

# Towards Ambient Projection for Intelligent Environments

Rahul Sukthankar

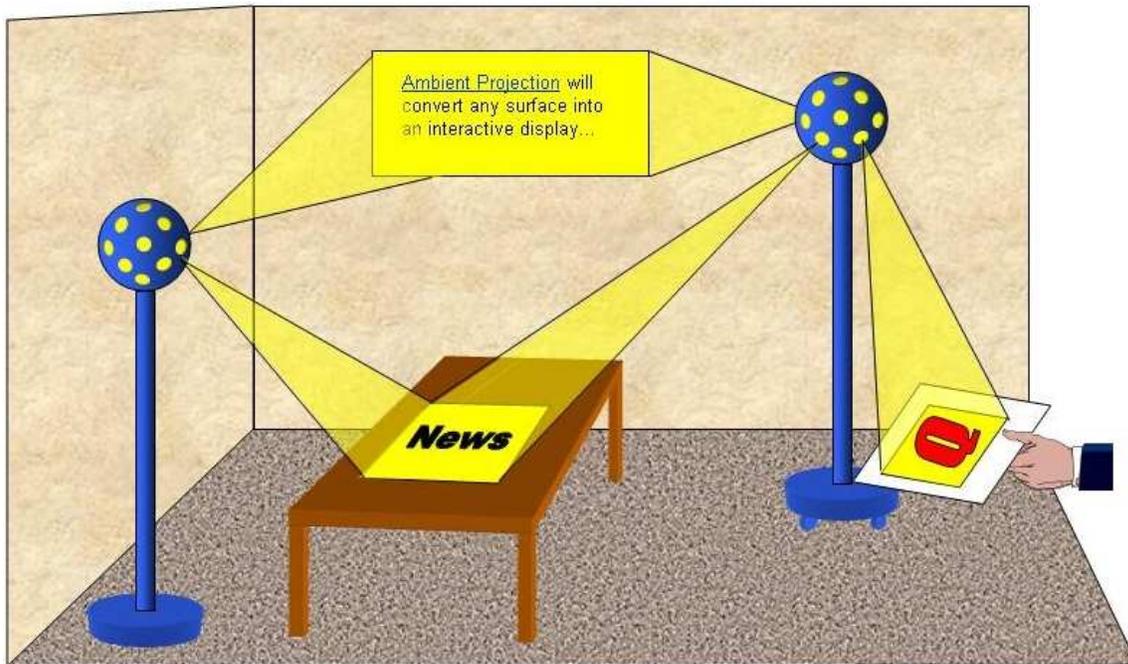
Intel Research Pittsburgh    The Robotics Institute  
4720 Forbes, Suite 410    Carnegie Mellon University  
Pittsburgh, PA 15213    Pittsburgh, PA 15213  
rahuls@cs.cmu.edu

## Abstract

*Traditional desktop computing paradigms provide a poor interface for interacting with intelligent physical spaces. Although handheld devices are an important platform for interface agents, their displays are inadequate for many pervasive computing tasks and need to be supplemented by larger high-resolution displays. We propose the notion of augmenting indoor intelligent environments with ambient projection, where large numbers of projectors simultaneously illuminate the environment from multiple directions — analogous to the way in which ambient lighting permeates a room. Ambient projection could enable any suitable surface in an environment to be employed as a display device. Using such displays, the intelligent environment could present high-resolution information, proactively alert users who are not carrying handheld devices and annotate objects in the environment without instrumentation. Several challenges must be solved before such projected displays become a practical solution. This paper provides an overview of our research in computer vision for enabling interactive ambient projected displays.*

## 1: Introduction

Interface agents serve two important functions in an intelligent environment: (1) they must obtain input from the user, either through an explicit user interface or implicitly by interpreting sensor data; (2) they must present output generated by the intelligent environment in a manner best suited to the user's needs. Handheld devices such as cell phones and PDAs are adequate for displaying short text messages but are unsuitable for high-resolution information, and information kiosks are only useful for user-initiated interactions. We suggest that ambient projected displays are an important addition for indoor intelligent environments. Ambient projected displays will allow multiple users to simultaneously collaborate and interact with high-resolution applications, such as manufacturing schematics, satellite maps, video-conferencing displays, and shared whiteboards. Individual users could use ambient projected displays to customize their office environment (*e.g.*, create a personal photo wall) without altering the physical space. Several challenges must be solved before such projected displays become a cost-effective and practical solution. This paper provides an overview of our research in pervasive projected display technology and discusses future research challenges.



**Figure 1.** We propose the idea of *ambient projection*, where several omni-directional projectors simultaneously illuminate objects in the environment from multiple directions. Interactive projected displays can be created on any suitable surface, including moving, handheld clipboards. Shadows are dynamically eliminated using redundant lighting, large-format displays are created by seamless tiling and projectors avoid casting light into users' eyes by sensing their presence.

## 2: Ambient Projected Displays

The motivation for ambient projected displays is that the increased availability and affordability of high-resolution consumer digital projectors can enable augmentation of indoor environments without additional instrumentation. Where previous researchers have advocated large plasma screens, tiled display walls or steerable projectors, we propose the notion of *ambient projection*: large numbers of projectors that simultaneously illuminate the environment from multiple directions — analogous to the way in which ambient lighting permeates a room. Ambient projection will enable any suitable surface in an environment to be employed as a display device. Not only does this include tables, walls and counter-tops but also objects such as ordinary clipboards held in one's hand, as illustrated in Figure 1.

This vision of an ambient display environment requires solving a number of challenging research problems. First, the various projectors must determine the geometric and photometric properties of suitable surfaces in their environment, including tracking moving surfaces. Second, the projectors must correctly compensate for the problems introduced by oblique projection — either in hardware using specialized optics or in software using pre-warping. Third, the displays must be made resistant to shadows cast by users and other objects in the space. Fourth, different projectors must cooperate to create large, seamless displays. Finally, the system must present natural interfaces for interacting with the agents that drive these displays.

Significant progress has been made on these problems by researchers in several communities

including computer vision [18, 20], computer graphics [6, 15], visualization [3, 10] and ubiquitous computing [12]. The following sections describe some of our computer vision efforts towards building ambient projected displays.

### **3: Oblique Projection**

Projectors are traditionally mounted so that they are as perpendicular to the display surface as possible. This placement minimizes the effects of keystone distortion and ensures uniform illumination of the screen. However, an ambient projection system must often display information on surfaces that are at arbitrary orientations to the projector. This requires us to solve two problems: (1) geometric calibration, which consists of determining the mapping between an oblique projector and the surface to eliminate perspective distortions; (2) photometric calibration, which addresses the issue of non-uniform illumination of the surface and methods for compensating for uneven albedo on the display surface.

#### **3.1: Geometric calibration**

Ambient projected displays demand accurate geometric calibration for two reasons. First, the projected display should be correctly aligned with the physical surface in the scene. For instance, when projecting on to a coffee table, the projected light should fill the tabletop without spilling on to the carpet. Second, when multiple projectors cooperate to illuminate a particular region (whether to form a tiled display or to provide redundant lighting), their outputs must be precisely aligned to avoid undesirable visual artifacts such as double-images or discontinuities.

If the geometric mapping between a projector and a surface is known, one can pre-warp the projected image in a manner that compensates for the oblique placement of the projector. For small numbers of projectors and surfaces, this mapping could be manually specified using appropriate interactive tools. However, it is more desirable to obtain this mapping automatically by observing the scene with a video camera. We apply techniques from multi-view geometry [7] to automatically recover this mapping (even when the positions of the projector and camera are unknown) by projecting a series of light patterns and observing their interaction with the scene. Figure 2 illustrates how our geometric calibration is employed to automatically align a projected image to a presentation screen in the room. Our method is detailed in [18] and related efforts are described in [13, 14].

Similar ideas have been applied in aligning large-format tiled projection displays [5]. One important challenge is achieving sub-pixel alignment accuracy using commodity video cameras. Our geometric calibration system (detailed in [3]) is able to do this by incorporating information from multiple uncalibrated camera views, each of which observes a small portion of the environment. Figure 3 shows a portion of the large-format projected display before and after geometric alignment. The alignment can be performed in a few seconds and a registration accuracy of 0.45 pixel is achieved.

We have implemented a Smart Presentation application to showcase the benefits of automatic geometric calibration. The motivation behind this application is that the presenter might occasionally need to make use of a non-conventional space, such as an oddly-shaped room or hallway, where the projector could not occupy a central location. Smart Presentation (shown in Figure 4) is composed of a laptop, portable projector and a camera. It performs a geometric calibration at the start of the presentation to adapt to the unique constraints of the environment. The user interacts with the presentation display by using a laser pointer to advance slides or sketch on the screen; the motion of the

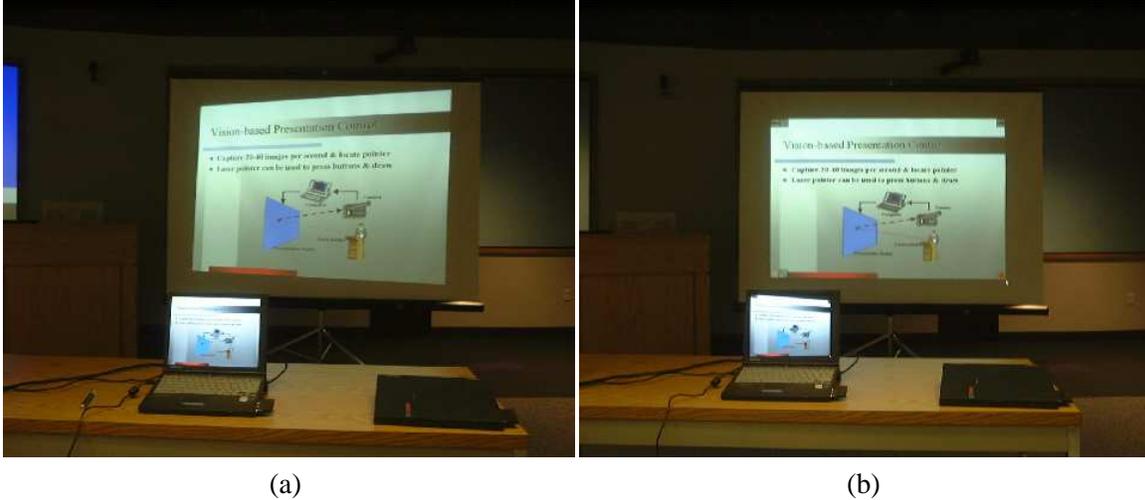


Figure 2. Geometric calibration for oblique projection. (a) An oblique projector normally creates a keystone image on the presentation surface. (b) Automatic calibration is performed using a camera (not visible) to remove the distortion. Note that the image sent to the projector (visible on the laptop screen) is pre-warped in a manner that precisely compensates for the distortions induced by the projector's position.

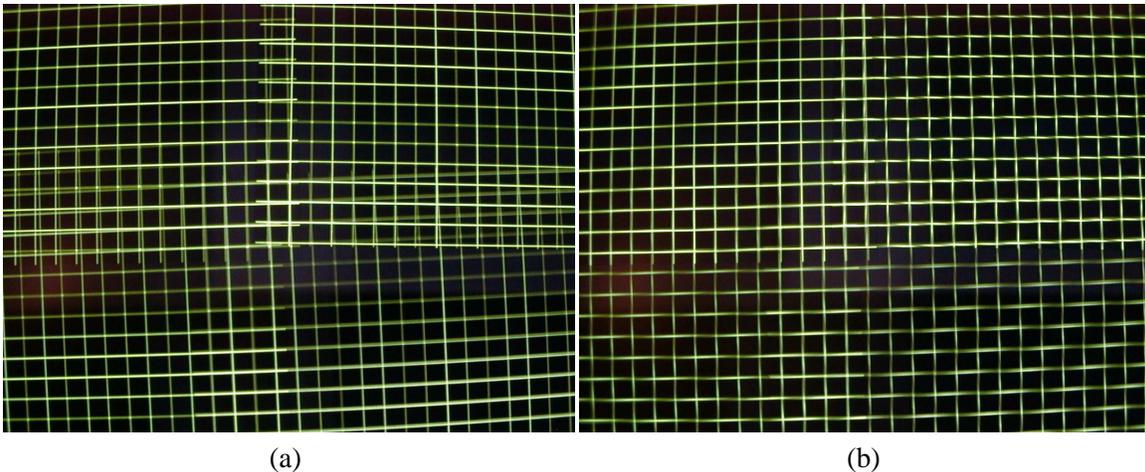
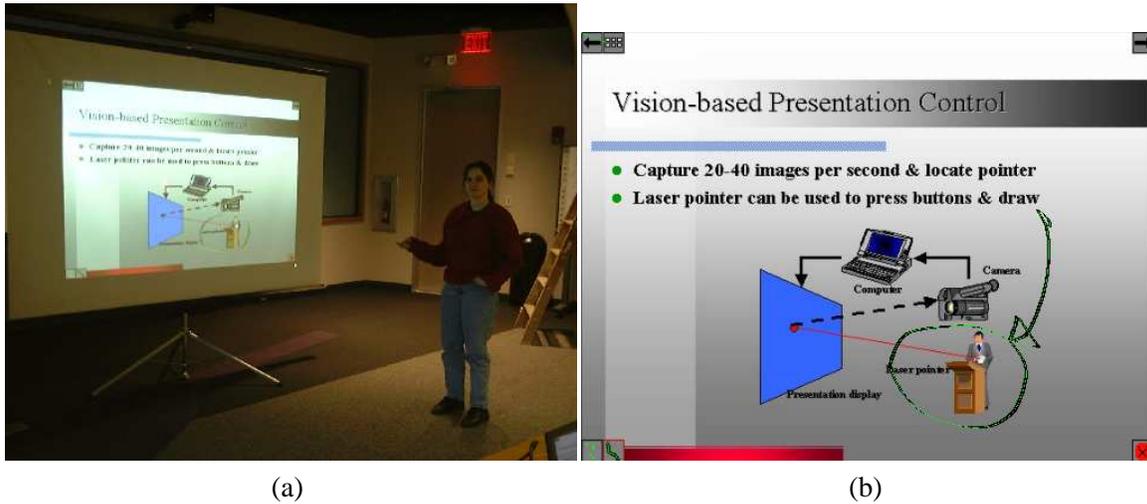


Figure 3. Geometric alignment for seamless tiling of multiple projectors. This figure shows a small portion of a 18'×8' projected display before (a) and after (b) automatic alignment. Information collected from several uncalibrated images taken by a commodity NTSC pan-tilt camera is used to achieve a local registration accuracy of 0.45 pixel.



**Figure 4. Smart Presentation application: the ambient projection system automatically locates the presentation screen and creates a display using an oblique projector (a). Cameras track a laser pointer, enabling the user to control the presentation by “clicking” on projected virtual buttons (b). Decoupling the presentation experience from the computer and projector allows the user to focus more on her audience.**

laser pointer is tracked by the same camera used to calibrate the display. We have also constructed an alternate user interface that allows the user to advance and preview slides on an iPaq handheld.

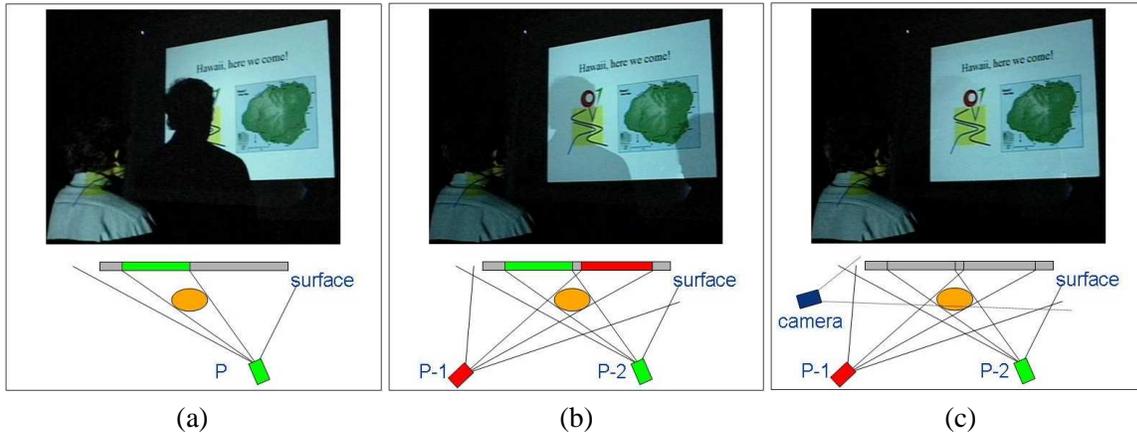
### 3.2: Photometric calibration

Geometric calibration of the projectors ensures that light is projected at the desired location. Photometric calibration addresses the problem of the amount and color of light that should be projected into the scene. It is a key component in the seamless blending of regions where projectors overlap [15], color matching for multi-projector displays [10] and even compensation for stains and colors on the projection surface [11].

Here, we outline the simplest application of photometric calibration: ensuring that a flat surface is uniformly illuminated by an oblique projector [21]. The intuition behind our approach can be stated as follows. If we assume that each pixel is equally illuminated inside the projector,<sup>1</sup> then the non-uniformity in luminance due to oblique projection is precisely related to the relative area of the projected pixel. In other words, pixels that subtend a large area are dim, while pixels that project onto a small area are bright. For the common case of oblique projection onto a planar surface, we can easily determine the ratios of subtended pixels from the geometric calibration, and this enables us to attenuate the brighter pixels appropriately to achieve luminance uniformity over the display surface.

In a more general setup, rather than explicitly performing photometric calibration, an ambient projection system may employ real-time feedback to constantly measure the light at various regions of the scene and adjust projectors to ensure that it is uniform. The next section extends this notion of feedback-based photometric calibration to address the problem of shadow removal.

<sup>1</sup>This is not precisely true for real projectors; however, the inherent non-uniformity between pixel illuminations can be addressed as a decoupled problem.



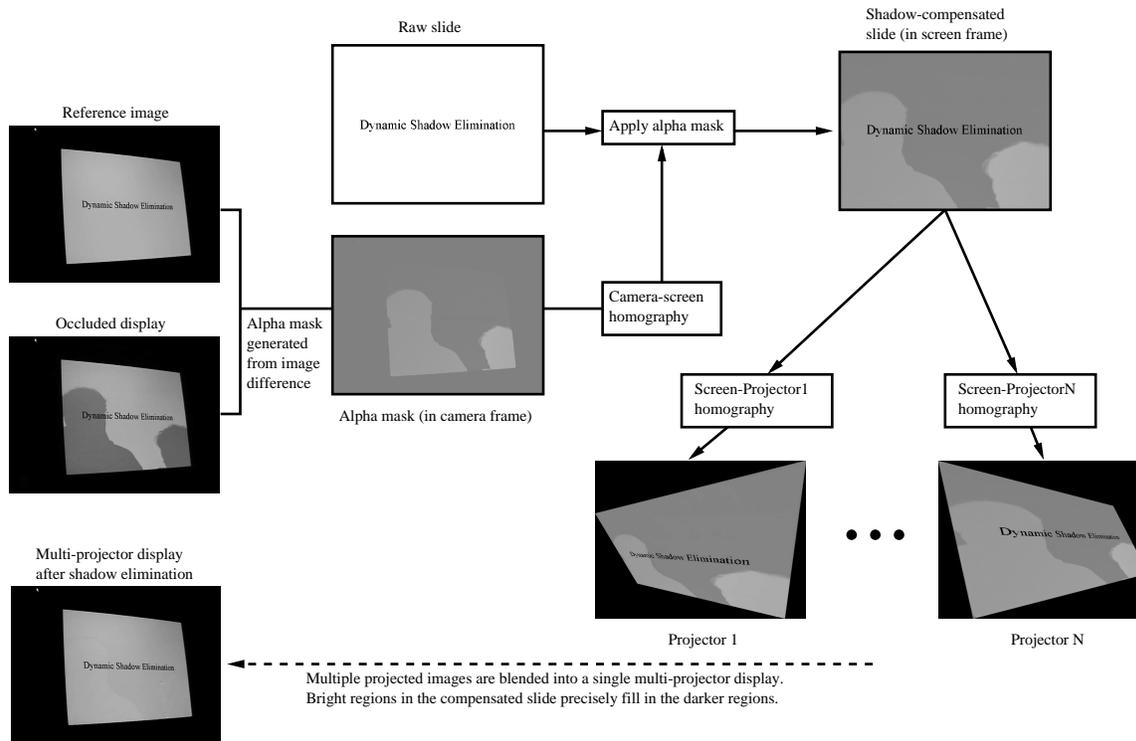
**Figure 5. Dynamic shadow elimination. Left: when an object comes between a projector and the display surface, harsh shadows are created. Center: by employing redundant illumination from multiple oblique projectors, these shadows can be muted. Right: dynamic shadow elimination actively monitors the display surface for potential shadows and adjusts the projector output so that shadows are removed in real time.**

#### 4: Shadow Elimination and Occluder Light Suppression

An obvious problem with front-projection setups is that the light path is blocked whenever an object comes between the projector and the display surface. A related issue is that users are blinded when they walk between a projector and the screen. In traditional presentation environments, these problems can be addressed either by mounting the projector from the ceiling (which reduces the potential for occlusion) or by using rear-projection (which eliminates the problem altogether). The latter is also commonly employed for CAVE environments [4], where the display surfaces are constrained to the walls of the environment.

Unfortunately, neither of these solutions are applicable for ambient projection since rear-projecting onto everyday objects is impossible, and one cannot eliminate the possibility of objects coming between a projector and a surface. Our approach is to: (1) redundantly illuminate display surfaces from multiple projectors; (2) detect the presence of potential shadows and eliminate them by appropriately adjusting the unoccluded projector outputs. This is shown in Figure 5 and the approach is summarized below.

Figure 6 illustrates the algorithm for dynamic shadow elimination. A camera observes the projection surface and detects the presence of shadows (by comparing the current image against a previously-acquired unoccluded reference image or against a predicted image). For example, the image shown in Figure 6 (center left) has two dark regions, corresponding to the two penumbrae cast by one person standing in front of the display (each projector creates one shadow). A pixel-wise image difference between current and reference camera images can be used to locate shadow and over-compensation artifacts. A negative value in a difference image pixel means that the corresponding patch on the screen was under-illuminated in the current image. This information is represented in terms of an alpha mask, which when applied to the current camera image, should bring it closer to the reference image. The alpha mask is updated at every time-step using a feedback system, and is shown to converge to a stable fixed point in a short period of time. Additional details about our system are available in [17] and an alternate approach to shadow removal is described



**Figure 6. System diagram for dynamic shadow elimination.**

in [8].

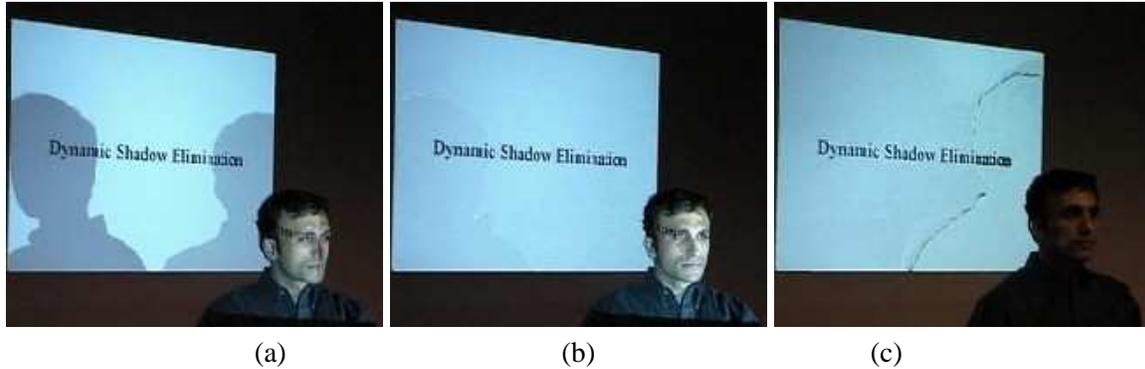
We have also developed a more sophisticated shadow elimination system [2] that simultaneously removes any light falling on occluded objects. This is particularly desirable since it enables the ambient projection system to create displays around users while refraining from casting light into their eyes. The benefits of occluder light suppression can be appreciated in Figure 7.

As an example application for this technology, we have developed a redundantly-lit front-projected poster display that automatically eliminates any shadows created by audience members. To do this we construct an ambient display system composed of two projectors, a camera, and use the shadow elimination and occluder light suppression technique described above. As the audience members move closer to read the “poster” they invariably create shadows by blocking portions of one or both projector beams. These are detected and corrected automatically by the Persistent Poster system.

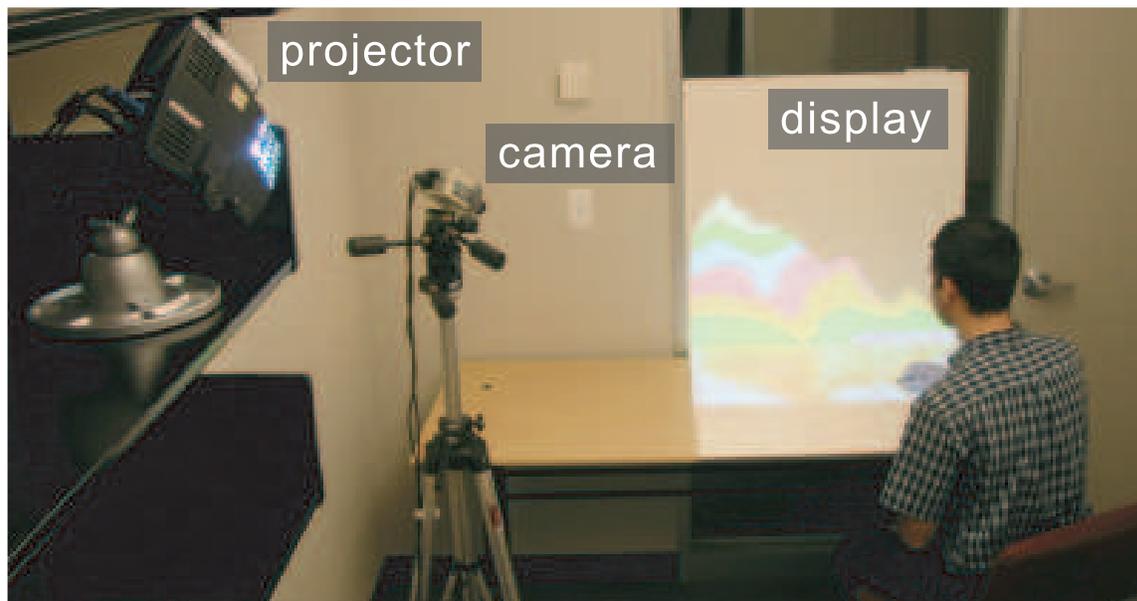
## 5: Multi-planar Displays

Unlike traditional presentation setups, where the projector illuminates a single presentation screen, ambient projected displays often require a single projector to simultaneously project onto multiple surfaces — either to optimize whatever display surfaces are available or to specifically exploit the properties of multi-planar surfaces such as corners and table/wall intersections for innovative applications. Projects such as the Office of the Future [15], Shader Lamps [16] and iLamps [14] have addressed related issues. Our research focuses on displays that consist of intersecting planes, as shown in Figure 8. Details are given in [1].

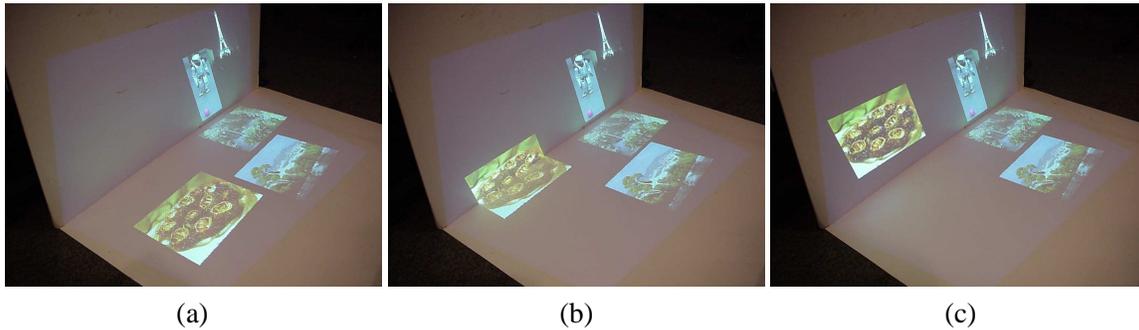
To showcase the benefits of multi-planar displays, we have developed several applications including a multi-planar photo board. This application demonstrates the use of ambient displays for



**Figure 7. Benefits of combining dynamic shadow elimination with occluder light suppression for ambient projected displays. (a) passive redundant projection using two projectors — muted shadows are visible; (b) dynamic shadow elimination alone — shadows have been removed but the user’s face is harshly illuminated with projected light; (c) dynamic shadow elimination with occluder light suppression — light is not projected on the user’s face.**



**Figure 8. A multi-planar projected display is a projector-camera system where one or more oblique projectors display content on several planar surfaces such as a table and a wall.**



**Figure 9. PhotoBoard: an application for ambient displays for environment customization. Note that the images are rendered so that they always appear rectangular, even as they are moved from one display to the other.**

environment customization. The system detects suitable surfaces, such as an area of wall adjoined by a desk, and employs oblique multi-planar projection, as shown in Figure 9. Family photos or artwork can be displayed on the wall, where they do not interfere with the user’s workspace. The user can, at any time, move images to the horizontal workspace for editing. Note that a single projector drives both virtual displays, and that the photos appear rectangular on both surfaces, and even as they move between the two surfaces. This enhances the illusion that the user is manipulating “real” photographs.

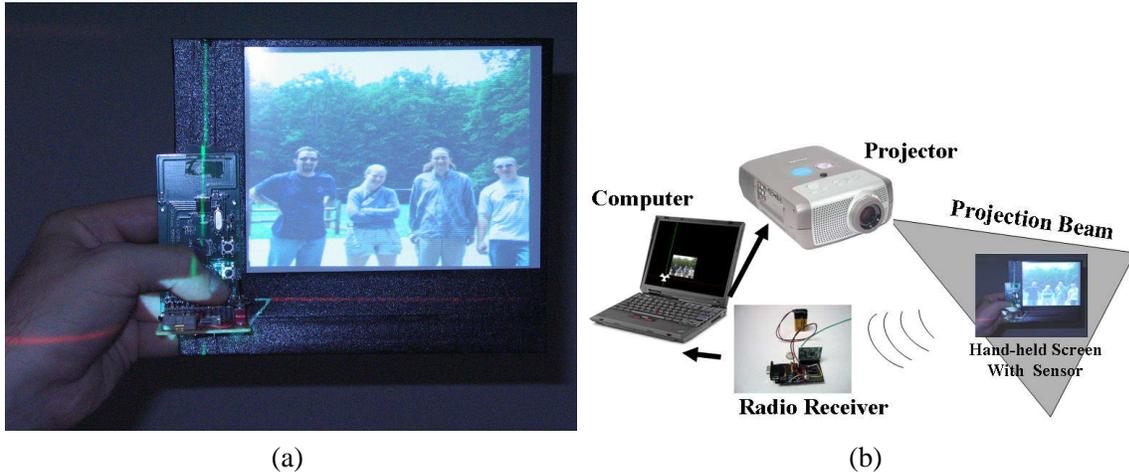
## 6: Projecting onto Moving Handheld Surfaces

An important difference between our vision for ambient projection and earlier pervasive computing display ideas, such as projector walls [9] and CAVEs [4] is that we wish to create displays *inside* the user’s space, not just on the walls. An important step towards this vision is the ability to project on small, moving, handheld objects such as clipboards.

Our approach employs inexpensive photo-sensitive sensors with radio transmitters in conjunction with temporal localization codes broadcast by the projectors in the environment. While we may ultimately prefer a solution that requires no additional instrumentation, our current solution provides robust and accurate localization at little cost. Each handheld object is tagged with active embedded light sensors and determines its own location (in a privacy-preserving manner) by observing patterns being projected in its vicinity. The projected handheld display is shown in Figure 10(a), and the method is summarized in the system diagram shown in Figure 10(b). Additional details are available in [19].

## 7: Conclusion

This paper proposes a novel vision for ambient displays, where several omnidirectional projectors and cameras cooperate to create interactive displays within the intelligent space. These displays can exploit any suitable surfaces in the environment, including moving, handheld objects. These displays are made resistant to shadows through the use of redundant lighting combined with dynamic feedback systems that monitor and eliminate the occurrence of shadows. We describe research that addresses four important problems that must be solved to make ambient projection feasible and demonstrate several prototype applications that showcase various aspects of this idea.



**Figure 10. (a) A handheld screen showing the projected display. The optical sensor is visible near the user's thumb, where faint red and green cross-hairs are visible, showing that the localization is accurate. (b) System diagram. While moving, a sensor on the hand-held object detects location information from the projected temporal code and broadcasts this using radio. This information is received by the ambient infrastructure, which adjusts the projector output so as to keep the projected image precisely on the handheld screen.**

## Acknowledgments

The geometric calibration techniques for oblique projector-camera systems were developed in collaboration with Robert Stockton and Matthew Mullin [18]. Large format display calibration was in collaboration with Han Chen, Grant Wallace and Kai Li [3]. The photometric compensation was developed in collaboration with Himanshu Arora and Robert Ulichney [21]. The shadow elimination and occluder light suppression research was in collaboration with Tat-Jen Cham, James M. Rehg, Gita Sukthankar, Matthew Flagg and Jay Summet [2, 17]. The multi-planar display was in collaboration with Mark Ashdown [1]. The system for localizing handheld projection surfaces was in collaboration with Jay Summet [19].

## References

- [1] M. Ashdown, M. Flagg, R. Sukthankar, and J. Rehg. A flexible projector-camera system for multi-planar displays. In *Proceedings of Computer Vision and Pattern Recognition*, 2004.
- [2] T.-J. Cham, J. Rehg, R. Sukthankar, and G. Sukthankar. Shadow elimination and occluder light suppression for multi-projector displays. In *Proceedings of Computer Vision and Pattern Recognition*, 2003.
- [3] H. Chen, R. Sukthankar, G. Wallace, and K. Li. Scalable alignment of large-format multi-projector displays using camera homography trees. In *Proceedings of Visualization*, 2002.
- [4] C. Cruz-Neira, D. Sandlin, and T. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Proceedings of SIGGRAPH*, 1993.

- [5] T. Funkhouser and K. Li. Large format displays. *Computer Graphics and Applications*, 20(4), 2000. (guest editor introduction to special issue).
- [6] M. Gross, S. Wuermlin, M. Naef, E. Lamboray, C. Spagno, A. Kunz, E. Koller-Meier, T. Svoboda, L. Van Gool, S. Lang, K. Strehlke, A. Vande Moere, and O. Staadt. Blue-C: A spatially immersive display and 3D video portal for telepresence. In *Proceedings of SIGGRAPH*, 2003.
- [7] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [8] C. Jaynes, S. Webb, R. Steele, M. Brown, and W. Seales. Dynamic shadow removal from front projection displays. In *Proceedings of IEEE Visualization*, 2001.
- [9] K. Li, H. Chen, Y. Chen, D. Clark, P. Cook, S. Daminakis, G. Essl, A. Finkelstein, T. Funkhouser, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, J. Singh, G. Tzanetakis, and J. Zheng. Building and using a scalable display wall system. *Computer Graphics and Applications*, 20(4), 2000.
- [10] A. Majumder, H. Zhu, H. Towles, and G. Welch. Achieving color uniformity across multi-projector displays. In *Proceedings of IEEE Visualization*, 2000.
- [11] S. Nayar, H. Peri, M. Grossberg, and P. Belhumeur. A projection system with radiometric compensation for screen imperfections. In *Proceedings of International Workshop on Projector-Camera Systems*, 2003.
- [12] C. Pinhanez. The Everywhere display. In *Proceedings of Ubiquitous Computing*, 2001.
- [13] R. Raskar and P. Beardsley. A self-correcting projector. In *Proceedings of Computer Vision and Pattern Recognition*, 2001.
- [14] R. Raskar, J. van Baar, P. Beardsley, T. Willwacher, S. Rao, and C. Forlines. iLamps: geometrically aware and self-configuring projectors. In *Proceedings of SIGGRAPH*, 2003.
- [15] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays. In *Proceedings of SIGGRAPH*, 1998.
- [16] R. Raskar, G. Welch, and K.-L. Low. Shader Lamps: Animating real objects with image-based illumination. In *Proceedings of Eurographics Workshop on Rendering*, 2001.
- [17] R. Sukthankar, T.-J. Cham, and G. Sukthankar. Dynamic shadow elimination for multi-projector displays. In *Proceedings of Computer Vision and Pattern Recognition*, 2001.
- [18] R. Sukthankar, R. Stockton, and M. Mullin. Smarter presentations: Exploiting homography in camera-projector systems. In *Proceedings of International Conference on Computer Vision*, 2001.
- [19] J. Summet and R. Sukthankar. Tracking locations of moving hand-held displays using projected light. In *Proceedings of International Conference on Pervasive Computing*, 2005.
- [20] N. Takao, J. Shi, and S. Baker. Tele-Graffiti: a camera-projector based remote sketching system with hand-based user interface and automatic session summarization. *International Journal of Computer Vision*, 53(2), 2003.
- [21] R. Ulichney and R. Sukthankar. Correcting luminance of obliquely-projected displays. In *Proceedings of IEEE International Conference on Consumer Electronics*, 2005.