixel (x_c, y_c) to a display pixel (x_d, y_d). The concatenation of these two warps defines $T_{P_i \to D}(x_{P_i}, y_{P_i})$, 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, denoted by I_i, for each projector P_i in the camera coordinate space.

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 $T_{P_i \to C}$ into the camera coordinate space and read the luminance at that transformed pixel from I_i. Hence

$$L_{P_i}(x_{P_i}, y_{P_i}) = I_i(T_{P_i \to C}(x_{P_i}, y_{P_i})) \tag{1}$$

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 $T_{P_i \to D}$ denoted by $T_{D \to P_i}(x_d, y_d)$ and add them up.

$$L_d(x_d, y_d) = \sum_{i=1}^{n} L_{P_i}(T_{D \to P_i}(x_d, y_d)) \tag{2}$$

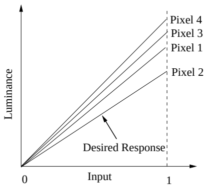

Figure 2: The common achievable response with four sample pixel response. The response with the least range is the common achievable response.

2.1.2 *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 $L_d(x_d, y_d)$ and 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 2 illustrates this.

2.1.3 *Generating the Luminance Attenuation Map*

The LAM, denoted by $A_d(x_d, y_d)$, is first generated in the display coordinate space and is given by

$$A_d(x_d, y_d) = \frac{L_{min}}{L_d(x_d, y_d)} \tag{3}$$

Thus A_d signifies the pixelwise scale factor (less than 1.0) by which L_d should be scaled down 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 $T_{P_i \to D}$, this is achieved by

$$A_{P_i}(x_{P_i}, y_{P_i}) = A_d(T_{P_i \to D}(x_{P_i}, y_{P_i})) \tag{4}$$

2.2 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 $T_{P_i \to D}$ is used to the generate $M_{P_i}(x_{P_i}, y_{P_i})$ which is the part of M that projector P_i should project.

$$M_{P_i}(x_{P_i}, y_{P_i}) = M(T_{P_i \to D}(x_{P_i}, y_{P_i})) \tag{5}$$

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

$$F_{P_i}(x_{P_i}, y_{P_i}) = M_{P_i}(x_{P_i}, y_{P_i}) \times A_{P_i}(x_{P_i}, y_{P_i}) \tag{6}$$

3. IMPLEMENTATION

In this section we will describe how our method is implemented. The implementation is done on two wall configurations. The first one is a wall of resolution 1200×800 made up of 2×2 array of four projectors. Later we extended this to a wall of resolution 4500×2000 made of 5×3 array of fifteen projectors.

3.1 Luminance Response Measurement

This section focuses on the luminance response measurement.

3.1.1 *Geometric Calibration*

We need an accurate geometric calibration algorithm for our photometric calibration. Several geometric calibration algorithms have been designed in the past [10, 12, 15]. Any geometric calibration algorithm that can define accurately the two warps, $T_{P_i \to C}$ and $T_{C \to D}$, can be used for our method. For our implementation, we use two *cubic nonlinear* warps to define $T_{P_i \to C}$ and $T_{C \to D}$. These non-linear warps include the radial distortion correction for both the camera and the projectors and can be implemented in real time on traditional graphics pipeline by using texture mapping. The details of our algorithm are available in [3].

3.1.2 *Data Capture for Luminance Correction*

As mentioned in the preceeding section, we need to capture images for every projector P_i when it is projecting the maximum input 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 1 of the color page.

In the preceeding section, we assumed linear devices for our algorithm. To satisfy this assumption, we find the camera's nonlinear response and linearize it using a color look-up-table. For our implementation we use a Fujifilm MX-2900 camera. We use the method presented in [2] to reconstruct its nonlinear response. This method generates a per channel color look-up-table that linearizes the per channel luminance response of the camera. Every image captured by the camera is linearized using this look-up-table.

3.1.3 *Generating the Luminance Surface*

Generating the luminance response surface for the display requires several steps.

Generating the Luminance Surface in Camera Coordinate Space: First, we find the luminance surface in the camera coordinate space corresponding to linearized images generated in the preceeding section. For this we use the standard linear transformation usually used to convert RGB colors to YUV space given by

$$Y = 0.299R + 0.587G + 0.114B \tag{7}$$

Generating the Per Projector Luminance Surface: In this step, we generate L_{P_i} for each projector P_i. For every pixel of the projector we find the corresponding camera coordinate using $T_{P_i \to C}$ and then interpolate bilinearly the corresponding luminance from the luminance of the four nearest neighbors in the camera coordinate space. An example of the luminance surface thus generated for a projector is shown in Figure 3.

Edge Attenuation: In most projection based displays, adjacent projectors are overlapped to avoid rigid geometric alignment. However, the luminance in the overlap region is much higher than the luminance in the non-overlapped region and this spatial transition is very sharp. Theoretically, to reconstruct this edge between the overlapped and non-overlapped 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 error tolerance to inaccuracies in reconstruction of the luminance surface in regions of sharp transition. We do this attenuation 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. (The width of this attenuation can be changed as long as it is less than the width of the overlap region. Similarly, a different function can be used e.g. a cosine ramp.) Figure 3 shows the luminance after such an edge attenuation. Note that we do not need information about the exact location of the overlap regions for this purpose but just an approximate idea about the width of the overlap so that the attenuation width is less than the width of the overlap. Further, this approach allows us to process the luminance of each projector independently, without explicitly considering geometric correspondences across the projectors.

Adding Them Up: Now, we have got the luminance image for each projector. The next step is to add them all up in the display coordinate space to generate L_d. For every projector pixel, we use $T_{P_i \to D}$ to find the corresponding display coordinate and then add the contribution of the luminance to the nearest four display pixels in a bilinear fashion. The generated luminance surface for the 2×2 array of four projectors and the 5×3 array of fifteen projectors is shown in Figure 4.

3.2 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

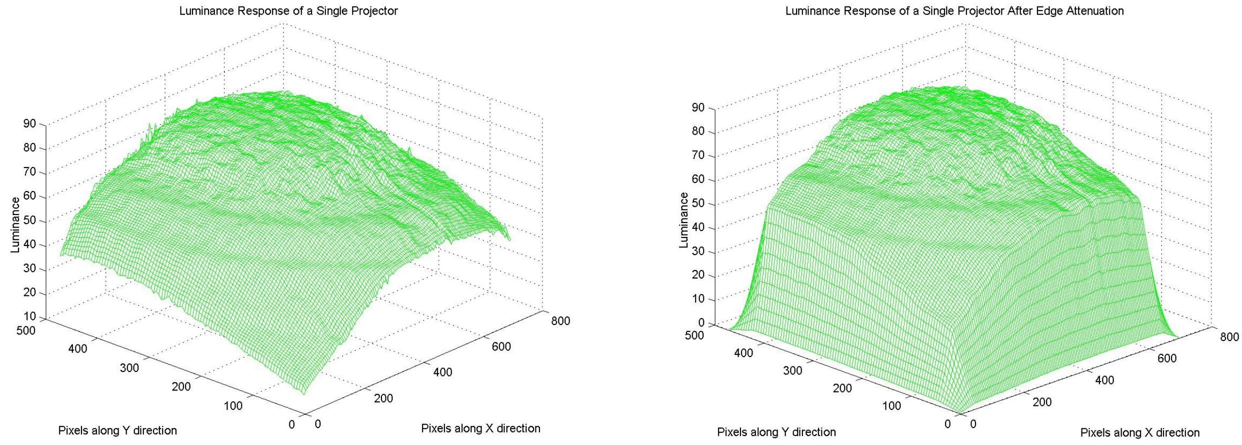

Figure 3: Left: The luminance surface generated for one projector. Right: The same luminance surface after edge attenuation.

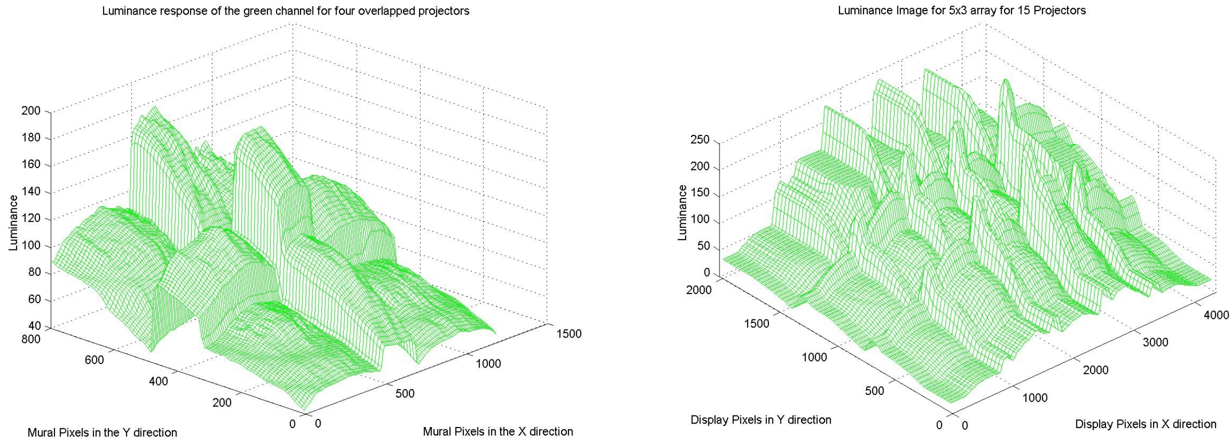

Figure 4: Left: The luminance surface generated for 2×2 array of four projectors. Right: The luminance surface generated for 5×3 array of 15 projectors.

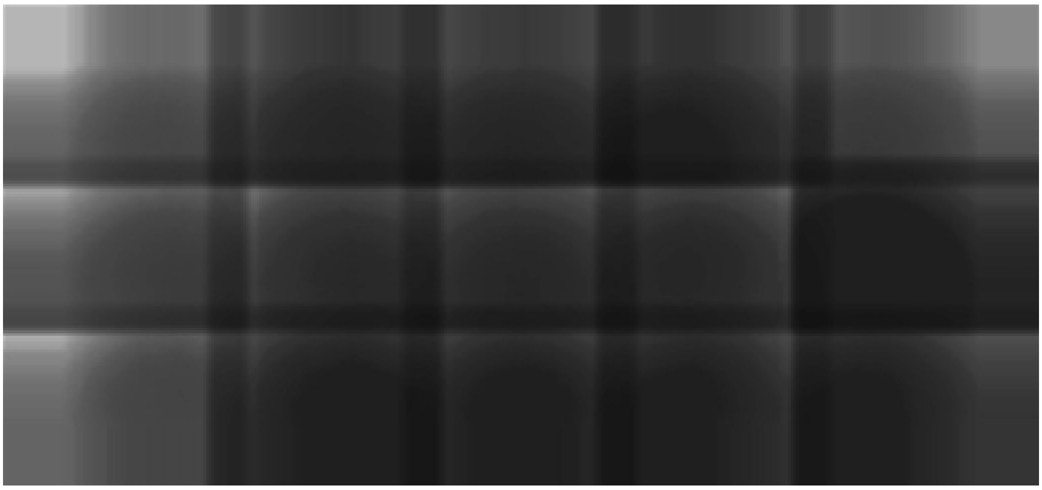

Figure 5: LAM for a display made of the 5×3 array of 15 projectors.

by dividing L_{min} by L_d. This is shown in Figure 5. Notice how the LAM is dimmer to compensate for the brighter regions of the luminance surface in the overlap regions and near the center of each projector. 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 leading to unacceptable picture quality. Hence, 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 bilinearly the value of A_d from the nearest four neighbors.

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 as was done while generating the luminance image in the preceeding section. Figure 6 shows an example LAM for one projector. The fifteen projector wall had larger luminance variation, with some of the projectors having very low luminance response. Hence the attenuation in the fifteen-projector display is higher than that in the four-projector display.

3.3 Image Correction

The image correction is done in two steps.

Image Attenuation: The LAM is multiplied with the image to be rendered. This can be extended to an interactive application using traditional graphics hardware, where the LAM can be used as an alpha mask that is blended with the rendered image.

Linearization of Projectors : Since we have assumed linear response for the projectors, we have to linearize the projectors. This is done by a look-up-table. These per projector look-up-tables are pregenerated. It is shown in [8] that the projector non-linearity response does not vary spatially. Hence, we use a photometer to measure the per channel nonlinear luminance response only at the center of every projector. Then we find a look-up-table that would linearize this luminance response and use this for all pixels of the projector.

4. RESULTS

In this section we present and discuss our results for the four and the fifteen-projector display walls. Figure 7 shows the results on the four projector wall. 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 2 of the color page shows our algorithm on such images. Note that our algorithm achieves uniformity even in such cases. The faint vertical line seen in the images is not the projector boundaries but is the physical crack between the vertical planks that make our display wall.

Figure 3 of the color page show the results for the fifteen-projector display. Two types of artifacts are faintly visible in these results: some contours and some of the projector edges. 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 the 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.

The LAM can be implemented by using the conventional graphics pipeline in real time by alpha blending. For the final linearization in the image correction step, however we need a look-up-table (LUT). Usually off-the-shelf projectors have in-built hardware LUT, which would be ideal for this purpose because it would not incur any extra computational overhead. However, most commercial projectors do not give the user complete access to this hardware LUT. Hence we had to implement this using the software LUT in OpenGL. Unfortunately, this becomes a bottleneck in terms of achieving interactive speeds. We can render a movie using OpenGL at 15 frames per second just with the alpha mask. (This speed is limited by the time required to load the movie and not to render it.) Without the final linearization, the movie however does not look right. But, if we use the OpenGL LUT for this purpose, it takes $2-3$ seconds per frame on nVidia GeForce cards. Currently we are trying to find some projectors that will give us complete access to their hardware LUTs so that we can implement an interactive version of this algorithm.

5. ISSUES

As a result of our work, we have identified several issues that we now comment on.

Accuracy of Geometric Calibration: Our geometric calibration algorithm gives us an accuracy of 0.2 pixels. Each display pixel is about $2.5mm$ is size. Even with this accuracy, however, a misalignment of even a couple of pixels in the reconstructed luminance response can cause perceived discontinuities without the edge attenuation. The edge attenuation alleviates the situation, and we can tolerate greater errors of about $5-6$ pixels in the display space.

Sampling Density: Sampling density decides the accuracy of the reconstruction of the luminance surface for the display. As is clear from the results, having two times the resolution of the display is ideal and would get rid of any sampling artifacts. More important however, is the minimum sampling density required to reconstruct the surface correctly. Obviously, this will be different from wall to wall. But to get an approximate idea, we did the following experiment. We reconstructed the luminance response of a four-projector region of the wall sampled at the ideal sampling density. The frequency content of the luminance of this region is representative of that of a larger display because the larger display is made of several such four-projector configurations. Fourier analysis of this luminance image after edge attenuation showed that the required sampling resolution is about one-fifth of the display resolution. In our fifteen-projector implementation, the wall is sampled at one-third the display resolution, and still we see some artifacts since there may be places in the wall that were not properly represented by the small region we used to decide on the minimum sampling density.

Dynamic Range of the Calibration Images: It is important for the brightness of each projector to be well within the dynamic range of the camera. This can be verified by

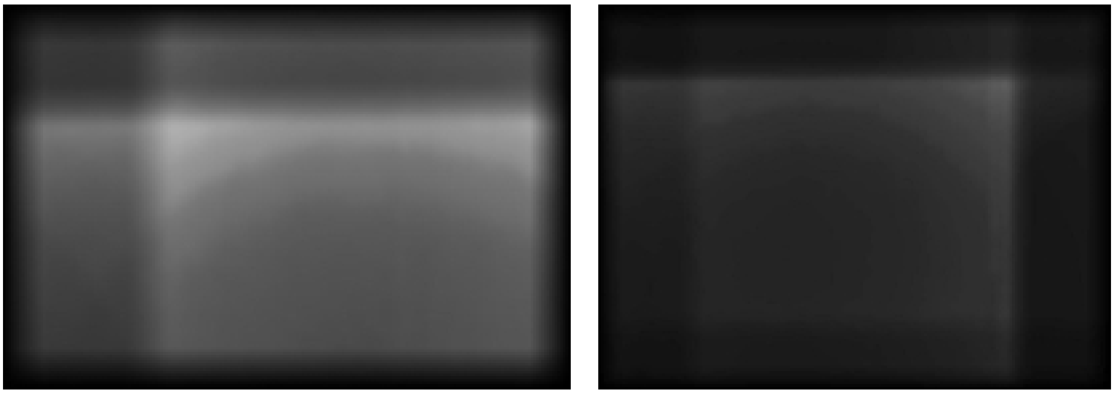

Figure 6: Left: LAM for a single projector in the four-projector display. Right: LAM for a single projector in the fifteen-projector display.

simple under or over-saturation tests of the camera images. In display walls made of many projectors there may be large brightness variation across projectors. In such cases, the camera exposure should be adjusted to accommodate for this variation. This change in exposure should be taken into account by appropriate scaling factors while generating the luminance surface [2]. Using same exposure for all projectors leads to contouring artifacts as seen in the right-most projector in Figure 3 of the color page.

Camera Properties: It is important for the camera not to introduce additional luminance variation beyond that is already present in the wall. Hence, the camera must produce flat fields when it is seeing a flat color. As is mentioned in [2], most cameras satisfy this property at lower aperture settings, especially below $F8$. Our camera had 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.

Black Offset: In our method we assume that black produces zero luminance. This is not true in case of the projectors. Because of several leakages in the light path, the projectors have a non-zero black luminance called the *black offset*. Hence, if the image content is near black, we can see faint seams. From our experience, we find that the black offset has less effect on images with high frequency contents.

White Balance: Our method generates a per channel LAM for every pixel. Since each channel may get attenuated differently, the grays may not be retained as grays when transformed by the LAM. Faint color blotches may therefore appear in the results. Hence, we use the LAM generated for the green channel for all channels. Since the nature of luminance variation is similar across the three channels, the small inaccuracy introduced by this does not show any visible artifacts.

6. CONCLUSION

In summary, we presented a camera based method to achieve photometric uniformity in multi-projector displays. Our one-time calibration procedure generates a luminance attenuation map that is then used to correct any image. The LAM achieves a luminance matching across all the display pixels.

We believe that this is the first step toward achieving photometric uniformity across projection based displays, but much more still needs to be done. Following are some of the areas we are working on.

- Although this method removes the seams, the dynamic range of the display reduces dramatically, since we are matching the response of all the pixels to the response of the worst pixel. This leads to under-utilization of system capabilities, especially in the overlap regions which have higher brightness and range. We are developing algorithms that can remove seams and at the same time make better use of the resources, thus leading to higher dynamic range displays.
- To evaluate the results of algorithms, we are designing a metric that quantifies the different display properties that are improved or degraded by the proposed algorithms.
- As we move towards bigger display walls, the limited camera resolution will be insufficient to sample the luminance surface adequately, leading to sampling artifacts in the corrected images. Hence, we are investigating scalable solutions that can correct parts of the wall at a time and then stitch together the results.
- Our method does not depend on the image content. But, if content of the image is considered as an input to the algorithm, the compression in the input range may be reduced, leading to higher dynamic range images. We are investigating such content based corrections, which may be more suited for canned movies as are used for entertainment purposes.

7. 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.

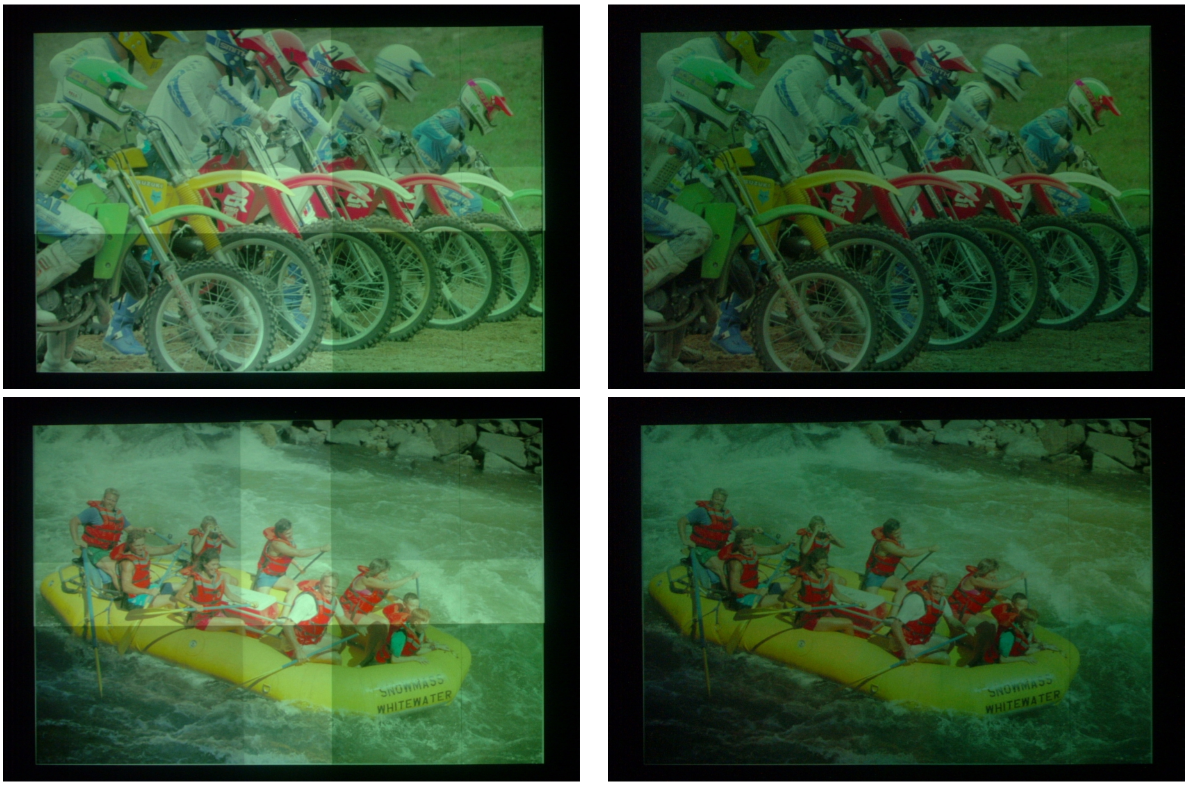

Figure 7: The left column shows the image before correction and the right column shows the image after luminance matching.

8. REFERENCES

[1] I. Buck, G. Humphreys, and P. Hanrahan. Tracking graphics state for networked rendering. *Proceedings of Eurographics/SIGGRAPH Workshop on Graphics Hardware*, pages 87–95, 2000.

[2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. *Proceedings of ACM Siggraph*, pages 369–378, 1997.

[3] M. Hereld, I. R. Judson, and R. Stevens. Dottytoto: A measurement engine for aligning multi-projector display systems. *Argonne National Laboratory preprint ANL/MCS-P958-0502*, May, 2002.

[4] G. Humphreys, I. Buck, M. Eldridge, and P. Hanrahan. Distributed rendering for scalable displays. *Proceedings of IEEE Supercomputing*, 2000.

[5] G. Humphreys, M. Eldridge, I. Buck, G. Stoll, M. Everett, and P. Hanrahan. WireGL: A scalable graphics system for clusters. *Proceedings of ACM SIGGRAPH*, pages 129–140, 2001.

[6] G. Humphreys and P. Hanrahan. A distributed graphics system for large tiled displays. In *Proceedings of IEEE Visualization*, pages 215–223, 1999.

[7] A. Majumder, Z. He, H. Towles, and G. Welch. Achieving color uniformity across multi-projector displays. *Proceedings of IEEE Visualization*, pages 117–124, Oct 2000.

[8] A. Majumder and R. Stevens. Effects of projector parameters on color variation in projection based displays. *Technical Report TR02-022, Department of Computer Science, University of North Carolina at Chapel Hill*, 2002.

[9] A. Majumder. Properties of color variation across multi-projector displays. *To appear in Proceedings of Eurodisplay*, 2002.

[10] R. Raskar, M. Brown, R. Yang, W. Chen, H. Towles, B. Seales, and H. Fuchs. Multi projector displays using camera based registration. *Proceedings of IEEE Visualization*, pages 161–168, 1999.

[11] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image based modeling and spatially immersive display. In *Proceedings of ACM Siggraph*, pages 168–176, 1998.

[12] R. Raskar. Immersive planar displays using roughly aligned projectors. In *Proceedings of IEEE Virtual Reality 2000*, pages 109–116, 1999.

[13] M. C. Stone. Color balancing experimental projection displays. *9th IS&T/SID Color Imaging Conference*, April, 2001.

[14] M. C. Stone. Color and brightness appearance issues for tiled displays. *IEEE Computer Graphics and Applications*, March, 2001.

[15] R. Yang, D. Gotz, J. Hensley, H. Towles, and M. S. Brown. Pixelflex: A reconfigurable multi-projector display system. *Proceedings of IEEE Visualization*, pages 167–174, 2001.