

Dynamic Shadow Removal from Front Projection Displays

Christopher Jaynes*, Stephen Webb, R. Matt Steele, Michael Brown, W. Brent Seales

Metaverse Lab, Dept. of Computer Science
University of Kentucky

ABSTRACT

Front-projection display environments suffer from a fundamental problem: users and other objects in the environment can easily and inadvertently block projectors, creating shadows on the displayed image. We introduce a technique that detects and corrects transient shadows in a multi-projector display. Our approach is to minimize the difference between predicted (generated) and observed (camera) images by continuous modification of the projected image values for each display device. We are unaware of any other technique that directly addresses this problem. Furthermore, we speculate that the general predictive monitoring framework introduced here is capable of addressing more general radiometric consistency problems such as display-surface inter-reflections and the changes in display color and intensity due to projector bulb temperature variation.

Using an automatically-derived relative position of cameras and projectors in the display environment and a straightforward color correction scheme, the system renders an *expected* image for each camera location. Cameras observe the *displayed* image, which is compared with the expected image to detect shadowed regions. These regions are transformed to the appropriate projector frames, where corresponding pixel values are increased. In display regions where more than one projector contributes to the image, shadow regions are eliminated. We demonstrate an implementation of the technique to remove shadows in a multi-projector front projection system.

Keywords: Large-scale display, shadow removal, immersive media, calibration.

1. INTRODUCTION

The large-scale front-projection display, built to operate as a single logical display system from multiple projectors, has emerged in recent years as a compelling and feasible technology. Immersive teleconferencing applications [1], virtual reality environments [2, 3], realistic simulation, and augmented reality [4, 5] have all

begun to explore the multi-projector paradigm. Researchers are systematically addressing the significant technical challenges associated with cooperative, projection-based rendering in somewhat unconstrained environments [6-9]. These challenges include calibration, automatic and dynamic re-calibration, geometric registration, intensity blending and color correction, and synchronization of the rendering pipelines among individual components [7, 8, 10]. Solutions to these challenges enable a loosely configured projection system to be rapidly arranged, calibrated and used as a coherent, scalable display to produce seamless, immersive imagery.

Despite advances, the goal of building a coherent and seamless display from a scalable number of individual projectors still faces significant obstacles. An important



Figure 1: Front projection displays suffer from shadowing artifacts when users occlude projector frustums. This is particularly a problem for projector-based immersive environments that involve large numbers of projectors and encourage users to move within the display environment.

remaining challenge is related to the difficulty in achieving a radiometrically constant synthesis of the underlying projectors. Certain radiometric artifacts have already been addressed by pre-calibration techniques, such as color balancing [11] intensity balancing [12, 13] and blending at overlap regions [9]; others can be controlled through a careful construction of the environment, such as using special display materials with uniform reflectance properties [14], and physically aligning the projectors [3]. But there are perceptually

* Email: jaynes@metaverselab.org, Phone: 859-257-6240

significant radiometric effects that have not yet been addressed.

In front-projection systems, *shadows* are easily created and, though transient, are extremely distracting (Fig. 1). Shadows, regardless of position, provide a perceptual cue that removes the user from the visually immersive experience. While back-projection can be used to avoid shadows, it introduces other problems including space considerations, intensity and sharpness attenuation, and mechanical complexity. Constraining user movement to prevent shadows is not acceptable for interactive display environments that adaptively render a model based on the user’s position. Requiring a user to move in order to avoid shadows forbids particular views of the model to be visualized.

Shadow correction takes place for display regions that are illuminated by at least two projectors and observed by at least one camera. Although several researchers have explored multi-projector environments for arbitrary configurations [2, 8, 9, 15-17], our constraints are reasonable and encountered in most common indoor architectural environments.

The technique requires that at least one camera is able to observe the screen surface at all times. This restricts the placement of the cameras in the display environment (for the experiments presented here a camera is mounted overhead to minimize the chance of occlusion by the user). We believe it is important that projectors can be placed arbitrarily, without regard for the potential for occlusion, so as to maximize the usefulness of the display environment (surface area coverage or resolution). Restricting the position of cameras, in order to support this type of reconfigurable immersive environment is appropriate.

2. CALIBRATION

Accurate calibration of the devices within the display environment is critical to both shadow detection and removal. Predicted images are constructed using the recovered position of the camera with respect to display projectors as well as a straightforward color transfer function. These predicted images can then be directly compared with observed images to detect unexpected artifacts. Detected shadows are warped from the camera frame to each contributing projector’s frame to determine the appropriate projector pixels for adjustment.

Calibration is a two phase process that is performed prior to use of the display system. The relative geometry of both projectors and cameras is recovered. Next, the color differences between each projector and camera are estimated by iteratively projecting different color intensities and measuring each camera response.

2.1 Geometric Calibration

Given a camera and projector pair, calibration determines the transform from pixels in the camera plane to their corresponding positions in the projectors’ frame buffers. Given this transform, regions in shadow, observed in a camera, can then be correctly adjusted in the projected imagery. There have been a number of researchers who have used the controllable nature of a projector and camera pair to recover calibration information in a number of contexts [18] and several different calibration techniques have been explicitly designed for front-projection display environments [12, 16, 19, 20]. In the interest of readability, we present one such calibration technique for the case in which the display surface is planar. The planar assumption is not a requirement, however, and other calibration techniques to derive a pointwise mapping between image and framebuffer pixels could be used [19]. These approaches involve an extra rendering pass to implement the transform, however, and slow overall system performance.

If we assume that the devices observe a plane, the calibration problem becomes a matter of finding the collineation \mathbf{A} such that:

$$\tilde{p}_j = \mathbf{A}p_i \quad \text{Equation 2-1}$$

for all points p_i in the camera and all p_j in the projector. Because \mathbf{A} is a planar projective transform (a collineation in P^2) it can be determined up to an unknown scale factor λ , by four pairs of matching points in general configuration [21, 22]. Iteratively projecting a random point from the projector onto the display surface and observing that point in the camera generates matching points. Each image point center p_i is computed by fitting a 2D Gaussian to observed greyscale response whose variance is related to expected image noise. The resulting center point of the function is then stored with its matching projector pixel p_j . Given at least four random pairs (for a set of degenerate cases see [21]), we compute \mathbf{A} up to an unknown scale factor λ . For the results shown in this paper, \mathbf{A} is computed using 10 matching pairs.

The accuracy of the recovered \mathbf{A} can be measured as a pixel projection error on the projector’s frame buffer for a number of matching points. Specifically, we make calibration error estimates by illuminating the scene with a known projector pixel p , observing its corresponding position in the camera, and then computing a (sub)pixel difference:

$$\mathbf{e} = \sum_i^N \|p - \mathbf{A}\hat{p}\|^2 \quad \text{Equation 2-2}$$

For the results contained in this paper, ϵ is measured by generating 50 points in the projector frame and calculating projection error in the camera as in Equation 2-2. To improve calibration accuracy, we employ a Monte Carlo technique that estimates \mathbf{A} over many trials of randomly generated match points and measures ϵ for each trial. The recovered \mathbf{A} that leads to the smallest ϵ is retained. Experimentation reveals that, for our situation, twenty trials are usually sufficient to recover accurate calibration. Mean re-projection error is reduced to sub-pixel accuracy, typically between 0.4 and 0.6 pixels.

2.2 Color Calibration

In addition to geometric calibration, the mapping between color values for each camera/projector pair must be estimated. Extensive research has been conducted to address problems associated with inter-projector color non-uniformity in multi-projector systems [11, 13, 23]. For the technique presented here, only a rough estimate of the color transfer function between camera and projectors is required. We model this transfer function in order to generate a more accurate predicted image, so that shadow regions can be more robustly detected. We do not model the color differences between projectors in order to render a uniform image as is the goal in [11]. Techniques that model the color differences between projectors to produce a more uniform color image can be directly incorporated with the framework presented here.

The true relationship between a projected color and the image captured by a camera pair is a complex multi-dimensional function including the projector's gamma curve, the camera's gain function, projection surface properties, and wavelength. For simplicity, the three-color channels (Red, Green, Blue) are assumed to be independent, and are calibrated separately by approximating this complex function for each color channel.

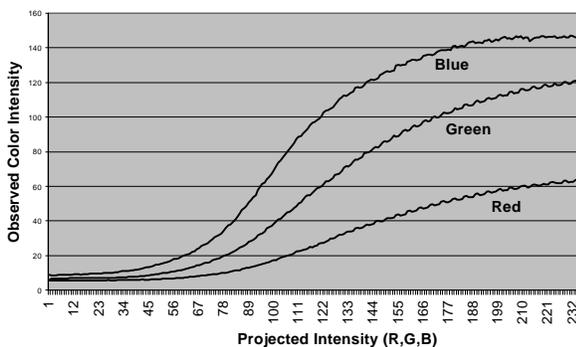


Figure 2: Measured transfer functions for each color channel. Projected intensity values versus observed intensity for each color. Both axes range from 0 to 255.

A given camera \mathbf{C} observes the display surface, while uniform color images of increasing intensity are iteratively projected from projector \mathbf{P} . For each projected color image, the mean color intensity is computed over the corresponding observed image. This is computed for each color channel, holding the other two color values constant at zero. The mean value over ten trials is computed for each color channel. Figure 2 shows the three transfer functions measured by the color calibration process for a single camera/projector pair.

Finally, a transfer function of the form:

$$f_c(x) = \frac{a}{1 + e^{-a(x-b)}} + k \quad \text{Equation 2-3}$$

is fit to the measured values for each channel, where $f_c(x)$ is the color transfer function for color channel C . The four parameters are fit to the measured datapoints using the nonlinear optimization Levenberg-Marquardt technique [24]. Initial estimates for the parameters are provided by the user, who is presented with the color response curves captured during the color calibration process. The mean squared fit error for all three color channels, for the results shown here, was 0.96 (on a 0-255 scale).

The resulting color transfer functions provide a straightforward way to predict how a color in projector space will appear in the camera image. By using this information, predicted color images can be more accurately compared with the observed images to detect shadow regions.

Figure 6a shows a predicted image in a two-projector, one-camera display. Figure 6b shows the image as seen by the camera. Using the transfer functions fit to the data shown in Figure 2, a predicted image based on the relative position of the camera in the environment (geometric calibration) as well as the color mapping between the projector and camera (color calibration) is computed. This image is shown in Figure 6c. Comparison of Figure 6c to Figure 6b versus Figure 6a to Figure 6b reveals that the color corrected predicted image is more accurate than the non-corrected image. Initial studies show that shadow detection is far more robust when comparing color-corrected imagery with observed imagery.

3. SHADOW DETECTION, REMOVAL

Once the display environment has been calibrated, shadow correction should take place continuously and automatically. Correction is a two-phase process. Initially, changes due to unexpected radiometric artifacts on the display surface are detected. Predicted imagery is constructed for a specific camera position and color

transfer function and compared to captured images. A filtering process produces a set of *delta pixels* that must be corrected in the next frame. Given a set of delta pixels, the system then determines corresponding projector pixels that are attenuated or intensified to bring the intensities of the observed image into agreement with the predicted imagery.

An alternative to this approach is a more straightforward, region based technique. Rather than perform a pixelwise compute a new region for which pixels should be either detection of the shadowed regions (represented in the delta image), a bounding box can be fit to detected shadowed regions. This bounding region can then be warped to the appropriate projector framebuffers to illuminated or darkened. This approach has an advantage in that it is straightforward and easy to represent [25]. In a multi-projector environment this more compressed representation can be efficiently transmitted to each rendering device in the display.

The pixelwise radiometric correction approach presented here has significant advantages. In contrast with region based shadow removal in which rectangular regions are either completely on (alpha=1) or off (alpha = 0), the pixel-wise approach operates on individual pixels and accommodates all intensity values by incrementally adjusting the alpha channel values. This more general framework enables future exploration of radiometric consistency problems such as ambient lighting, display inter-reflections, non-uniform inter-projector intensities, and subtle intensity variation due to non-uniform reflectance properties of the display surface.

3.1 Image Processing

It is important that cameras are able to detect radiometric changes on the display surface efficiently and accurately. We assume that the camera is able to observe the display surface without a significant chance of occlusion by the user. For example, the camera may be mounted overhead and oriented down onto a display wall. Ideally, each display camera should observe as large a screen area as possible.

An important aspect of our approach is that radiometric changes are detected directly in camera space. This removes the need for an explicit model of the occluding object. In addition, image-based change detection removes the requirement for full Euclidean calibration of the camera and the position of the occluding object does not need to be known.

A predicted image must be made available to each camera so that it can be compared to the currently observed camera view. In situations where the image is fixed (exploration of a high-resolution still, for example), the predicted image for each display camera can be pre-

computed prior to running the shadow removal algorithm. Predicted imagery may also be updated according to known mouse and keyboard input for scenarios where the system does not correct for underlying surface distortions. This approach is useful if the projected imagery is a simple desktop environment, and is easily extended to interactive display environments.

3.1.1 Camera View Prediction in a Dynamic Environment

In a dynamic display the imagery may change in an unpredictable way (user movement, simulations, video data). The predicted imagery must account for the changing display. This is accomplished by warping the rendered projector framebuffer to camera coordinates (through the recovered homography). Our linear mapping technique for determining both what projectors should render and what cameras should observe based on user position is related to an insightful technique introduced by [1, 9, 15] for view warping under full perspective.

Given a desired image I , each projector in the display environment computes an appropriate framebuffer for projection in order to correctly contribute to the overall image I . This can be accomplished in a number of ways [8, 12, 13, 19, 20, 26] and, for the results shown here, is accomplished in a manner that accounts for display surface geometry, relative projector positions, and the user's headtracked viewpoint in the display. For more details on how the individual projector images are constructed in the multi-projector environment used here, the reader is referred to [10, 23].

Given the recovered collineation between the camera c and a projector p , \mathbf{A}_p^c a predicted image is recovered by warping all projector pixels into the camera frame. For a single projector, the predicted image is given by:

$$\tilde{I} = \mathbf{A}_p^c I \quad \text{Equation 3-1}$$

Because the predicted image is the basis for subsequent modification of projector frame buffer pixels in the display environment, it is important that the image is as accurate as possible. In practice, we super-sample the predicted image by computing the position of each pixel corner on the projector plane to recover an interpolated subpixel estimate of the pixel value at its center. That is, for a pixel $\mathbf{a}=[i \ j \ 0]^T$ on the camera plane, we compute

$$\mathbf{b} = \left[\mathbf{A}_p^c \right]^{-1} \mathbf{a} \pm d\mathbf{a} \quad \text{Equation 3-2}$$

where δa is the effective size of half a pixel in the camera. Each \mathbf{b} is a vertex of a quadrilateral on the image plane of projector.

The correct pixel value $I(i,j)$ for the predicted image I , is estimated as a weighted average of the pixels contained within the quadrilateral, weighed by the percentage of each pixel that is contained in the back-projected region:

$$\tilde{I}(i, j) = \frac{1}{N} \sum_{k=0}^{N-1} p_k \cdot I(p_k) \quad \text{Equation 3-3}$$

where p_k is a projector pixel contained in the quadrilateral and $\lambda(p_k)$ is a weighting factor equal to the total percentage of pixel p_k contained in the quadrilateral. Finally, in the case where more than one projector contributes to a single camera pixel, the mean of all corresponding projector regions, computed using Equation 3-3, is stored in the predicted image.

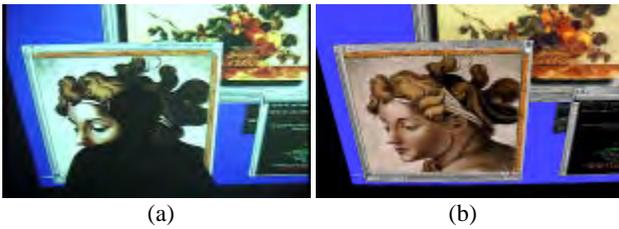


Figure 3: Actual camera view versus predicted image. (a) Actual camera view. (b) Predicted image before color correction using calibrated camera information and the image interpolation technique discussed in this section. Obvious differences include the shadow on the screen surface. Other differences include sensor gain and color differences between the camera and projector (see Text).

Figure 3a shows the image captured by a camera hanging on a ceiling grid, observing a projection display. Figure 3b depicts the predicted image using the technique discussed in this section. In cases where the camera and projector were both mounted upside down (as was the case here), images have been flipped vertically for ease of understanding.

Note that the predicted image differs by the image actually captured by the camera due to sensor nonlinearities, and properties of the projection surface. The transfer functions discovered in the color calibration phase (see Equation 2-3) are applied to the predicted image \tilde{I} to recover a color corrected, predicted image, $\tilde{\tilde{I}}$, that can then be compared directly to the captured imagery. Each color component in the predicted image is adjusted according to:

$$\tilde{\tilde{I}}(i, j, C) = f_c(\tilde{I}(x, y, C)), \quad C = \{R, G, B\} \quad \text{Equation 3-4}$$

Color corrected predicted images are compared to captured imagery by a subtraction of color components to derive two delta images, each representing pixels that are

either too dark and should be intensified (the Δ^+I image) or pixels that are too bright and should be attenuated (Δ^-I).

Each delta image, ΔI , is then filtered with a 3x3 median filter to remove spurious pixels that may emerge due to sensor noise. The size of the filter is directly related to the expected calibration error and expected image noise. Empirical studies have shown that, for the display environment used in this paper, a calibration error of less than a pixel and individual pixel variance can be eliminated with this simple filter. Finally, these delta pixels are converted to an alpha mask in each projector frame. Figure 4a depicts the Δ^+I image for the example shown in Figure 3.

3.2 Alpha Mask Generation

Delta images are directly related to the difference of the observed and predicted imagery for each camera in the display system. Therefore they are computed in the coordinate frame of the camera and must be warped to the reference frame of each projector for correction. This is accomplished using the recovered homography between the camera and projector.

In practice, camera devices and projectors are loosely coupled in the display environment through a communication protocol such as TCP/IP, and therefore recovered delta images must be transmitted as efficiently as possible to projectors for rendering.

In light of the need for transmission efficiency, a single delta image is constructed from the two images using a representation scheme that encodes the sign of a pixel value in the high order bit of a single byte. The remaining seven bits encode the difference (after the lambda gain has been applied) that should be integrated into the next display image. This image represents the alpha values that should be added to and subtracted from the current alpha mask to correct for observed differences. This delta image is then run-length encoded to improve transmission rates and reduce overhead bandwidth required by the radiometric correction algorithm. Due to the typical structure of a difference image (see Figure 4a), the encoded image is typically reduced from approximately 500k bytes to less than 3k bytes. The encoded mask is multicast to all rendering client projectors that have an overlapping field of view with the camera that has detected the radiometric variation.

Each rendering client that receives a delta image, ΔI , first decodes the image and then warps the resulting alpha mask based on the relative position of the projector to the camera to recover the delta image in the projector coordinate frame, ΔI^p . Because the homography between the projector and camera has been recovered in

the calibration phase, this warping is straightforward and is implemented as:

$$\Delta I^P = [\mathbf{A}_p^c]^{-1} \Delta I \quad \text{Equation 3-5}$$

Once a delta image has been aligned to a projector, an appropriate alpha mask is computed as follows:

$$\mathbf{a}^P(i, j) = \mathbf{r} \left[\Delta^+ I^P(i, j) - \Delta^- I^P(i, j) \right] \quad \text{Equation 3-6}$$

where ρ is the maximum allowed intensity change between any two frames and is used to avoid rapid fluctuations on the display surface. For the experiments presented here a ρ value of 25 was used. Although it may seem important to change the projected imagery as quickly as possible, sensor noise may lead to over-correction for shadowed regions. Potentially, this can result in the system iteratively over- and under-correcting the shadowed region. Due to minor calibration error, this feedback may propagate the region of fluctuation on the display. The ρ parameter acts as a dampener to avoid this situation.

The alpha blending process takes into account whether incoming alpha values should be added or subtracted from the alpha channel currently being projected. Figure 4a shows the positive difference image, ΔI^+ . Figure 4b shows the resulting alpha mask, applied to a second projector for correction.

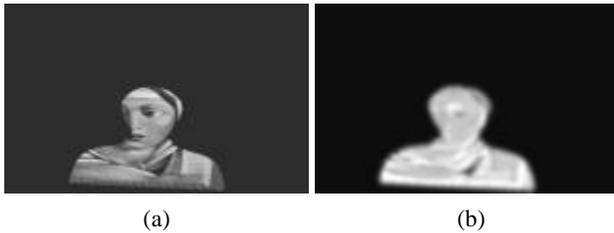


Figure 4: Difference image and resulting alpha mask from the shadow event depicted in Figure 3. (a) Pixels detected in the $\Delta^+ I$ image are the basis for brightening regions in display projectors. (b) Alpha mask after warping the difference image into the display projector’s frame buffer.

4. EXPERIMENTAL RESULTS

The radiometric correction algorithm has been integrated with a front-projection research display environment. The display is composed of four high-resolution projectors, two infrared head-tracking units and two digital video cameras. Each projector is connected to a single PC that contains a commodity, OpenGL-compliant graphics accelerator, and a standard network interface. The cameras are connected to a single PC that contains two frame-grabbers and a network interface. The PC

rendering clients are connected via a 100Mb network hub. The display runs on the Linux operating system.

4.1 Calibration Accuracy

The system was calibrated using the techniques described above and then used to demonstrate the shadow removal technique. In order to estimate the accuracy of our calibration approach, the mean error contained in the recovered homographies between each device and all other devices was estimated. This was accomplished by selecting 10 matching pairs for all devices with a view frustum overlap. Using these matching pairs, the mean pixel error for a particular device was computed by projecting known match points from other devices to the image plane and measuring a pixel disparity. Although we do not calibrate to a Euclidean coordinate system, metric errors on the screen can be estimated by back-projecting a line of known pixel length for each device and measuring the pixel sample distance on the screen for each device. The mean pixel error can then be multiplied by this scale to arrive at an estimate for calibration error in millimeters.

Table 1 reports these errors for all rendering clients in the display environment.

Rendering Client	Mean Pixel Error	Mean Screen Error (mm)
Projector 1	0.583	1.23
Projector 2	0.603	1.61
Projector 3	0.616	1.64
Projector 4	0.664	1.72
Camera A	0.782	1.02
Camera B	0.793	1.18

Table 1: Calibration accuracy. Mean pixel error was measured by reprojecting 10 known points in all other devices to rendering clients through recovered homography and measuring pixel disparity. Screen error was measured by back-projecting error to display surface for each device independently.

4.2 Shadow Removal Results

We have tested the system under a number of different conditions for a number of different applications. These include still objects placed in the path of a single projector, a person moving in front of the display, an object mounted to a controllable pan-tilt unit that is casting a shadow, and the casting of other light sources on the display (such as a flashlight) during its use. Due to space constraints, we only present the system being used by a single user who occludes a projector.

Figure 7 shows a user in a front-projection display environment. Clearly, both the subject’s body and arm are casting on a shadow onto the image. The shadow precludes the user from interacting with the imagery in a

meaningful manner. Note also that an object (a stepladder) partially occludes the lowermost portion of the screen surface.

The system builds a predicted image and detects changes on the screen surface as they occur. Each incoming frame is then compared with the predicted image and differences are translated into modifications of the alpha mask that will correct the display. Figure 5a shows the ΔI^+ image for the occlusion event depicted in Figure 7. The subtraction image is empty in this case and is not shown. Figure 5b shows the ΔI^P image for a projector that is currently not occluded by the user. This image is then directly blended with that projector's current alpha mask to attenuate or intensify the image for radiometric correction.

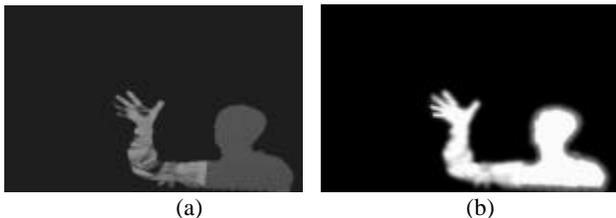


Figure 5: Delta pixels used to correct for observed shadow changes. (a) Computed intensity differences based comparison of predicted versus observed image. (b) Alpha pixels warped to corresponding pixels in the projector's framebuffer.

The newly blended alpha mask is rendered in the projector in concert with all other rendering devices on the next synchronized rendering pass. Typically, the system requires 3-4 frames to converge to a blended image. However, in an occlusion event such as the one depicted here, the algorithm will significantly brighten the shadowed region in the first frame after it is detected. Figure 8 shows the result of blending based on the generated delta images from Figure 5. The information, previously lost in shadow, is now available to the user for inspection despite significant occlusion.

A close-up view of the corrected image is shown in Figure 9. Due to reprojection error as well as inter-frame latency between the detected position of the shadow and the update time, a small boundary between corrected shadows and blended image may exist as is seen in Figure 9. We have measured the border on the display surface by halting the system and inspecting the observable gaps. The widest point on the surface is approximately 11.5mm. Although a shadow boundary is visible and color differences between the different projectors make the shadowed region apparent, the information previously lost to the user is now present. Blending techniques, commonly used to remove the artifacts between adjacent projectors in multi-projector systems such as [9] can be easily modified to smooth the observable shadow boundary.

5. CONCLUSIONS

We have shown that the framework introduced in this paper can be used to remove the appearance of shadows in front projection displays. Our technique does not depend on discovering a complex radiometric model of the scene, nor is it necessary to recover geometric or radiometric models of the objects or phenomena in the display area. Instead, the problem is solved in screen space using cameras, which observe the display and compare the observed image to an expected one. The primary constraint is that screen points are illuminated by more than one projector, and that the relative geometric relationships between cameras and projectors be known. Although the current implementation of the system is not at interactive rates. The algorithm can be implemented in real-time with appropriate rendering hardware. A similar algorithm, that does not require pixelwise updates has been implemented and runs at interactive speeds [25].

We are exploring the use of video cards that provide three alpha channels, one for each color band. This capability will allow correction for color band intensity with control over individual color bands. A second area of exploration is with algorithms for baseline color correction via observed color test patterns. We will also extend the framework to more general cases of radiometric inconsistencies such as dynamic ambient illumination, surface inter-reflection and non-uniformities due to projector and display surface variances.

6. REFERENCES

1. Chen, W., *et al.* *Toward a Compelling Sensation of Telepresence: Demonstrating a portal to a distant (static) office.* in *11th Ann. IEEE Visualization Conference (Vis) 2000.* 2000.
2. Schaufler, G., T. Mazuryk, and D. Schmalstieg. *High Fidelity for Immersive Displays.* in *Short paper, ACM SIGCHI'96 conference companion.* 1996.
3. Cruz-Neira, C., D. Sandin, and T. DeFanti. *Surround-screen Projection-based Virtual Reality: The Design and Implementation of the CAVE.* in *SIGGRAPH'93.* 1993.
4. Ramesh, R., G. Welch, and F. Fuchs. *Spatially Augmented Reality.* in *Augmented Reality; Placing Artificial Objects in Real Scenes—Proceedings of the First IEEE Workshop on Augmented Reality (IWAR'98).* 1998. San Francisco, CA, USA.

5. Underkoffler, J., B. Ullmer, and H. Ishi. *Emancipated Pixels: Real-World Graphics In The Luminous Room*. in *SIGGRAPH'99*. 1999.
6. Herald, M., I.R. Judson, and R.L. Stevens, *Introduciton to Building Projector-Based Tiled Display Systems*. IEEE Computer Graphics and Applications, 2000(July/August 2000): p. 22-28.
7. Samanta, R., *et al*. *Load Balancing for Multi-Projector Rendering Systems*. in *SIGGRAPH/Eurographics Workshop on Graphics Hardware*. 1999.
8. Humphreys, G. and P. Hanrahan. *A Distributed Graphics System for Large Tiled Displays*. in *IEEE Visualization 1999*. 1999.
9. Raskar, R., G. Welch, and H. Fuchs. *Seamless Projection Overlaps Using Image Warping and Intensity Blending*. in *Fourth International Conference on Virtual Systems and Multimedia*. 1998.
10. Webb, S., *et al.*, *Synchronization with Low Cost Networks in a Collaborative Rendering Environment*, .Technical Report 2000, University of Kentucky.
11. Majumder, A., *et al*. *Achieving Color Uniformity Across Multi-Projector Displays*. in *11th Ann. IEEE Visualization Conference (Vis) 2000*. 2000.
12. Surati, R., *Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays.*, 1999, Massachusetts Institute of Technology.
13. Hereld, M., *et al.*, *Developing Tiled Projection Display Systems*, 2000, Mathamatics and Computer Science Division, Argonne National Laboratory.
14. Powerwall, <http://www.lcse.umn.edu/research/powerwall/powerwall.html> .
15. Raskar, R., *et al*. *Efficient Image Generation for Multiprojector and Multisurface Displays*. in *Proceedings of the Eurographics Workshop*. 1998. Vienna, Austria.: Addison-Wesley.
16. Seales, W.B., G. Welch, and C. Jaynes. *Real-Time Depth Warping for 3-D Scene Reconstruction*. in *IEEE Aerospace*. 1998. Boulder, CO.
17. Schikore, D., *et al.*, *High-Resolution Multiprojector Display Walls*. IEEE Computer Graphics and Applications, 2000. **20**(4): p. 38-44.
18. De Piero, F. and M. Trivedi, *3D Computer Vision using Structured Light: Design, Calibration, and Implementation Issues*. *Advanced in Computers*, 1996. **43**: p. 243-278.
19. Raskar, R., *et al*. *Multiprojector Displays using Camera-based Registration*. in *IEEE Visualization*. 1999. San Francisco, CA.
20. Chen, Y., *et al.*, *Automatic Alignment of High-Resolution Multi-Projector Displays Using an Un-Calibrated Camera*, in *TR--618-00*. 2000, Dept. of Computer Science, Princeton University.
21. Faugeras, O., *Three-Dimensional Computer Vision*. 1993, London, England: The MIT Press.
22. Kanatini, *Computational Projective Geometry*. CVGIP, 1991. **54**(3).
23. Jaynes, C., W. Stephen, and B. Seales, *Radiometric Correction in the CoRE Display*, 2000, Technical Report TR313-01, University of Kentucky, Dept. of Computer Science.
24. Press, W.H., *et al.*, *Numerical Recipes*. 1988, Cambridge, MA: Cambridge University Press.
25. Webb, S., C. Jaynes, and B. Seales, *Efficient Region-Based Shadow Correction in Multi-Projector Displays*, 2001, Technical Report, University of Kentucky, Dept. of Computer Science.
26. Raskar, R., *Oblique Projector Rendering on Planar Surfaces for a Tracked User.*, 1997, Technical Report, University of North Carolina at Chapel Hill: Chapel Hill, North Carolina.

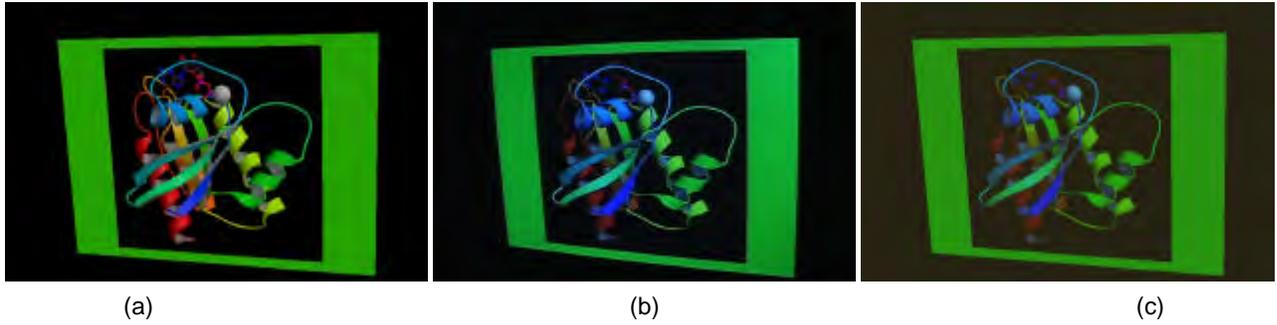


Figure 6: Comparison of predicted image with and without color correction to the observed image. (a) Predicted image without color correction. (b) Image captured by camera while the display environment is in use. (c) Predicted image after three channel transfer functions have been applied. Predicted image is closer in color-space to the captured image than the predicted image without. Predicted imagery is the basis for detection of significant radiometric anomalies such as shadows.



Figure 7: User occludes a region of the desktop while using the large-scale visualization system.



Figure 8: View of the display surface with the radiometric uniformity algorithm running.



Figure 9: Close-up of the corrected image. Other than subjective analysis of the results, accuracy can be analyzed through direct measurement of remaining shadows on the display surface (see Text).