Shadow Elimination and Blinding Light Suppression for Interactive Projected Displays

Jay Summet, *Member*, *IEEE*, Matthew Flagg, *Member*, *IEEE*, Tat-Jen Cham, *Member*, *IEEE*, James M. Rehg, *Member*, *IEEE*, and Rahul Sukthankar, *Member*, *IEEE*

Abstract—A major problem with interactive displays based on front projection is that users cast undesirable shadows on the display surface. This paper demonstrates that shadows can be muted by redundantly illuminating the display surface using multiple projectors, all mounted at different locations. However, this technique alone does not eliminate shadows: Multiple projectors create multiple dark regions on the surface (penumbral occlusions) and cast undesirable light onto the users. These problems can be solved by eliminating shadows and suppressing the light that falls on occluding users by actively modifying the projected output. This paper categorizes various methods that can be used to achieve redundant illumination, shadow elimination, and blinding light suppression and evaluates their performance.

Index Terms—Shadow elimination, blinding light suppression, projector calibration, multiple projector display.

1 Introduction

The increasing affordability of high-quality projectors has allowed the development of multiprojector systems for video walls [1], [2], [3], [4] and immersive 3D virtual environments [5]. In most of these systems, cameras are used to coordinate the aggregation of multiple projectors into a single large projected display. In constructing a video wall, for example, the geometric alignment of overlapping projector outputs can be accomplished by using a camera to measure the keystone distortions in projected test patterns and then appropriately prewarping the projected images.

While scalable projector-based displays are quite attractive for visualization applications, interacting with a projected light display currently involves some unpleasant trade-offs. Rear-projection systems support seamless interaction (for example, the Xerox Liveboard [6] or the Stanford Interactive Mural [7]), but can be prohibitively expensive and difficult to deploy because of the need for custom installation of a large display screen and the significant space required behind the screen.[1]

Nonprojected display technologies (such as LCD and plasma display panels) can be used to produce small to medium-sized displays, but become prohibitively expensive as the display size increases. Additionally, front-projected displays offer two novel advantages over rear-projected and nonprojected display technologies. First, they can be moved or steered (via mobile projectors or mirrors) to provide multiple asynchronous displays at several different locations in the environment [9]. Second, they can be used to augment objects in the environment for applications ranging from retail sales to projector-guided painting [10], [11].

Using front projection reduces the system's overall cost, increases portability, and allows easier retro-fitting of existing spaces, but suffers from two problems when a user (or inanimate object) moves between the projectors and the display surface: 1) shadows are cast on the display surface and 2) computer graphics and bright light are projected on the user, which is often a source of distraction for the audience and discomfort for the user.

In this paper, we examine five front-projection systems and demonstrate that these problems can be solved without accurate 3D localization of projectors, cameras, or occluders. We discuss multiple methods we have developed to mitigate and eliminate shadows as well as blinding light. We present a range of algorithms with different hardware and processing requirements that vary from simple passive solutions to fully reactive systems.

We begin with Warped Front Projection (WFP), a single projector passive solution that moves shadows to the side of users, allowing them to interact with graphics directly before them on an interactive surface (Fig. 1). Passive Virtual Rear Projection (PVRP) incorporates a second projector calibrated to overlap with the first, which adds redundant illumination to fill in shadows (Fig. 2). Other than some initial calibration, these passive solutions require minimal additional computation while running, but have the drawback that they project blinding light onto users and do not completely eliminate shadows.

By observing the screen with a camera, Active Virtual Rear Projection (AVRP) can detect shadows cast by users and correct for them by boosting the illumination levels, but still casts blinding light on users. AVRP with Blinding Light Suppression (AVRP-BLS) detects when projectors are

1. Average cost of US office space is $77 per square foot [8].

- *J. Summet, M. Flagg, and J.M. Rehg are with the GVU Center & College of Computing, 85 5th Street NW, 3rd Floor, Atlanta, GA 30332-0760. E-mail: {summetj, mflagg, rehg}@cc.gatech.edu.*
- *T.-J. Cham is with the School of Computer Engineering, Nanyang Technological University, Blk N4 2A-32, Nanyang Avenue, Singapore 639798, Singapore. E-mail: astjcham@ntu.edu.sg.*
- *R. Sukthankar is with Intel Research Pittsburgh, 4720 Forbes Ave, Suite 410, Pittsburgh, PA 15213. E-mail: rahuls@cs.cmu.edu.*

Manuscript received 16 Feb. 2006; revised 22 July 2006; accepted 16 Oct. 2006; published online 2 Jan. 2007.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCG-0014-0206.
Digital Object Identifier no. 10.1109/TVCG.2007.1007.

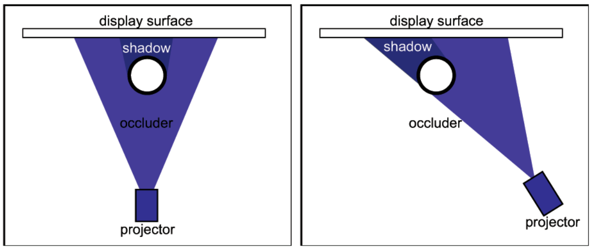

(a) (b)

Fig. 1. Front Projection and Warped Front Projection.

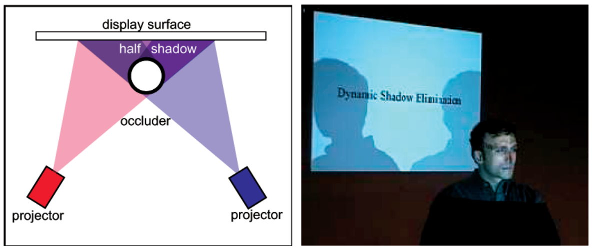

(a) (b)

Fig. 2. (a) Passive Virtual Rear Projection. (b) Graphics within the penumbral shadows are clearly visible, but the shadows are uncorrected. Additionally, the user is subject to blinding light.

shining on users and eliminates blinding light in addition to shadows. The final system we present, Switching, reduces the computational complexity of the AVRP-BLS system and a version implemented on the Graphics Processing Unit (GPU) allows for real-time operation. In addition to the techniques and algorithms, we present an empirical evaluation of the differing systems' relative performance and give adoption guidelines for different deployment needs.

The techniques presented in this paper are particularly useful for projecting interactive surfaces at various locations throughout an environment from steerable projectors. Our solutions provide robustness in the face of occlusions, which will be especially important for steered or reconfigurable projected displays, as these projectors may not always be positioned in an optimal position with respect to the display surface.

2 WFP

The simplest method to minimize shadows on the display surface and reduce the amount of blinding light being cast on users is to move the projector to a position near the display surface and project at an extreme acute angle. An obliquely mounted projector can have two benefits: 1) it moves the location of the light beam so that the user is less likely to be in it and 2) it may shift the shadow away from the important areas of the display surface. A standard data projector can be mounted at a moderately acute angle, and commodity 3D video cards can be used to prewarp the projected image to compensate for keystone distortions, but the projector must remain positioned such that all portions of the display remain within its field of focus.

Commercial projectors designed to be mounted within 1 m (3 feet) of the display surface use specialized optics such as aspheric mirrors to warp the projected image [12], [13]. In addition to image warping to compensate for keystone distortions, these optics also have appropriately varying focal lengths for the varying lengths of the beam path.

Even with a very acute projection angle provided by expensive optics, these warped front-projection systems suffer from occlusions whenever the user comes close to or touches the display, making them unsuitable for interactive applications. The areas of occlusion can be filled in by using a second projector to provide redundant illumination.

3 PVRP

By adding more projectors, it is possible to create a display that is more robust to occlusions. We use the general term *Virtual Rear Projection* (VRP) to describe the class of display systems that use multiple redundant front projectors to approximate the experience of a rear-projected surface. Most areas that are occluded with respect to the first projector can be illuminated by a redundant projector with an unoccluded view. Shadows resulting from all of the projectors being occluded are termed *umbral*, and those where at least one projector is not occluded are termed *penumbral*. By definition, the system cannot control lighting within an umbra, so we strive to avoid umbral occlusions by positioning the projectors so that the display is illuminated from several different directions. The greatest challenge to providing passive redundant illumination (Fig. 2) is for the system to accurately align the projected images on the display surface. Our system uses computer vision to calculate the appropriate homographies (a projective transform) to align the projected images to within one pixel. More information on how the homographies are calculated by observing calibration patterns projected by the projectors can be found in Section 2 of the supplementary material online and in our prior publications [14], [15], [16]. Once the projectors are calibrated, the output of each projector is warped (using the 3D graphics hardware) so that the images are aligned on the display surface.

Mounting the projectors off-axis and warping their output has several consequences. First, having to warp the output slightly reduces the number of pixels available for display, as some pixels around the edge of the warped image will not be used. Second, if the warping function is not performed carefully, the resulting image may appear jaggy or aliased. Finally, as the projectors are moved off-axis, the length of their optical path to each side of the display surface changes, and one side may be more "in-focus" than the other. We have found that mounting generic projectors 35° off-axis (for 70° of total angular separation) gives good redundant light paths while reducing resolution by less than 20 percent and keeping both sides of the display surface within the projectors' depth of focus. Because we are performing bilinear interpolation on the 3D graphics hardware, the warped images do not suffer from aliasing.

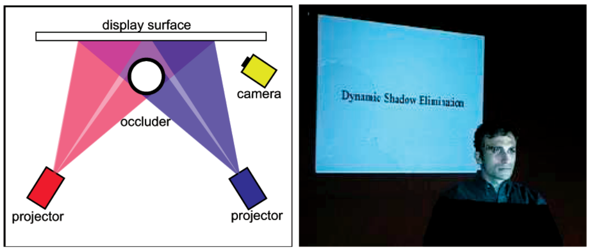

(a) (b)

Fig. 3. (a) Active Virutal Rear Projection. (b) Penumbral shadows are eliminated, but the blinding light remains.

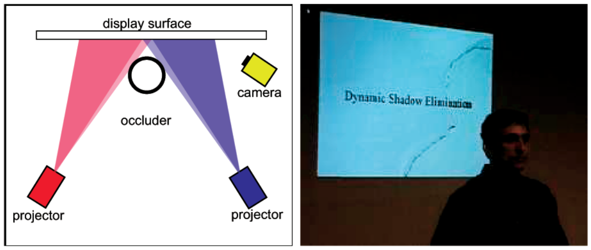

(a) (b)

Fig. 4. (a) AVRP with Blinding Light Suppression. (b) Light is kept off the occluder's face.

4 AVRP—Shadow Elimination

By adding a camera or other sensor that is able to detect the shadows on the display surface (Fig. 3), it is possible to dynamically correct penumbral shadows by projecting additional light into the region from one of the nonoccluded projectors. This AVRP system must precisely adjust projector output to compensate for each occlusion. If too little light is added, the shadow will remain visible; if too much light is used, overillumination artifacts will be created. The shadow boundaries must be treated carefully since humans are very sensitive to edge artifacts. A detailed description of the AVRP system has been published in CVPR 2001 [15] and is included in the supplementary material online in Section 3.

When initialized, the system projects each image it wishes to display and creates a reference image, representing the desired state of the display. AVRP uses a camera to detect artifacts on the display surface, caused either by uncorrected shadows or overillumination on the display surface. A feedback loop is used to correct the artifacts by comparing the current system's state with the reference image. This system is robust and efficient, able to correct for fuzzy occlusions and diffuse shadows without requiring an explicit model of the shadow formation process. However, it suffers from several drawbacks. It requires an unoccluded view of the display surface and only works with known reference images, making it unsuitable for most interactive displays. Additionally, while in the process of correcting shadows, the system actually *increases* light projected on the occluder. In the case of humans, this blinding light is distracting [17].

5 Blinding Light Suppression

To suppress blinding light, the AVRP-BLS must detect which pixels in each projector are shining on an occluder and, thus, have no effect on the display surface (Fig. 4). We determine if a projector pixel is occluded by determining if changing it results in a change on the display. Unlike the AVRP system, which modifies all projectors simultaneously, the AVRP-BLS system modifies each projector in turn, watching for changes on the display surface. If changes are not detected for pixels from a specific projector, those pixels are marked as occluded, and their output is driven to black, removing light on the occluders. Note that the probe technique must be employed during shadow elimination as well. In particular, the system must be able to discover when the light from a suppressed pixel is available again. This constraint is smoothly incorporated into our algorithm. A full discussion of the AVRP-BLS algorithm is available in Section 4 of the online supplementary material and was published in CVPR 2003 [14].

The following synthetic example illustrates the algorithm. For a particular screen pixel at a typical steady state when shadows have been eliminated, suppose the corresponding source pixels (1 and 2) from different projectors have an alpha mask value of $\alpha = 0.5$.

When source pixel 1 is suddenly occluded (source pixel 2 remains unoccluded), our system increases the α value for source pixel 1. When no change is observed, the α value for source pixel 1 is driven to zero, and the α value for source pixel 2 is raised and subsequently adjusted until the display matches the stored reference image.

Note that even when source pixel 1 becomes unoccluded, nothing changes if source pixel 2 remains unoccluded since the shadows have already been satisfactorily eliminated (Fig. 5). This particularly illustrates the *hysteresis effect* in which source pixels are not boosted or blanked until new shadows are created.

The drawback to such an iterative technique is that the alpha mask can require several iterations to converge; in

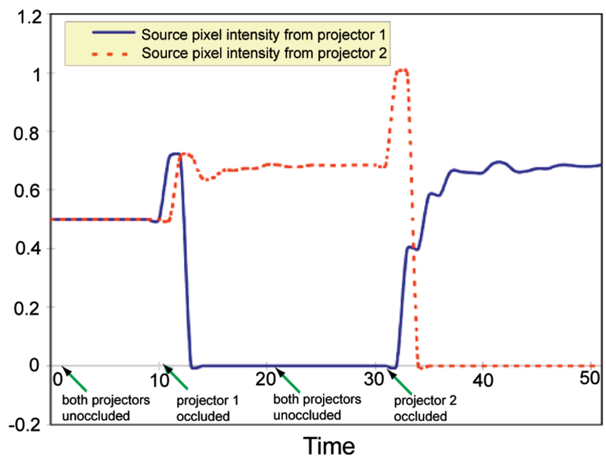

Fig. 5. Synthetic example of transitions in projector source pixel intensities. This graph shows the intensity transition of two corresponding projector source pixels over time, subject to four events of occlusions and deocclusions. Note the hysteresis effect; the source pixels are not boosted or blanked until new occlusion events occur.

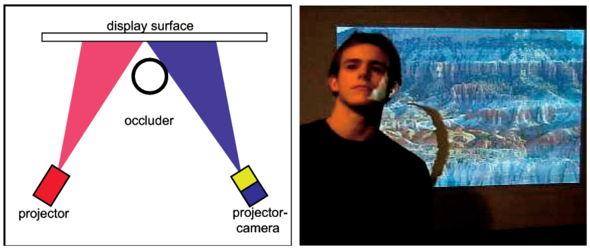

(a) (b)

Fig. 6. (a) Switching VRP. (b) Shadows are eliminated, and blinding light is suppressed with a moving user. The transient shadow created by the moving occluder is corrected in one iteration.

(a) (b)

Fig. 7. **Seam blending:** (a) before and (b) after.

practice, shadows are eliminated in approximately five to seven iterations.

In our software-only implementation, the AVRP-BLS system is able to calculate 1.6 iterations per second. Even assuming advances in processing power, when using commodity projectors that are limited to 60 or 85 fps, a series of five to seven iterations would produce a visual artifact for up to one-tenth of a second.

There are two possible solutions to making the changes to the display unnoticeable to humans. The first method is to greatly increase the speed of the entire feedback loop. This would require projectors and cameras that operate at 120 fps or faster. The second method is to detect the occluder instead of the occlusion (shadow) and use that knowledge to correct the occlusion as (or before) it occurs.

6 Switching VRP

The previous systems provide redundant illumination from multiple projectors, dynamically adjusting the amount of illumination from each projector on a per-pixel basis based upon the feedback provided by a camera observing the projected display. The downside of these approaches is that they assume that the camera has an unoccluded view of the display surface. We can relax this assumption by detecting the occluder instead of the occlusion (shadow). However, as we would no longer have an unobstructed view of the display, we would have to correct the projector's output blindly, without feedback. To do this successfully, each pixel on the display surface must be illuminated by only one projector at a time. As the projector illuminating a pixel becomes occluded, the responsibility for illuminating that pixel is shifted to another (unoccluded) projector (Fig. 6). This presents several challenges:

1. The system must know which pixels are occluded for at least $N-1$ of the N projectors in the system so that it can correctly assign pixel regions to unoccluded projectors to ensure that a complete image appears on the display surface, regardless of occlusions that may partially block portions of each projector.
2. The output from all projectors must be photometrically uniform so that any projector can "fill in" for any other projector without a noticeable change in intensity or color.
3. The subimages projected from each projector must overlap in such a way as to produce a uniform output image without visible seams or intensity/color shifts. To achieve this, the edges of each image must be blurred so that they blend together imperceptibly.

6.1 Occlusion Detection

In our approach, we position a camera close to the projector lens of the primary projector so that detected occluder silhouettes align with corresponding projector mask silhouettes with little to no parallax effects caused by projector-camera disparity. If the optical axes of the projector and camera are aligned by means of a beam splitter, parallax effects are eliminated [18]. To simplify the detection of occluders, the camera is filtered to detect only infrared light, and the display surface is illuminated with infrared lights. Background subtraction of the infrared (IR) camera images is not affected by light projected from the projectors and, as shown in Fig. 10b, the backlit silhouette of occluders creates a strong contrast between foreground and background.

Because we are detecting occluders (instead of shadows), we do not need to preshoot background plates for each expected frame [14] nor predict the expected appearance of each image when projected onto the display surface [19], [20].

For each iteration of the algorithm, the IR camera image must be used to calculate the final alpha masks for each projector through a series of steps illustrated in Fig. 10. First, the acquired image is warped to align with the display surface using a camera-surface homography, which can be calculated in the same manner as the homographies of Section 3. Second, the image is segmented into occluder and nonoccluder regions using simple background subtraction. In some cases, median filtering is needed for noise removal, but in our experiments, the backlit occluders were easily segmented without noise. Third, the occluder regions are dilated to allow a region of tolerance for occluder movement between each compensation step. Finally, the mask is blurred to blend seams between projectors. Fig. 7 illustrates the need for blending to avoid distracting seams.

6.2 Photometric Uniformity

The projected display from each projector must appear photometrically uniform to ensure visual consistency. Calibration for photometric uniformity is necessary to make the hand-off of a pixel from one projector to another imperceptible.

Majumder and Stevens have found that the major source of apparent color variation across multiple projectors is due to luminance variation, and that the chrominance of projectors (of the same brand) are very similar [21]. Their work has

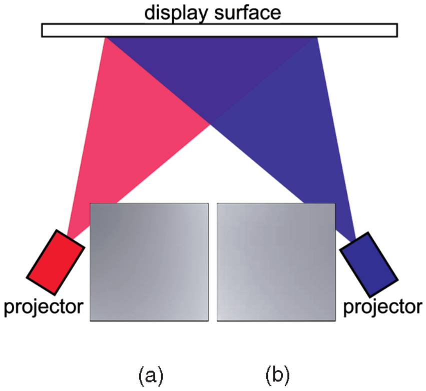

Fig. 8. **LAMs:** (a) LAM for projector positioned to left of projection surface. (b) LAM for projector positioned to the right of the projection surface. Note that the dark regions of each LAM correspond with the shortest projection distance to the display surface.

focused on tiled multiprojector displays where the projectors are oriented perpendicular with the display surface.

In a VRP system, the projectors are oriented as much as 50° from the normal, with a 30° to 45° off-axis orientation being typical. This extreme angle causes drastic changes in the level of illumination from each projector across the display surface. The side of the display surface closer to the projector is overilluminated, whereas the far side is underilluminated. This angle-induced ramping function is in addition to the variations in projector illumination found by Majumder and Stevens.

To correct for the intensity variance in our VRP system, we use luminance attenuation (alpha) masks, which modify the intensity of each projector pixel so that all pixels are evenly illuminated, regardless of their location on the display surface or which projector is currently illuminating them.

The method we use to generate the attenuation maps is similar to that used by Majumder and Stevens for their Luminance Attenuation Maps (LAMs) [22], except that it does not require a calibrated projector or camera. The darkest intensity measured when projecting white from each projector independently is set as a target. All pixels are iteratively reduced in intensity (to account for nonlinear projector and camera responses) until the target intensity is uniform across the display. Fig. 8 shows two example LAMs, and Fig. 9 describes our algorithm for LAM creation.

```
CREATE-LAMS:

for each projector p
  1. project white for p and
     black for all other projectors
  2. capture image
  3. if darkest intensity d for projector p
     is darker than overall darkest intensity d*,
     d* = d
  4. initialize LAM(i,p) = white for all pixels i
end for

for each projector p
  initialize j = UPDATE_LIMIT
  project black for all other projectors
  while j > 0
    project LAM(*,p) and capture image
    for each pixel i
      if (intensity(i) > d*)
        LAM(i,p)--
      end if
    end for
    j--
  end while
  low-pass filter LAM(*,p)
end for
```

Fig. 9. LAM creation pseudocode.

6.3 Seam Blending

We assume that the output image from each projector is already geometrically aligned on the display surface and the output of each projector has been modified to be as photometrically uniform as possible. Our goal is to project portions of the image from different projectors while retaining a final displayed image that appears uniform and without edge artifacts. This can be achieved by using edge-blended alpha masks to limit the output of each projector, which are generated as follows:

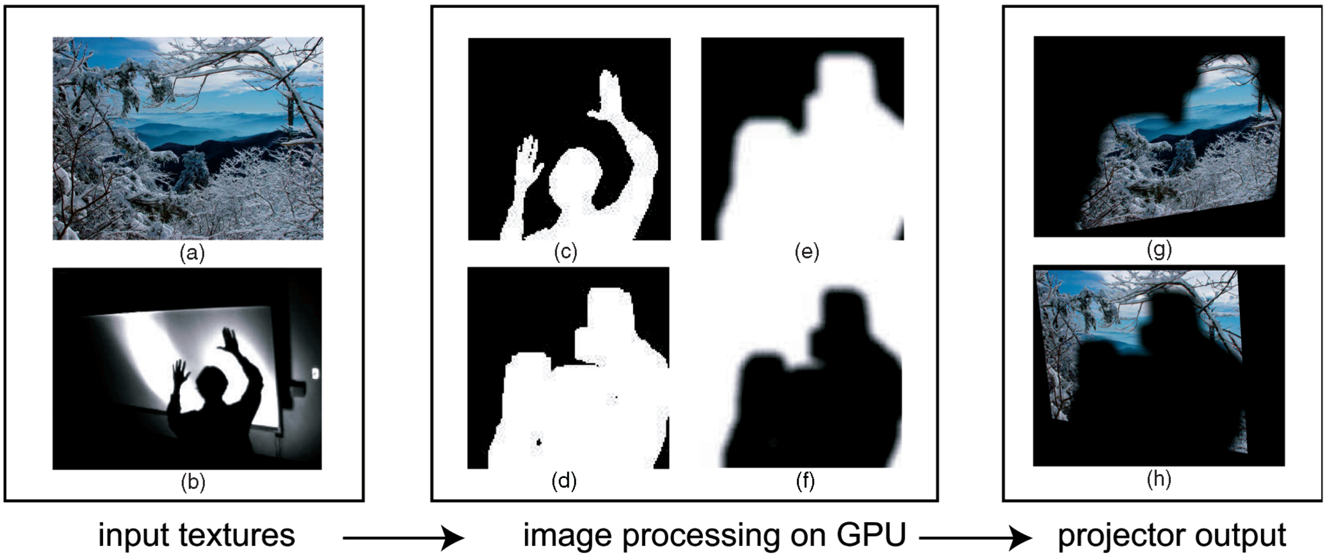

Fig. 10. **GPU-centric architecture:** (a) display texture, (b) IR camera frame, (c) occluder mask texture, (d) dilated mask to tolerate interframe occluder movement, (e) blurred mask for projector 1 blending, (f) blurred mask for projector 2 blending, (g) keystone-corrected projector 1 output, and (h) keystone-corrected projector 2 output.

1. Order the projectors from $P_0 \ldots P_{N-1}$. Projector P_0 will be initially responsible for the whole display. As it is occluded, projector P_1 will be used to fill in occluded regions. Any regions occluded in both projectors P_0 and P_1 will be handled by projector P_2, and so on. Associate an initially zero alpha mask with each projector $\alpha_0 \ldots \alpha_{N-1}$, which will be used to control the active output pixels.
2. Using the cameras attached to projectors $P_0 \ldots P_{N-2}$ generate an occlusion mask $O_0 \ldots O_{N-2}$ for all but the last projector, indicating which projector pixels are occluded.
3. For the alpha mask of the ith projector α_i, turn on all pixels that are not occluded in the occlusion mask O_i and that have not already been turned on in any previous alpha masks $\alpha_{0 \ldots i-1}$. For the last alpha mask α_{N-1}, turn on all pixels that are off in all alpha masks $\alpha_{0 \ldots n-2}$. This results in a set of mutually exclusive alpha masks that favor projectors based on their ordering. Note that a camera on P_{N-1} does not improve system performance but allows the detection of any umbral shadows.
4. We then perform the following operations on each alpha mask to add a feathered edge that hides the seam:
   - Filter each alpha mask $\alpha_0 \ldots \alpha_{N-1}$ with a 3×3 median filter to remove noise.
   - Dilate each alpha mask three times.
   - Blur the expanded alpha masks with a Gaussian filter to feather their seams.

When the occluders are correctly detected, the result of using these alpha masks to control the output of the projectors is a projected display that appears seamless and shadow free.

6.4 Improving Performance Using the GPU

As users move in front of an active VRP display, they may cast new shadows by moving faster than the system can update the screen. This occurs when the users move outside of the region of tolerance created by the dilation operation before the display is updated. Increasing the system frame rate and decreasing system latency enables users to make quick natural movements such as emphasizing a point with a fast hand gesture. The image processing steps needed for Switching VRP may be optimized by exploiting today's programmable graphics cards (GPUs). Image processing on the GPU shifts the speed limit of Switching VRP away from computation on the CPU to capture and display rates of the camera and projector. Fig. 10 illustrates our image processing pipeline using the GPU.

There are three capabilities of GPUs and DirectX 9.0 that we exercise in order to eliminate the bottleneck of image processing: 1) multiple render targets (MRTs), 2) pixel shaders, and 3) multihead resource sharing. First, the MRTs' capability provided with Direct3D version 9.0 enables us to store the results of each image processing step in an off-screen rendering surface for succeeding filter operations to use as input. By setting the texture coordinates (u, v) of a screen-aligned quadrilateral to correspond with the camera image coordinates (x, y) of the projected display, the camera-surface warp may be performed by rendering the quadrilateral texture-mapped with the camera image. The warped texture is now available on an off-screen surface for subsequent filtering using pixel shaders.

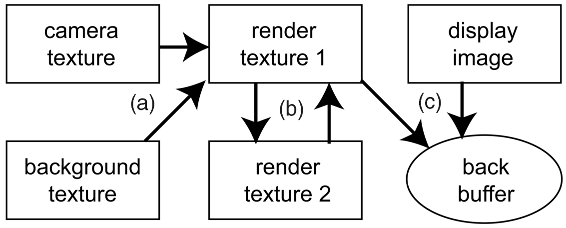

Fig. 11. **Pixel shader pipeline:** Boxes represent textures, and arrows denote texture sampling operations used in pixel shaders. (a) Background subtraction shader stores result in render texture 1. (b) Render textures 1 and 2 are used as sampling buffers for dilation and blurring operations, each of which require two independent shaders. (c) The final occluder mask is composited with a display texture and rendered into the DirectX back buffer for display.

The second capability provided by GPUs is fast image processing using pixel shaders. Background subtraction, dilation, median filtering, and blurring may be implemented as pixel shader programs [23]. These pixel shaders were written in DirectX High-Level Shader Language (HLSL). Using two texture samples and a threshold, the result of a background subtraction shader is stored in the first of two off-screen render targets. Next, dilation is performed using two separate pixel shaders. The first shader dilates the result of background subtraction using 1D texture samples horizontally, and the second dilates the resulting texture vertically. Separating dilation into two operations decreases the number of required texture samples and improves performance from $O(n^2)$ to $O(n)$. To further improve processing time, the two off-screen render textures were reduced to a resolution of 128×128 pixels (to be subsampled during compositing operations). Following dilation, blurring is performed in a similar manner using two separate shaders. Finally, the resulting occluder mask is composited with the display frame using one pixel shader. The interaction between each pixel shader and the input/output textures they use is illustrated in Fig. 11.

Finally, multihead resource sharing in DirectX 9 makes it possible to use one rendering device across multiple display heads. Previously, each head required its own device and therefore needed separate sets of textures and pixel shader computations. By using one device instead of two, some of the pixel shaders need only be executed once, saving time and texture memory. A background subtraction and dilation pixel shader computation is removed. An initial dilation of n pixels is performed to permit sufficient occluder movement within frame updates.

A second dilation of k pixels is needed to overlap projector masks before blending. Before multihead resource sharing, one display device performed $2n$ texture samples, and the other sampled $2(n+k)$ pixels ($4n+2k$ total samples). After multihead sharing, a dilation using $2n$ texture samples is shared among both display heads, and the remaining $2k$ pixels are sampled for the overlapping region ($2n+2k$ total samples), saving $2n$ texture samples per pixel. Following dilation, blurring and compositing operations must be

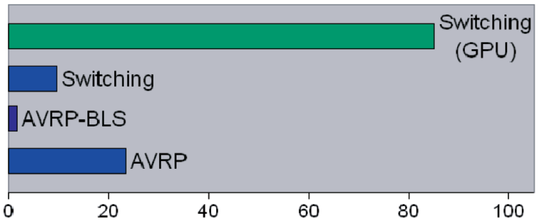

	AVRP	AVRP-BLS	Switching	Switching (GPU)
FPS	23.3	1.6	9.5	85

Fig. 12. **Speed of active systems.** Relative speed in frames per second of the active systems. The GPU implementation of the Switching algorithm is provided for comparison. The 85-fps speed is limited by the output speed of the projectors; the actual update rate on the GPU is 101 fps.

performed for each display head separately due to differences between the occluder masks.

7 Related Work

The 3M Digital Wall Display and the NEC WT-600 projector are commercial products that implement WFP using special optics, but do not allow the user to position the display arbitrarily. Many projectors include a keystone correction feature implemented by a physical “lens shift,” which allows the image to be warped to a limited degree. New advanced video processing chips (such as the Silicon Optix Realta, including eWARP-2 Geometry Processing) are being integrated into high-end projectors emerging in the market, which support off-axis projection of up to $\pm 45°$ (vertical) and $\pm 35°$ (horizontal). Modern graphics cards by nVidia support the nvKeystone feature, which allows users to warp a projected display arbitrarily. None of these commercially available solutions deal with the calibration of multiple displays to provide redundant illumination.

Much work has been done with multiple projectors to build tiled display walls [1], [2], [4], and Chen et al. have used homographies to align multiple tiled projectors in a large-format display wall [24]. These display wall projects did not overlap projectors to provide redundant illumination.

Tan and Pausch used IR backlighting and an IR camera to detect a person and create a black “mask” over projected graphics, which solves the blinding light suppression problem [25]. A commercial appliance sold by iMatte uses a similar technique of masking out projected graphics, but has a more sophisticated segmentation algorithm, allowing it to work with a front projected IR pattern, eliminating the need for backlighting. Both of these systems suppress blinding light, but do not compensate for shadows occurring on the display. Indeed, they actually slightly increase the size of the shadows on the display, as the masks they generate are slightly larger than the users. They are most useful for applications where the user is not trying to interact with the displayed image (for example, while giving a presentation).

Our previous work, as described in Sections 4 and 5, eliminated shadows, but required the collection of a reference image with no occluders present and assumed an unoccluded view to the display [15]. Researchers at the University of Kentucky developed a photometric model, which they use to generate a reference image of arbitrary graphics, predicting how it should appear when projected [19]. But, their system was too slow for interactive use, retained the assumption of an unoccluded view to the display, and did not solve the blinding light problem. Jaynes and colleagues enhanced this work to increase the speed to approximately 9 fps by updating bounding regions instead of individual pixels [20]. Similar to AVRP or AVRP-BLS, their system requires numerous frames to converge to a stable display. Their updated system still requires that cameras have an unoccluded view of the screen and does not eliminate blinding light.

Our work, as described in Section 5, solves the blinding light problem, but was also slow and relied on an unoccluded view of the display. The Switching system presented in Section 6 removes the unoccluded camera view requirement and is currently limited by the projector to 85 fps (Figs. 12 and 13).

Majumder and Stevens used LAMs in multiprojector tiled displays on a per-channel basis to correct for color nonuniformity [21], [22], but their method for generating the LAMs required a full photometric model of the camera and projector(s), which required significant manual processing.

8 Evaluation

To evaluate our systems' relative performance, we performed an empirical evaluation of each of the algorithms discussed previously. In this experiment, each algorithm was run on the same hardware setup. After the algorithms

	Effects on shadows	Effects on blinding light	2nd projector required	Camera	Camera requires unoccluded view	IR filter or dedicated IR camera
WFP	moves shadows aside	none				
PVRP	reduces shadows	none	✓	✓		
AVRP	fills in shadows	none	✓	✓	✓	
AVRP-BLS	fills in shadows	removes blinding light	✓	✓	✓	
Switching	fills in shadows	removes blinding light	✓	✓		✓

Fig. 13. System features and hardware requirements.

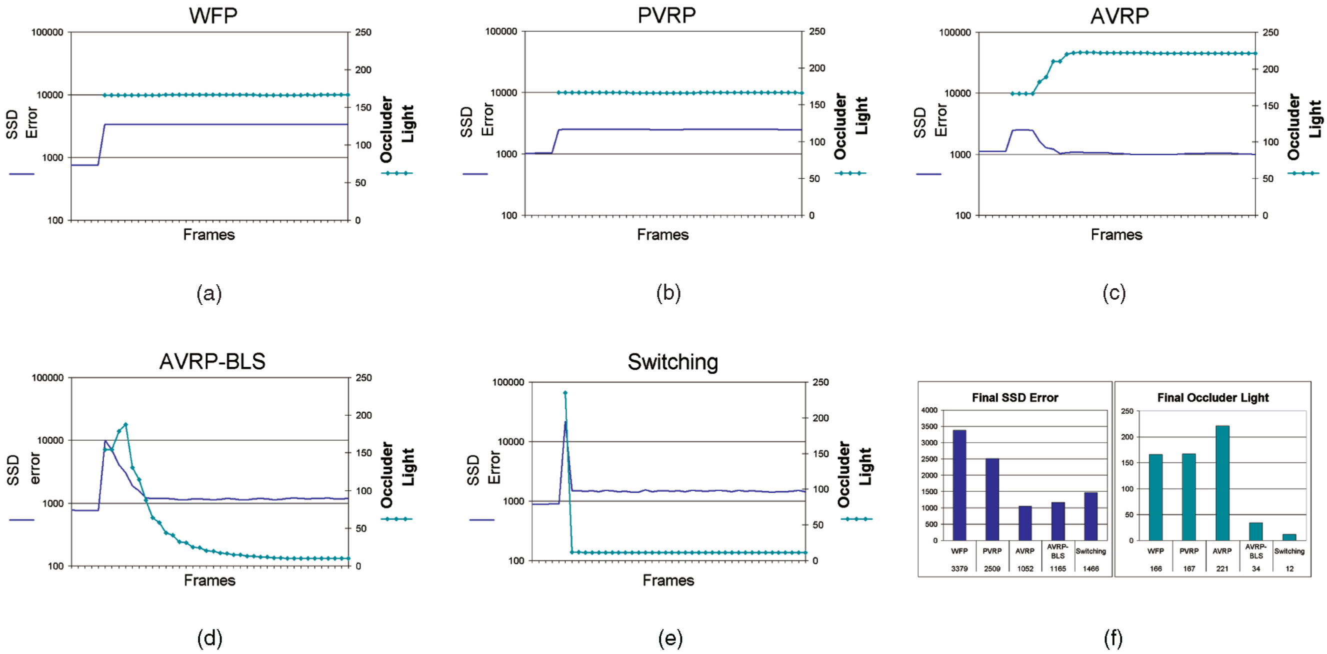

Fig. 14. (a) Warped Front Projection. (a) Passive Virutal Rear Projection. (c) Active Virtual Rear Projection. (d) Active Virtual Rear Projection with Blinding Light Suppression. (e) Switching VRP. (f) Final sum squared distance (SSD) and occluder light measures.

had initialized, we collected a reference frame consisting of the average pixel values on the display with no occluders and then paused the algorithm. We then introduced an occluder into the beam path of one projector and restarted the algorithm.

We used a static occluder that appeared (to the algorithms) instantaneously so that each algorithm would be measured under identical conditions. Because the tests cannot be performed in a simulated environment, we were unable to physically replicate the motion of a dynamic occluder in our lab with sufficient precision to ensure repeatability.

As each algorithm reacted to the occluder (WFP and PVRP took no action), the sum squared difference in pixel values of the camera image from the reference image was recorded on each iteration of the algorithm. A second camera recorded the relative light levels falling on the occluder. The results are presented in Fig. 14 and Table 1.

8.1 Experimental Setup

Each algorithm was run on a dual processor Intel Pentium 4 Xeon 2.2 GHz Dell Precision workstation with 2 Gbytes of RAM. An nVidia GeForceFX 5800 Ultra graphics card on an AGP $4\times$ bus drove two Hitachi CP-SX 5600 LCOS projectors. The projectors were mounted 430 cm apart on a bar 360 cm from the display surface and 240 cm above the floor. The display surface was 181 cm wide and 130 cm high, mounted so that its bottom was 63 cm from the floor. Each projector was 34° off of the projection surfaces' normal, giving a total angular separation of 68° between the projectors.

A Sony N50 3CCD progressive scan camera was used to measure the SSD pixel error seen with respect to a reference image captured before the occluder was introduced. Each algorithm was allowed to initialize normally with no occlusions. The system was then paused, and a static occluder was introduced, partially blocking the beam of one projector. The occluder was a 40.6 cm wide by 50.8 cm high white painters canvas, mounted on a tripod 150 cm from the screen.

After the occluder was introduced, the system was restarted. To the algorithms, this gave the appearance of an instantly appearing occluder, which blocked approximately 30 percent of one projector. In Fig. 14, the occluder appears in frame 5. At this point, the algorithms were allowed to run normally until they had stabilized.

TABLE 1
Algorithm Performance Measures

Condition	Frames to Converge	SSD Error	Occluder Light	F.P.S.	Time to Converge (Seconds)
WFP	n/a	3379	166	N/A†	N/A†
PVRP	n/a	2509	167	N/A†	N/A†
AVRP	7	1052	221	23.3	0.3
AVRP-BLS	7	1165	34	1.6	4.4
Switching	1	1466	12	9.5‡	0.1

† WFP and PVRP do not actively compensate for shadows.

‡ We evaluated a CPU only version of the Switching algorithm so that the FPS numbers are an accurate representation of the relative computational complexity of the active algorithms. The GPU version of the Switching algorithm runs at 85fps, limited by the refresh rate of our projectors.

8.2 Results

In the simple cases of WFP and PVRP, the system performed no compensation, and the light on the occluder and errors in the displayed image are immediately stable. As seen in Table 1 (SSD error) and the graphs in Fig. 14, PVRP improved the image quality over that achieved by a single projector solution (WFP) despite taking no explicit compensatory action.

AVRP, which attempts only to minimize the error of the displayed image, required seven iterations to converge, or 0.3 seconds in real time. After convergence, the SSD error was effectively the same as before the occluder was introduced, although the light cast on the occluder was more than in the nonactive cases. This is due to the fact that the AVRP algorithm increases light output from *both* projectors when attempting to correct a shadow, leading to increased light cast on the occluder.

AVRP-BLS also took seven iterations to converge, but due to the increased processing required, this equated to 4.4 seconds in real time. The benefit of the additional computational time is shown in the amount of light remaining on the occluder, which was reduced significantly when compared to the previously described algorithms.

The Switching VRP system is able to compensate immediately after detecting the occluder (one iteration, or 0.1 seconds). Because it does not employ a feedback loop, the SSD error after compensation is larger than in the AVRP or AVRP-BLS case, but the subjective image quality is good. Occluder light suppression is excellent, with the amount of light cast on the occluder lower than any other algorithm. Additionally, it has the fastest real-time performance of the algorithms discussed.

8.3 Qualitative Results with Moving Occluders

When we evaluated the systems with humans performing interactive tasks, WFP provided benefits over a standard front projector [17]. PVRP and AVRP were able to almost completely eliminate umbral shadows with a single user and performed well with multiple users. The AVRP-BLS system was too slow to be usable, and the software version of the Switching system suffered from a noticeable lag. The GPU implementation of the Switching system is fast enough to compensate for all but the fastest user motion. Readers are referred to the online supplementary material for a video of the GPU version of the Switching system in operation.

9 Conclusion

Front projectors are best suited to applications where a very large display is needed, or where the display will be deployed in an ad hoc and temporary manner. The most extreme example of this is a steerable projector, where the projected image can move in real time under programmatic control.

This paper presents potential solutions to one of the drawbacks of front projection, occlusions and the resulting blinding light. The presented techniques work with generic front projectors. We expect the cost for front projectors to continue to drop, whereas their resolution and brightness will continue to increase.

The introduction of LED projectors will reduce their noise and heat output while reducing power consumption, allowing for the development of mobile battery-powered devices with integrated projectors. Although thin-screen rear projection TVs, plasma screens, and LCD display panels are also dropping in price, due to their integrated display surface, they will never be as small and mobile as projectors, and will have difficulties scaling to wall-sized displays. As electronics become commodity items, their production costs begin to be limited by the cost of raw materials, and the smaller volume of projectors will continue their cost advantages over other display technology.

We recommend system adopters use WFP when limited to a single projector or looking to purchase an off-the-shelf solution. When two projectors are available, PVRP is easy to implement and provides good robustness in the face of occlusions. The Switching form of VRP is the best method that makes use of active compensation for shadows due to the ability to implement it in GPUs of commodity graphics cards and its good balance of speed and image quality.

The techniques presented here enable the projection of interactive displays to arbitrary positions in the environment from steerable or manually redeployable projectors. The suppression of blinding light and elimination of shadows allow for displays to be positioned where they are needed and repositioned on-the-fly, regardless of the position of potential occluders in the environment.

We intend to continue this work by 1) demonstrating VRP on steerable projectors, 2) locating projectors with lower latency to improve our overall performance, and 3) improving the Switching algorithm to implement optical flow on the GPU to predict where occlusions will occur and eliminate shadows before they are cast.

References

[1] R. Yang, D. Gotz, J. Hensley, H. Towles, and M. Brown, "PixelFlex: A Reconfigurable Multi-Projector Display System," *Proc. IEEE Conf. Visualization,* pp. 168-174, 2001.
[2] M. Hereld, I. Judson, and R. Stevens, "Introduction to Building Projection-Based Tiled Displays," *Computer Graphics and Applications,* vol. 20, no. 4, pp. 22-28, 2000.
[3] G. Humphreys and P. Hanrahan, "A Distributed Graphics System for Large Tiled Displays," *IEEE Visualization,* pp. 215-223, 1999.
[4] K. Li et al., "Building and Using a Scalable Display Wall System," *Computer Graphics and Applications,* vol. 20, no. 4, pp. 29-37, 2000.
[5] R. Raskar, M. Brown, R. Yang, W. Chen, G. Welch, H. Towles, B. Scales, and H. Fuchs, "Multi-Projector Displays Using Camera-Based Registration," *IEEE Visualization,* pp. 161-168, 1999.
[6] S. Elrod, R. Bruce, R. Gold, D. Goldberg, F. Halasz, W. Janssen, D. Lee, K. McCall, E. Pedersen, K. Pier, J. Tang, and B. Welch, "Liveboard: A Large Interactive Display Supporting Group Meetings, Presentations, and Remote Collaboration," *Proc. Conf. Human Factors in Computing Systems,* pp. 599-607, 1992.
[7] B. Johanson, A. Fox, and T. Winograd, "The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms," *IEEE Pervasive Computing,* vol. 1, no. 2, 2002.
[8] *Building Construction Cost Data 2002*, P.R. Waier ed. Robert S. Means Co., 2002.
[9] C. Pinhanez, "The Everywhere Display," *Proc. Int'l Conf. Ubiquitous Computing (UbiComp),* pp. 315-331, 2001.
[10] C. Pinhanez and M. Podlaseck, "To Frame or Not to Frame: The Role and Design of Frameless Displays in Ubiquitous Applications," *Proc. Int'l Conf. Ubiquitous Computing (UbiComp),* pp. 340-357, 2005.
[11] M. Flagg and J.M. Rehg, "Projector-Guided Painting," *Proc. User Interface Software and Technology (UIST),* pp. 235-244, Oct. 2006.
[12] 3M, "Digital Walldisplay," http://www.3m.com/meetings/product_catalog/digitalwalldisplay/index.jhtml, 2007.
[13] NEC, "Wt-600," http://www.nec-pj.com/, 2007.
[14] T.J. Cham, J.M. Rehg, R. Sukthankar, and G. Sukthankar, "Shadow Elimination and Occluder Light Suppression for Multi-Projector Displays," *Proc. Computer Vision and Pattern Recognition,* pp. 513-520, 2003.
[15] R. Sukthankar, T.J. Cham, and G. Sukthankar, "Dynamic Shadow Elimination for Multi-Projector Displays," *Proc. Conf. Computer Vision and Pattern Recognition,* pp. 151-157, 2001.
[16] R. Sukthankar, R. Stockton, and M. Mullin, "Smarter Presentations: Exploiting Homography in Camera–Projector Systems," *Proc. Int'l Conf. Computer Vision,* p. 247, 2001.
[17] J. Summet, G.D. Abowd, G.M. Corso, and J.M. Rehg, "Virtual Rear Projection: Do Shadows Matter?" *Proc. Conf. Human Factors in Computing Systems (CHI '05) Extended Abstracts,* pp. 1997-2000, 2005.

[18] S. Nayar, H. Peri, M. Grossberg, and P. Belhumeur, "A Projection System with Radiometric Compensation for Screen Imperfections," *Proc. First IEEE Int'l Workshop Projector–Camera Systems (PROCAMS '03),* 2003.

[19] C. Jaynes, S. Webb, R.M. Steele, M. Brown, and W.B. Scales, "Dynamic Shadow Removal from Front Projection Displays," *IEEE Visualization,* pp. 175-182, 2001.

[20] C. Jaynes, S. Webb, and R.M. Steele, "Camera-Based Detection and Removal of Shadows from Interactive Multiprojector Displays," *IEEE Trans. Visualization and Computer Graphics,* vol. 10, no. 3, pp. 290-301, 2004.

[21] A. Majumder and R. Stevens, "Color Nonuniformity in Projection-Based Displays: Analysis and Solutions," *IEEE Trans. Visualization and Computer Graphics,* vol. 10, no. 2, pp. 177-188, 2004.

[22] A. Majumder and R. Stevens, "LAM: Luminance Attenuation Map for Photometric Uniformity in Projection Based Displays," *ACM Virtual Reality and Software Technology,* 2002.

[23] W.F. Engel, *Shaderx2: Shader Programming Tips & Tricks with DirectX 9.* Wordware Publishing, 2003.

[24] H. Chen, R. Sukthankar, G. Wallace, and K. Li, "Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees," *IEEE Visualization,* pp. 339-346, 2002.

[25] D.S. Tan and R. Pausch, "Pre-Emptive Shadows: Eliminating the Blinding Light from Projectors," *Proc. Conf. Human Factors in Computing Systems (CHI '02) Extended Abstracts,* pp. 682-683, 2002.

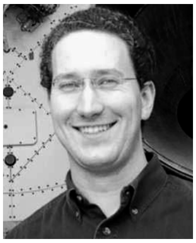

Jay Summet received the MS degree from Oregon State University, specializing in end-user programming, where he was supported by a NASA Space Grant. He is currently working toward the PhD degree at the Georgia Institute of Technology, where his thesis work involves projector-camera systems to enable ubiquitous projected displays. His research interests involve user evaluation, mobile computing, and projection systems. He is a member of the IEEE.

Matthew Flagg received the MS degree in computer science from Georgia Tech in 2003, and specializes in computer vision and graphics. He is currently working toward the PhD degree under the advisement of Dr. James Rehg at the Georgia Institute of Technology. His current research focuses on projector-camera systems for artistic applications. He is a member of the IEEE.

Tat-Jen Cham received the BA and PhD degrees from Cambridge, after which, he worked at Jesus College and Compaq CRL. He is an associate professor in the School of Computer Engineering, Nanyang Technological University, Singapore. He received best paper prizes at PROCAMS '05, ECCV '96, and BMVC '94 based on his research in computer vision and projector-camera systems. He is a member of the IEEE.

James M. Rehg received the PhD degree from Carnegie Mellon University. He is an associate professor in the College of Computing at the Georgia Institute of Technology. He is a member of the GVU Center and codirects the Computational Perception Laboratory. His research interests are computer vision, computer graphics, machine learning, and robotics. He is a member of the IEEE.

Rahul Sukthankar received the PhD degree in robotics from Carnegie Mellon. He is a principal research scientist at Intel Research Pittsburgh and an adjunct research professor in the Robotics Institute at Carnegie Mellon. His current research focuses on computer vision and machine learning in the areas of projector-camera systems, information retrieval for nonindexed data, and object recognition. He is a member of the IEEE.