Increasing the Usability of Virtual Rear Projection Displays

Jay Summet, Mathew Flagg, James M. Rehg, Gregory M. Corso, Gregory D. Abowd

11th July 2003

Abstract

Rear projection is often preferred over front projection for creating large-scale interactive displays because it does not suffer from shadows. However, rear projection is not always feasible due to space and cost reasons. Recently, Virtual Rear Projection (VRP) techniques have been developed which coordinate the outputs of multiple overlapping projectors to compensate for shadows and occlusions. We present the results of a user study which compares VRP approaches to traditional rear and front projection. One conclusion from this study is that active compensation schemes are justified by user preference. We also present new results for configuring projected displays using markers, which increases their usability in large interactive displays.

1 Introduction

Virtual Rear Projection (VRP) is an exciting new technology for large-scale interactive displays. By overlapping multiple aligned projectors and using visual feedback, it is possible to eliminate shadows cast by occluders[10, 4] and suppress light cast on the occluders[11]. These features are ideally suited for large interactive displays where users approach the display surface. We refer to this technology as *virtual* rear projection because it provides some of the benefits of rear projection displays using inexpensive, flexible front projection technology. Figure 1 illustrates the basic concept. In contrast to VRP, conventional rear projection is costly from a space, display material, and installation standpoint. In addition, the resulting display area is constrained by the limited extent of the installed projection surface.

Recent work on electronic whiteboards[5], digital tape drawing[1], and focus plus context displays[2] have demonstrated a variety of applications for large interactive displays. We are particularly interested in deploying large VRP displays in a classroom setting. While our work is based on coordinating the outputs of multiple conventional projectors, there has been much recent interest in customized projector systems. The Everywhere Displays projector[7] allows interactive displays in the foot to yard scale to be front projected onto arbitrary planar surfaces, greatly increasing their ubiquity. Recently, Raskar demonstrated a portable projector-camera system to augment reality[8].

In this paper, we describe two extensions of VRP technology which are designed to enhance its usability for interactive displays. The first is a user study which benchmarks the effects of occlusions and shadows on end users and assesses the relative value of several different VRP approaches. This study is the first step in an effort to quantify the benefit of sophisticated compensation schemes under realistic application conditions. We believe this is the first user-study conducted for VRP technologies.

The second extension of VRP technology described in this paper is a real-time calibration technique which gives greater flexibility and ease in defining multiple display surfaces. Our technique makes it possible to easily co-opt any existing wall area into multiple displays simply by placing fiducial markers and pointing a projector-camera system towards them. In particular, we describe a simple form of steerable display which can be constructed with this approach.

2 Study Design

The study described in this paper is designed to: 1) Determine the extent to which shadows on a front projected

Figure 1: (a) Plan view of VRP system. (b) Passive VRP approach using "half shadows". Active VRP approaches include (c) Shadow elimination and (d) Occluder light suppression.

surface affect user task performance. 2) Investigate user strategies for coping with imperfect display technology (which allow occlusions). 3) Evaluate **Warped Front Projection (WFP)** and **Passive Virtual Rear Projection (PVRP)** in comparison to standard **Front Projection (FP)** and true **Rear Projection (RP)** in terms of human performance and preference.

WFP and PVRP are simple compensation techniques that do not require real-time feedback control of multiple projector outputs. Measuring their effectiveness provides a benchmark which more sophisticated active compensation schemes can be compared against. One conclusion from this study is that active compensation schemes are justified by user preference.

The study compares the four conditions listed below for a *single user* working with a large scale interactive surface. Participants were asked to perform interactive tasks on a SmartBoard which utilized a contact sensitive film (touch screen) on the display surface for input. Our study presented participants with four counterbalanced conditions**: FP, WFP, PVRP, RP,** which we describe in the following list:

- *Front Projection (FP)* - A single front projector is mounted along the normal axis of the screen. Users standing between the projector and the screen will produce shadows on the screen. This is a setup similar to most ceiling mounted projectors in conference rooms.

- *Warped Front Projection (WFP)* - A single front projector is mounted off of the normal axis of the projection screen, in an attempt to minimize occlusion of the beam by the user. The output is warped to provide a corrected display on the screen. Examples are new projectors with on-board warping functions, such as used by the 3M IdeaBoard[3], or the Everywhere Displays Projector[7]. Additionally, the latest version of the nVidia video card drivers includes a "keystoning" function which allows any Windows computer to project a warped display.

- *Passive Virtual Rear Projection (PVRP)* - Two front projectors are mounted on opposite sides of the normal axis to redundantly illuminate the screen. Output from each projector is warped (as with WFP) to correctly overlap on the display screen. This reduces the number, size and frequency of occlusions.

- *Rear Projection (RP)* - Using a single projector mounted behind the screen, so that it is not possible to occlude the projection beam or cause shadows.

2.1 Equipment Setup

Care was taken to adjust the output of all projectors so that the intensity on the screen was equal between the different conditions (as measured by a Sekonic Twinmate L-208 light meter). For all conditions the output resolution was adjusted to provide an apparent resolution of 512x512, covering the entire SmartBoard screen, which measures 58" (1.47m) diagonally.

For the front projection conditions (FP,WFP,PVRP) three matched projectors were mounted 7'1" (2.16m) high on a uni-strut beam 10' (3.05m) from the SmartBoard. The rear projection (RP) condition used a projector mounted behind the SmartBoard screen. The projector used for WFP was mounted to the user's right (all participants were right handed) when facing the SmartBoard, 27 degrees off-axis. The pair of projectors used for the

PVRP condition had 48 degrees of angular separation as measured from the screen.

Two video cameras were used to document each session. One camera was mounted behind the SmartBoard screen and was used to measure occlusions caused by the user in the front projection cases (FP,WFP,PVRP), while the other camera recorded the participants' interaction with the display surface.

2.2 Study Participants and Tasks

Our study participants were seventeen (17) college students, 9 males and 8 females, mean age of 21.3 (σ=1.77), from the School of Psychology's experimental pool. Participants were selected to be right-handed, and used their right hand exclusively for interacting with the screen. A photographic image, used to evaluate subjective image quality, and three tasks were presented to the participants. These tasks (especially the second and third) exercise the basic operations (searching, selecting, dragging and tracing) that a user performs with an interactive surface, and are the low level operations needed to perform such UI interactions as button pushing, slider movement, icon dragging, etc. Although they do not directly simulate the use of real applications, we feel that the tasks are relevant for many standard UI interactions and hence, many applications.

Accurate Selection Task (Crosses Task) - Twenty crosses were displayed in a grid over the display surface. The user was instructed to tap as close to the center of each cross as possible, taking as much time as necessary. Accuracy measurements (X and Y offset from the actual center) were made for each cross.

Fast Search, Selection, and Dragging Task (Box Task) - Boxes with 2" sides appeared pseudo-randomly in one of 8 positions around the perimeter of the screen (Figure 2), while a 4" target was placed in the center of the screen. The user was instructed to drag each box into the target. The user moved eighty (80) boxes (ten boxes from each of the eight positions) for each condition.

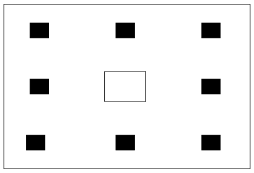

Figure 2: Target and box starting positions.

For each box, the search/select (acquire) time, drag time, and total time were recorded, as well as the number of drags/touches needed to move it into the target. For analysis of the three front projection conditions (FP,WFP,PVRP), data from the video camera behind the SmartBoard was used to determine if the box was initially visible or occluded. A box which was in a half-shadow (in the PVRP condition), and visible with a lower level of contrast, was considered to be unoccluded.

Fast Tracing Task (Spiral Task) - An Archimedes' spiral with three revolutions ($\theta = 6\pi$), designed to test non-linear dragging as an approximation to activities such as tracing and writing, was presented to the participants who were instructed to trace it as quickly as possible. While the user traced sufficiently close to the spiral, their finger would erase it. If their path deviated significantly from the spiral it would cease to respond (erase) and they would have to re-trace from their point of deviation. The error metric allowed for fast tracing, but was strict enough to discourage wild gesturing.

3 Results

Tables 1 & 2 summarize our results discussed in the following sections. Results from the subjective measures (Table 1) indicate significant differences for all four conditions. *[Image Quality: F(3,48) = 9.755, p < 0.001; Preference: F(3,48) = 20.812, p < 0.001; Acceptance: F(2.156,34.5) = 17.366, p < 0.001]*[1]

3.1 Subjective Measures: Image Quality, Preference & Acceptance

Image Quality - As expected, rear projection had the highest reported image quality ("*How would you rate the image quality of the display technology? [Poor Quality =*

---

[1] We conducted a repeated-measures ANOVA to analyze our data. To correct for a potential violation of the sphericity assumption in the acceptance case we applied a Greenhouse-Geisser correction. Post-hoc analysis used paired-samples t-tests. Values of p lower than 0.05 were deemed significant.

Condition	**Image Quality**	**Preference**	**Acceptance**
Front Projection (FP)	4.52	3.35	3.82
Warped Front Projection (WFP)	3.29	*3.18*	*3.47*
Virtual Rear Projection (VRP)	3.70	***4.65***	*4.88*
Rear Projection (RP)	**5.88**	**6.18**	**6.47**

Table 1: Mean subjective measures from 7 point scales. RP scores (in **bold**) are significant when compared to all other conditions ($p<0.05$). User preference of PVRP is also significant. The scores of WFP and PVRP (in *italics*) are significant in relation to each other in the user preference and acceptance categories. The other scores report trends in the data that does not fall under the $p<0.05$ significance criteria.

Condition	Box Aquire Time (sec.)	Crosses Error	Spiral Time (sec.)
Front Projection (FP)	1.25 (0.49)	0.0074 (0.0121)	13.75 (4.10)
Warped Front Projection (WFP)	1.12 (0.26)	0.0082 (0.0033)	13.15 (4.00)
Virtual Rear Projection (VRP)	1.15 (0.28)	0.0084 (0.0088)	13.06 (3.90)
Rear Projection (RP)	1.07 (0.23)	0.0081 (0.0183)	12.27 (3.81)

Table 2: Quantitative measures - Mean (Standard Deviation)

1 2 3 4 5 6 7 = Excellent Quality]”). To control independent variables we used the SmartBoard’s rear projection surface for all conditions. Projecting onto the front of the surface (as FP, WFP, and PVRP do) causes a “ghosting” of the image due to multiple reflections from the front and back faces of the surface and the touch sensitive overlay used for input. WFP and PVRP, which both use off-axis projectors, were at a distinct disadvantage, as the rear projection display surface is specifically manufactured to be used in an on-axis configuration, and off-axis projection results in a visible blurring of the image due to the “across-the-grain” projection. In the post session interview we found that the factor leading to the image quality score was primarily the sharpness (or blurriness) of the image (100%-P: 1-17) with some of the participants citing intensity or color saturation (29%-P: 4,7,8,13,16) and shadows (6%- P: 5) as additional factors. Some participants mentioned multiple factors and were counted in each category for factors leading to their image quality, preference and acceptance rankings.

Preference - Rear projection was the overall favorite on the preference question *(“Please rate the display technology on the following scale for the tasks performed. [Definite dislike = 1 2 3 4 5 6 7 = Liked very much]”)*. When asked to volunteer what factors they considered when making their preference judgments, about half of the participants mentioned image quality (65%-P: 1,3,5,6,7,9,10,12,13,16,17) and an equal number mentioned shadows (65%-P: 2,3,4,5,6,8,10,11,13,14,15), or lack thereof.

Acceptance - The user acceptance question (*“Please rate your willingness to use this display technology on the following scale: [Absolutely unacceptable = 1 2 3 4 5 6 7 = Completely acceptable]”*) was designed to determine if users would be willing to use a display technology, even if it was not their first choice (preference). Trends followed the preference rating question with slightly higher differences. When asked to volunteer what factors contributed to their acceptance rating, more than half mentioned image quality (53%-P: 2,3,4,5,6,9,14,16,17), and shadows (53%-P: 4,6,8,9,11,12,13,14,15). Ease of performing the task (P: 6,9), touch-screen problems (P: 7,12), unspecified reasons (P: 10) and “just kind’a a gut reaction” (P: 1) made up the remainder of responses.

3.2 Quantitative Measures: Speed & Accuracy

The Box task was specifically designed to generate output that would be likely to fall within (and be hidden by) the user’s shadow. We measured the difference in acquire time between occluded and unoccluded boxes, as well as observed the behaviors they adopted to compensate for shadows (see section 3.3). Figure 3 shows the time difference between occluded and unoccluded boxes, demonstrating the performance penalty experienced by users under occluding conditions. WFP (with 66 occluded - 4.9% of all boxes) and PVRP (with 4 - 0.3%) lower the number of occlusions dramatically in comparison to FP (with 178 - 13.1%). The majority of occluded boxes fell in the bottom left and bottom center quadrants of the screen. However, the number of occluded boxes was insufficient to significantly affect the overall task completion time.

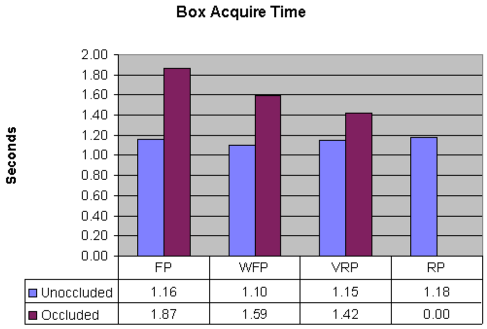

Figure 3: Acquire time for occluded vs. unoccluded boxes.

We found no significant difference between the four conditions for accurate selection, as measured by the crosses task.

The Spiral task measured the user's ability to quickly trace a curve, exercising muscle motions similar to free form drawing or writing, in a more controlled setting. Conditions which eliminated or reduced shadows (RP & PVRP) had slightly faster mean completion times than conditions which did not (FP & WFP), but these trends are not statistically significant.

3.3 Observations

3.3.1 Occlusion & Shadow Coping Strategies

Behavior in the PVRP and RP cases (minimal to no occlusions) were identical, with almost all participants standing near the center of the screen with feet shoulder-width apart ("a-frame" stance), moving only their arms to reach around the screen. When compensating for occlusions in the FP and WFP conditions, participants generally used one of the following four strategies. Almost all participants settled into a single strategy fairly quickly (within 10 boxes).

- *Edge of Screen* (7 of 17 participants) - Participants stood at the edge of the screen. Participants 2,9,13, and 15 would lean inward to move boxes, immediately returning to their home position to insure that they were not occluding the next box. Participants 1,8, and 14 stood slightly in from the edge, so they would occasionally occlude boxes on the left edge. When unable to find a box, they would sway their upper body from the waist until the box they were occluding became visible.

- *Near Center* (7 of 17 participants) - These participants would stand near the center of the screen (usually with their right shoulder directly above the target). Participants 5, 12 and 16 were short enough to occlude few boxes, while participants 6,7,10, and 17 would occlude boxes and use the above "sway" strategy to find occluded boxes.

- *Move on Occlusion* (3 of 17 participants) - These participants (P4,P9,P11) would move to a new position whenever they occluded a box, and stay there until they occluded another box, at which point they would move again.

- *Dead Reckoning* (1 of 17 participants) - Participant 3 stood near the center of the screen so that his shadow would occlude only a single box (position #5, lower left). Whenever he did not see a box, he would blindly select the area in his shadow where the box should be located (with an impressive degree of accuracy) and drag it to the target. (When performing the spiral task, participant 3 would "drag through" his shadow along the curve, also with impressive accuracy.)

For the Crosses task, most participants would work around their shadows, usually standing to the left of the cross they were currently working on. For the Spiral task, all participants (other than P3) would sway their body out of the way of the portion of the spiral they were currently tracing, giving a "tree swaying in the wind" appearance.

3.3.2 Participant Awareness of Shadow Coping Strategies & Preference Ratings Factors

About half of the participants (47%-P:2,4,6,8,9,13,14,15) volunteered that they developed strategies to cope with occlusions, (*"Where their any specific strategies you used to perform the tasks?"*) while others (43%-P:1,3,7,10,11,12,16,17) only recognized that they had done so when asked by the interviewer (*"Did you have any problems with shadows in any of the conditions?" / "How did you deal with them?"*) and one participant (6%-P: 5) who had only occluded 3 boxes (the average participant occluded 14.6 boxes) declared that they had no problems with the shadows.

Interestingly, of the eight participants who volunteered that they had developed strategies to deal with the shadows, seven (P: 2,4,6,8,13,14,15) stated that shadows were a factor in their preference ratings, while one (P: 9) only reported having considered image quality. Of the eight who only recognized their shadow coping behavior after being prompted by the interviewer, three (P: 3,10,11) cited shadows as a factor in their preference ratings, while five reported using image quality exclusively (P: 1,7,12,16,17).

4 Flexible Placement of Projected Displays

We have developed a real-time projector-camera calibration system[9] that could add flexibility to virtual rear projection. With this system, a user may specify the location of the display by placing four fiducial markers to define the corners of the desired projection area. Through visual tracking of both the positions of the four markers and the corners of the the quadrilaterals formed by each projector's output, the appropriate warping parameters can be updated in real-time. Recent improvements to the robustness of our calibration system enable a user to steer a display to multiple display regions in the spirit of IBM's steerable interfaces[6]. The user may move a projected display to any pre-defined region to have the display "snap" into place in a stable manner (see Figure 4). We now describe these improvements and present preliminary results for steerable projected displays.

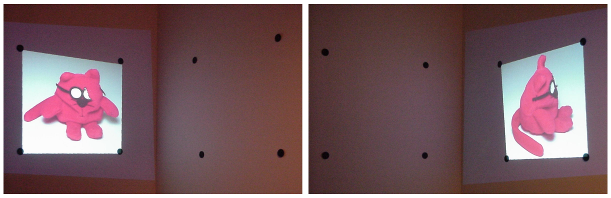

Figure 4: **Left**: Initial display calibrated to fit inside region defined by poker chips. **Right**: Corrected display after steering it from a region defined on a front-facing wall to one defined on a side wall. Note that the view of Catbert has switched from the front to the side as it is mapped to the side wall.

4.1 Real-Time Calibration

A key issue for the robustness of the projector-camera system is the ability to re-calibrate the homography mapping the camera's image to the projected image if either the camera or the projector are moved. In addition, a basic question is how to specify the location of the display. We now describe a real-time calibration system which addresses both of these concerns.

The system uses a set of four fiducial markers on a display surface such as a wall or table to define the four corners of the desired projected area. Since walls tend to be light colored, we have found that any small dark target, such as a poker chip, can serve as a good fiducial marker. By positioning the targets appropriately on the display surface, the user can identify the desired display area. We register the display to the defined projection region by texture-mapping an appropriately warped and positioned quadrilateral. The parameters used to position this quadrilateral are calculated by tracking both the positions of the four markers and the corners of the quadrilateral formed by the projector's output.

In some applications the positions of the projector or camera, as well as the positions of the markers, may change over time. We can view each of these changes as disturbances that perturb the calibrated relationship between the cameras, projectors, and display. In this instance, disturbance rejection can be easily accomplished by tracking the quadrilateral corners and marker positions in real-time, and updating the warping parameters appropriately. Note that the dynamics in this case are extremely fast, since the only limit on the speed at which the projector output can be changed is the overall system bandwidth.

We tackle the problem of incorrect warping due to inaccurate projection corner and marker detections by inspecting a buffer of previous detections. The system only permits an update of the warping parameters if the buffer of coordinates passes two simple tests: (1) for each marker and projection corner point, the distance between the median coordinates in the buffer and the currently stable coordinates must be greater than a chosen distance threshold; (2) the variance of the coordinates in the buffer from the mean coordinate must be less than a chosen variance threshold. The first test accounts for noisy coordinate location detections within a small neighborhood due to jittery, yet correct detections. The second test adequately handles incorrect detections because false detections should occur sporadically while true positives should remain fairly constant over time within a small neighborhood.

4.2 Poker Chip Tracking

The problem of tracking the fiducial markers is challenging when the projection image is constantly changing, such as during movie playback, or when room lighting conditions vary. Basic thresholding is obviously inadequate because it assumes the intensity distribution will not deviate from a model distribution used for picking a static threshold. We present a simple algorithm for tracking fiducial markers that adapts to illumination changes in real-time.

The algorithm is based on the assumption that the fiducial markers will always be among the darkest objects in the field of projected light. We tested a number of potential markers, such as bike reflectors and pieces of black felt, and found blue poker chips to be a good choice since they contrast well with most projected images and are convenient for hand manipulations due to their rigidity and form factor.

The algorithm performs threshold segmentation across a user-specified range of threshold intensities centered on the mean intensity. The system rejects all connected components that do not follow two heuristic rules: (1) connected component has square aspect-ratio within a user-specified tolerance; (2) area of connected component is within user-specified range. These heuristic rules are based on the assumption that connected components representing the poker chips will have a square aspect-ratio and occupy a "natural" percentage of the total image area. All connected components across the range of threshold images are binned into a grid that divides the image into 40 rows and columns. The mean positions recorded in the four peaking bins of the histogram are then returned as the detected poker chips.

4.3 Experimental Results

Previously[9], we reported results of three experiments to evaluate the ability of the visual feedback loop to compensate for disturbances to the projector and camera positions and the positions of the fiducial markers. We summarize these results in Figure 5. Since then, we have improved the robustness of the system. A user may define multiple display regions using multiple sets of poker chips and move the projector to target any region visible to the camera. This “steerable” feature of our auto-calibration system is possible due to the improvements to the poker chip detector, which make it robust to extreme lighting variations, and also due to the added stability of warping parameter updates. We tested projector retargeting by pointing the camera at the corner of a room and defining a projection region on each of the intersecting walls. Our system successfully re-warped the output to fit within the new region when moved from one wall to the next as illustrated in Figure 4. The retargeting of the display happened immediately and in a stable manner (display only moved once the projector covered all chips and remained static for one second).

5 Summary

In this paper we reported on an empirical study comparing front projection (FP), warped front projection (WFP), passive virtual rear projection (PVRP), and rear projection (RP) and demonstrated a method to control the placement of projected displays.

We feel the fact that users did not use occlusion coping strategies when using passive virtual rear projection is an important indication of the benefits it provides over FP & WFP. A rear projected display was prefered by users. Passive virtual rear projection, although ranked lower than

Figure 5: (a) and (b): The effect of a change in the marker configuration on the system is shown at two different time instants. The four markers define an interface by which the user can control the size and location of the display. (c) and (d): The effect of a change in the projector position is shown at two different time instants. The projector quadrilateral changes while the display defined by the markers does not. (e) and (f): The effect of a change in the camera position is shown at two different time instants. The entire image is distorted, but the display continues to fill the region defined by the markers.

rear projection, was ranked higher than warped front projection and front projection.

Users preferred rear projection over passive virtual rear projection, indicating that passive VRP is not yet good enough to replace rear projection. We plan to continue developing active virtual rear projection with occluder light suppression to make it more indistinguishable from true rear projection.

References

[1] R. Balakrishnan, G.W. Fitzmaurice, G. Kurtenbach, and W. Buxton. Digital tape drawing. In *ACM Symposium on User Interface Software and Technology*, pages 161–169, 1999.

[2] P. Baudisch, N. Good, and P. Stewart. Focus plus context screens: combining display technology with visualization techniques. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 31–40. ACM Press, 2001.

[3] 3M IdeaBoard. http://www.3m.com/.

[4] C. Jaynes, S. Webb, R.M. Steele, M. Brown, and W. B. Seales. Dynamic shadow removal from front projection displays. In *Proceedings of Visualization 2001*, pages 175–182. IEEE Press, 2001.

[5] E.D. Mynatt, T. Igarashi, W.K. Edwards, and A. LaMarca. Flatland: New dimensions in office whiteboards. In *CHI'99 Proceedings*, pages 346–353, 1999.

[6] G. Pingali, C. Pinhanez, A. Levas, R. Kjeldsen, M. Podlaseck, H. Chen, and N. Sukaviriya. Steerable interfaces for pervasive computing spaces. In *Proceedings of PerCom*, 2003.

[7] C. Pinhanez. The everywhere displays projector: A device to create ubiquitous graphical interfaces. In *Proceedings of Ubiquitous Computing (Ubicomp)*, pages 315–331, 2001.

[8] R. Raskar. ilamps: Geometrically aware and self-configuring projectors. In *Proceedings of SIGGRAPH*, 2003.

[9] J. Rehg, M. Flagg, T.J. Cham, R. Sukthankar, G. Sukthankar. Projected light displays using visual feedback. In *Proceedings of ICARCV*, 2002.

[10] R. Sukthankar, T.J. Cham, G. Sukthankar. Dynamic shadow elimination for multi-projector displays. In *Proceedings of CVPR*, 2001.

[11] T.-J. Cham, J. Rehg, R. Sukthankar, G. Sukthankar. Shadow elimination and occluder light suppression for multi-projector displays. In *Proceedings of CVPR*, 2003.