VIRTUAL REAR PROJECTION: IMPROVING THE USER EXPERIENCE WITH MULTIPLE REDUNDANT PROJECTORS

A Thesis
Presented to
The Academic Faculty

by

Jay W. Summet

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
College of Computing

Georgia Institute of Technology
December 2007

VIRTUAL REAR PROJECTION: IMPROVING THE USER EXPERIENCE WITH MULTIPLE REDUNDANT PROJECTORS

Approved by:

Professor Gregory D. Abowd
College of Computing
Georgia Institute of Technology

Professor James M. Rehg
College of Computing
Georgia Institute of Technology

Professor Elizabeth Mynatt
College of Computing
Georgia Institute of Technology

Professor Gregory M. Corso
School of Psychology
Georgia Institute of Technology

Dr. Jeffrey S. Pierce
Almaden Research Center
IBM

Dr. Claudio Pinhanez
T.J. Watson Research Center
IBM

Date Approved: 31 July 2007

To my parents, who made sure I had everything I needed to succeed, and to my sister, le Petit Chaperon rouge.

ACKNOWLEDGEMENTS

Many people have helped me along the way, but my advisers, Jim and Gregory, have always been in the forefront. I am thankful to Jim for introducing me to an exciting research topic and guiding the technical development and Gregory for his advice on evaluation and the PhD program in general. I am especially grateful for the time and effort my external committee members, Claudio Pinhanez and Jeff Pierce spent working with me on my research and the document. Other professors at Georgia Tech have helped me both with my thesis and with other interests. Greg Corso encouraged and improved my user evaluations even before he was on my committee. Beth Mynatt provided guidance on balancing the technology and human side of the research, as well as encouragement throughout. John Stasko provided valuable advice and feedback on publications. Jim Foley, Mark Guzdial and Thad Starner provided advice, support and encouragement in non-thesis areas of my academic career. I wish to express my gratitude to my wife Valerie for her continuous moral support. My peers at the College of Computing provided advice and guidance, and were great people to spend a large part of my life with. To those who went before me, thanks for the encouragement and guidance. To those who traveled with me, thanks for all the fun. And for those who are following, good luck!

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Front projection is an economical method to produce large displays. However, the twin problems of occlusions, which create shadows on the screen, and light projected onto users near the screen, potentially blinding them, makes front projection a poor fit for large upright interactive surfaces. Virtual Rear Projection (VRP) uses multiple redundant front projectors to provide the user experience of using a rear projected display. By using a projector-camera system to mitigate shadows and blinding light, a virtual rear projected display significantly improves upon the user experience of a traditional front projected display, allowing it to replace a rear projected display. In this thesis we characterize the problems caused by shadows and occlusions and develop projection technologies that mitigate shadows and blinding light. We also present a laboratory performance evaluation, and a user evaluation of the technology showing that VRP improves the user experience with respect to traditional front projection.

Chapter I

DISPLAY TECHNOLOGIES FOR INTERACTIVE SURFACES

Front projection is an economical method to produce large displays. Utilizing inexpensive display screens and easily installed projectors, front projection is ideal for information presentation activities. However the twin problems of occlusions and light projected onto users near the screen make front projection a poor fit for interactive surfaces. Occlusions create shadows and projected light may blind users.

Currently, rear projection is the accepted method for delivering digital output on large scale interactive surfaces such as electronic whiteboards due to its ability to produce a shadow free display. Unfortunately, rear projection is expensive. Expensive transmissive screens, the costs for installing these screens, and the cost of space for the projector rooms behind the screens make rear projection installations cost prohibitive.

The cost of digital projectors has fallen significantly in the last decade, and we foresee continued price and size decreases as Micro Electrical Mechanical Systems (MEMS) technology such as Digital Light Projection (DLP) replaces Liquid Crystal Displays (LCDs), and Light Emitting Didoes (LED) lighting replaces the short lifespan and hot-running incandescent bulbs used in current projectors. However, when compared to the space, display surface, and installation costs of a rear projected display, the projector makes up only a small amount of the total cost of ownership. Conversely, the projector cost is a significant percentage of the total cost of a front projected display, which typically includes only a screen and the projector. This projector price trend is already at the point where adding a second projector to a front projection installation is cheaper than building a rear-projection display into a room.

In this document, we use the term *Virtual Rear Projection* (VRP) to refer to systems which use multiple redundant front projectors to provide the user experience of using a rear projected display (See Figure 1). There are three challenges to overcome when using redundant front projectors to build a virtual rear projection display:

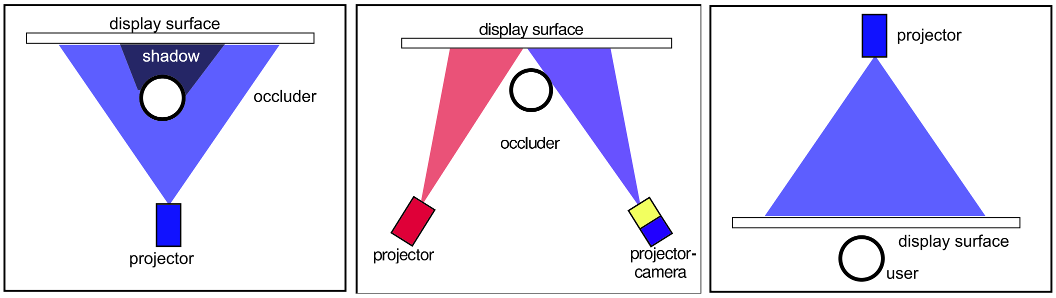

Figure 1: *left to right:* Front Projection, Virtual Rear Projection, Rear Projection

* Calibration - The output of the projectors must be precisely warped to correct for perspective distortion so that the multiple projected images perfectly overlap on the display surface.

* Shadow Elimination - Partial shadows caused by users or objects occluding some of the projectors should be corrected by enhancing the light originating from the unoccluded projectors.

* Blinding Light Suppression - Light that is "blocked" by a user or object before reaching the display surface can be annoying (to onlookers) or blinding (to users), and should be suppressed.

These problems can be solved using computer vision technology, which allows us to calibrate multiple projectors, detect occluders, and prevent shadows and blinding light. *By using a projector-camera system to mitigate shadows and blinding light, a virtual rear projected display improves upon the user experience with respect to a traditional front projected display.*

In this thesis we will discuss the technology developed to provide virtual rear projection displays, an initial evaluation, and plans for future evaluation of the technology. We make the following contributions with this work:

1. Technology development to support passive and active front projection technologies for interactive surfaces (Chapters 3 & 5).

2. A software toolkit (PROCAMS) and example applications enabling others to experiment with virtual rear projection technology and replicate our work without having to re-create our implementation (Chapter 6).

3. User evaluations of passive and active front projection technologies for interactive surfaces in

controlled laboratory experiments (Chapters 4 & 7).

We will present the contributions listed above in detail in the remainder of this document. In Chapter 2 we will discuss work related to the technology of virtual rear projection along with application areas for large interactive surfaces. In Chapter 3 we will discuss passive projection technology used to improve the front projected experience, while Chapter 4 will discuss a laboratory evaluation of this work [63]. Chapter 5 discuses technological enhancements made to improve passive virtual rear projection as a result of the initial user evaluation, and Chapter 6 describes the PROCAMS toolkit. Chapter 7 reports on the evaluation of the technology, consisting of controlled laboratory studies of user preference and behavior. Finally, Chapter 8 concludes with a summary of findings, suggested directions for future work, and recommendations for implementers and system builders.

1.1 Overview of Display Technologies

The technology developed and evaluated in this thesis is a display, or output, technology which projects images and graphics for users to view. Specifically, it is a projection technology, as opposed to a direct image or eye-coupled display. The following sections give an introduction to these different display technologies and highlight their relative benefits and drawbacks.

1.1.1 Direct Image Display Technologies

A direct image display is one where a physical object emits or reflects light in a computationally controlled way to generate a user perceivable image. Although a piece of printed paper from a teletype or printer is a form of static direct image display, we are limiting this discussion to displays that have the ability to dynamically update the displayed image.

The earliest widely used computer controlled direct image display was the CRT, or Cathode Ray Tube, monitor. The CRT operates by directing a ray of electrically charged particles via computer controlled electromagnets to illuminate luminous phosphors on a screen to produce text and graphics. By using multiple colors of phosphors, multi-color images could be displayed. The ray of electrically charged particles must travel through an evacuated vacuum, and have a minimum beam length based upon the size of the screen, so as CRT screens increase in size, they become deeper. A material of suitable strength and air-tight properties to maintain the vacuum (typically glass) is used for their construction, which makes CRTs large and relatively heavy. However, after several decades of development, and the economies of scale generated by the production of billions of televisions, CRT technology is mature, and CRT tube displays can be manufactured relatively cheaply.

Liquid Crystal Displays (LCDs) are light filters that can be electrically controlled. By selectively passing or blocking light, LCDs can display graphics or text. Low resolution special purpose LCDs are often used in digital watches or appliances when an inexpensive and low-power display is needed. These passive LCD displays typically do not generate any of their own light, instead reflecting or absorbing ambient light to create their display. Because it takes very little power to turn an LCD's filter on or off, passive LCDs can be powered by a battery for several years, but they are not readable in low light situations. Most current computer displays combine a LCD panel with a backlight to produce an image comparable to a CRT display. The backlight produces white light that is filtered by an array of very small LCD pixels, and then passed through filters of various colors, allowing the LCD display to generate a full color image. Because each layer of an LCD display (backlight, LCD matrix, color filters) is relatively thin, LCDs are much thinner and lighter than a comparable CRT tube based display. Manufacturing LCD display panels requires a complex assembly line similar to semiconductor manufacturing, and the size of the produced display is limited by the glass substrate size that the assembly line or plant can process. When first introduced, LCDs were physically small and had low resolution, but as demand grew and the economies of scale increased, glass substrate sizes and LCD sizes increased [45]. Today LCD displays (especially Thin-Film Transistor LCDs, or TFT-LCD) have overtaken CRTs as the computer display of choice. LCD displays are now commercially available in sizes that range up to 65 diagonal inches and HD resolutions.[1] Larger LCD displays have been demonstrated and will eventually reach the consumer market. LG Philip's has demonstrated a TFT-LCD display panel that measures 100 diagonal inches which used the "maximum efficiency" of LG Philips' seventh generation manufacturing line [31]. The current world leader in LCD display panel size is Sharp, which showed off a 108 diagonal inch LCD display (1920x1080 pixels, or High Definition resolution) at CES in 2007.

Light Emitting Diodes (LEDs) are solid state devices that convert an electric current directly

[1] In 2007, a consumer television with a 65 inch LCD display costs over $7000, consumes 610 watts and weighs 140 lbs.

into light. One of the first consumer displays made which used LEDs was the 1970 Pulsar digital watch [57]. Because LEDs use more power than LCDs, these early digital watches had a button that the user had to press when he or she wanted to view the display. LCDs quickly overtook LEDs for watch displays, and LEDs were not used for high resolution displays until decades later, when OLED (Organic Light Emitting Diode) displays entered the market. OLED displays have higher power efficiency, and can produce bright displays with higher contrast than an LCD matrix and backlight [48]. They are currently only economical to produce in small (2 to 4 inch diagonal) form factors, and are used in consumer devices such as cameras and cell phones. Manufacturers are currently attempting to modify the active matrix (TFT) substrate production technology used for LCD displays to make them compatible with OLED displays, allowing OLEDs to use the same production hardware that has received heavy investment for the production of large sized LCD screens. Samsung Electronics LCD R&D Center has demonstrated a 14.1 inch OLED display in the laboratory [16].

Plasma displays, or Plasma Display Panels (PDP) contain tiny chambers of inert noble gases sandwiched between pieces of glass. To produce an image, the chambers are electrically charged and converted to plasma, which excite phosphors and release light. Each pixel is made up of three gas chambers (for the three primary colors). PDP's can produce brighter displays than LED panels, and have been manufactured as large as 103 diagonal inches unveiled at the 2006 International Computer Electronics Show (CES), although consumer plasma TV's are only easily available up to 65 diagonal inches.[2]

Surface-conduction Electron-emitter Displays (SED), a prototype technology that is close to being marketed, are a mix of cathode ray and plasma display technology. Instead of having a single bulky ray tube for an entire screen, a SED display has an individual ray tube and phosphor screen for each color sub-pixel. Unlike LCDs and PDPs, the emitter matrix of a SED displays can be manufactured using a technology similar to ink-jet printing, theoretically allowing large displays to be manufactured more cheaply than LCDs or PDPs. The SED display has only a single electron emitter per color sub-pixel, and is a simpler version of the more general Field Emission Display (FED) technology. True FED displays use multiple redundant nano-wire emitters per pixel, but are

[2] In 2007, a 65-inch plasma TV cost over $8000, uses 675 watts, and weighs 70 lbs.

more difficult to manufacture with current technology. Because FED and SED displays require an extreme vacuum, they must be manufactured on and protected by rigid surfaces, usually glass, which prevents them from being used to build flexible displays. SED displays have been demonstrated in laboratories to have lower power consumption than LCD and PDP displays [78]. Currently, no SED displays are available to consumers.

A special type of direct image displays which use rear projection, including self-contained rear projection TV's, is discussed in Section 1.1.2.

All of the above mentioned display technologies require power to maintain an image, and with the exception of reflective LCD displays, are emissive, in that they generate light to produce a visible image. Other display technologies exist that are bistatic, and reflective, which means that they have two (or more) stable states and reflect different amounts or colors of light depending upon their current state. These displays can be changed by applying an electric charge but will maintain their current state without power. The best known of these technologies is electronic ink, which uses tiny magnetically charged spheres that are half white and half black to produce gray scale displays [21, 26]. Current e-ink displays use conventional TFT arrays to selectively distribute charge (flipping the orientation of the spheres from white to black), limiting the size of such displays to that obtainable by conventional TFT/LCD manufacturing processes. To date, high resolution (1024x768) glass substrate e-ink displays have been used in e-book readers (such as the Sony Librié and Reader, the Hanlin eBook, and the iRex iLiad) and a plastic (although not flexible) substrate low resolution display was used on the Motofone F3. These displays are smaller than 8 inches diagonally, but if printed driver electronics can be brought to market, eInk technology could potentially be used to produce rollable wall sized displays. Flexible organic semiconducting polymers have been used to manufacture flexible active matrix (TFT) arrays and eInk displays with 50ppi resolution in small quantities [9]. Current laboratory efforts in producing flexible displays using roll-to-roll manufacturing have been limited to monochrome displays with low resolution (10-50ppi) [43].

1.1.2 Projection Technologies

Projection display technologies create an image on a passive screen using projected light. Front projectors bounce light off a reflective screen, while rear projection displays transfer light through a transmissive screen. Ignoring obsolete technologies, such as Eidophor oil-film projectors, four major technologies: Cathode Ray Tubes, Liquid Crystal Displays (LCDs), Liquid Crystal on Silicon (LCOS), and Digital Micro-mirror Displays (DMD's) are used to produce projected images commercially, while a fifth, laser projection, is being developed in laboratories.

Cathode Ray Tube projectors simply take the light produced by a CRT and focus it through a projection lens system. Typically three CRTs are used, one for each primary color to produce more light and a brighter image. Although large and bulky, these projectors have long lifetimes because of the longevity of the base CRT technology. CRT projectors produce light as part of the image generation process, but the other three types of projection technologies (LCD, DMD, and LCOS) simply modify existing light to produce an image. LCD, DMD, and LCOS projectors typically use an incandescent or high intensity gas-discharge bulb to produce light, which is then filtered by the imaging element (LCD, DMD, or LCOS imager) to produce an image.

LCD projectors focus the light from a high-intensity discharge or incandescent lamp through a liquid crystal display, which modulates the light forming an image, and then out to the display surface via a projection lens. Because the maximum efficiency of an LCD display is 50%, LCD projectors are inherently less bright than their LCOS or DMD counterparts. Digital Micro-mirror Displays (DMD) (a.k.a Digital Light Projection or DLP) and LCOS projectors selectively reflect light from an imaging chip towards the screen or a trap within the projector. In a DMD projector, the imaging chip has millions of tiny mirrors manufactured using MEMS techniques. Each mirror can be electrostatically controlled to direct light towards a trap or the screen on a per-pixel basis. LCOS projectors use a liquid crystal to modulate light by reflection instead of transmission. An LCOS imaging chip reflects light based upon the state of the individual LCD pixels on its surface. Both chips are manufactured on a silicon substrate using processes developed in the semiconductor manufacturing industry.

All of the previous projection technologies use white light sources which produce many wavelengths of light. By contrast, a laser video projector uses one or more lasers to produce a coherent beam of light that is rapidly raster scanned across the display and amplitude modulated to produce a raster image. A similar system is used at laser light shows but they typically scan a laser in a vector pattern and do not modulate the light output. By using three different lasers to produce primary colors (red, green and blue) and mixing their intensity at each pixel location, a laser projector can produce the illusion of a full color image. These projectors can scan their laser beams by using mechanically oscillating or rotating mirrors, or with smaller MEMS mirrors [80]. Because they use coherent beams of light, laser projectors do not need projection optics to focus the image on the display screen. This means that they could theoretically be miniaturized much smaller than other optical projectors, and that they have no limit to their depth of focus. One drawback of laser projectors is that their coherent light causes a subjective speckle pattern when it illuminates any surface that is not perfectly smooth. Because any variations in the surface that are larger than one wavelength of the laser light (typically 300-600nm) causes speckle, it is incredibly difficult to produce a display surface that completely eliminates this visible speckling pattern.

Any of these projectors can be used in a front or rear projection configuration. Typically, in a front projection configuration, the projector and screen are separate. The screen can be rolled or folded for transport. Specialty paint can be applied to a suitably flat wall to produce a high quality projection screen. In some cases, a light colored wall or cloth is used as an ad-hoc projection screen with no modification. The projector may be mounted to a ceiling or wall, on a portable tripod or stand, or simply placed on a suitable table or bookcase.

Rear projection configurations typically come in one of two configurations: permanent, or a mobile self-contained unit. In a permanent installation, a transmissive screen is built into a wall, and the projector is permanently mounted behind the wall. In a self-contained unit, the projector and screen are built into a rigid housing. Some self-contained units have a hinge or folding mechanism to allow them to be folded for transport. Rear projection DLP televisions fall into the self-contained rear-projection category. Because they require a light path from the projector to the rear surface of the screen, rear projection displays are typically thicker than their flat panel (LCD, PDP) counterparts, although new optical systems (including aspheric mirrors and defractive gratings) are increasingly reducing the depth needed behind a rear projection screen for the beam path.

LCD, DMD, and LCOS projectors have traditionally used high intensity gas-discharge of incandescent bulbs to produce lights. These lights are power hungry and convert a large amount of the power they consume into heat instead of usable light. This extra heat has lead to the need for office projectors to have exhaust fans to cool the projectors, releasing heat (and noise) into the environment. As the efficiency of light emitting diodes increases, they have begun to replace traditional gas-discharge or incandescent bulbs, first in decorative lighting applications, and lately in small low-powered projectors. These LEDs replace the traditional incandescent or gas-discharge light source in a projector, although the imaging chip (typically DMD) remains the same. Currently LED projectors are relatively low power (20-50 lumens) when compared to their incandescent counterparts (where 2000 lumens is common in a consumer model) and are limited to short throw (small screen) and dark room applications. LED light projector models include the Toshiba TDP-FF1AU, the Mitsubishi PK-20 and the Samsung SP-310. LED projectors have the advantages of the ability to turn on and off instantly, reduced power consumption per lumen, less excess heat production, and longer bulb life. The reduced power requirements of LED light sources allows projectors to be battery powered and operate almost silently with slower and quieter exhaust fans. As LED lighting technology improves, more and more DMD/DLP projectors will use LEDs for their light source.

1.1.3 Eye-Coupled Display Technologies

Eye-coupled display technologies are worn by the user on or over their eye, as opposed to user worn or held direct image displays such as a digital watch face or video iPod. The eye-coupled display presents the user with a display that appears to be in the distance and relatively large, although the actual optical hardware is small enough to be head mounted. Eye-coupled displays can produce a large image that moves with the user. Two basic technologies exist for eye-coupled displays. The first is to project an optical image (focused in the distance) before the user's eye via a prism or other optical element. When the user focuses at the appropriate distance to focus the display, they view it as if it were in physical space before them [10].

The other technique, called a virtual retinal display (VRD) or retinal scan display (RSD), is to trace an image onto the users retina using a laser. VRD technology has much in common with the laser video projector, in that it requires little optics other than a scanning mechanism and beam amplitude modulator[15]. Because the image is traced on the users retina, it always appears to be in focus regardless of what the user is looking at. VRDs can produce a bright image with relatively little power, but suffer from a public relations standpoint due to their use of lasers directed into the eye.

Eye-coupled displays can be used to create the illusion of a large display that floats at a distance before the user and moves with the user's head. Alternatively, if the motion of the user and their head can be tracked with sufficient speed and accuracy, an eye-coupled display can be used to generate an image that moves as the user does, making the image appear to be fixed on a specific object in the real world. In such an augmented reality approach, an eye coupled display could be used to emulate any number of direct view displays scattered throughout the environment. In the long term, this may be the easiest and most cost effective method for achieving the effect of distributing displays throughout an environment, but current tracking technology is unable to work quickly or accurately enough to maintain the illusion. Additionally, eye-coupled displays have not yet been miniaturized and produced cheaply enough to be widely accepted.

1.2 Why Virtual Rear Projection?

Virtual rear projection provides space and cost benefits over traditional fixed rear projection installations. Even in new construction, rear projection is an expensive option. The average cost to build a square foot of office space in the United States is $77 USD [72]. A five foot (1.52m) wide rear-projection surface using traditional projectors will require a clearance of about three feet (0.91m) behind the screen, even when using a space saving twin mirror design. This fifteen square foot (1.39 m^2) area behind the screen will cost $1155 USD, approximately the cost of an inexpensive projector. A rear projection display also requires a specialized projection surface, which can cost thousands of dollars, significantly more than an equivalent front projection surface. In addition, these rear projection screens are usually mounted in custom built walls, requiring specialized construction. Compared to the minor ceiling mounting required by most front projection systems (which can usually be accomplished by an organization's existing facilities personnel), installation of a fixed rear projection display can be an expensive proposition. As current trends continue, and projector prices continue to decline, the cost of a virtual rear projection system will become significantly cheaper than a comparably sized rear projection system.

In an effort to bring interactive rear projected displays to the market, products such as the Xerox/Liveworks LiveBoard [12] and the rear projected SmartBoard [58] introduced portable "rolling cabinet" rear projection displays. Although they enjoyed limited success, these products were large and bulky, limiting their portability and dominating the spatial layout of rooms in which they were placed. They are currently being replaced with display solutions using warped front projection such as the 3M IdeaBoard [25] and touch sensitive overlays on plasma displays. Large format displays, such as plasma, LCD, and thin format DLP rear projection compete with virtual rear projection for producing a large format display suitable for interactive use. Plasma displays are still much more expensive than an equivalent dual projector display and are limited in size. Although there have been trade show demonstration models built with diagonal sizes of up to 103 inches, these behemoth plasma displays are still impractical to build in quantities due to the economic difficulties of scaling production lines to produce them at a price consumers are willing to pay.

Thin format DLP rear-projection displays, which use a form of warped (rear) projection to achieve thinness as small as seven inches, cost about the same as two projectors, and weigh around two hundred pounds. Because they are not subject to the same production line scaling and yield issues as plasma or LCD displays, these rear projection DLP based displays are the closest competitor to virtual rear projected displays.

All plasma, LCD, and rear projected displays have issues of size, weight, and cost which can be solved by a VRP display. Large plasma, LCD and rear projection DLP displays weigh several hundred pounds, and must be transported in large crates, leading to significant shipping and installation costs. Additionally, there are some public environments (subway stations, parks, etc.) where an expensive and relatively fragile display accessible to the public would be in danger of being stolen or vandalized. A virtual rear projection display can be shipped as a rolled screen (optionally touch sensitive) and two projectors. The system can be mounted by a single workman with the projectors located safely overhead. An interactive display surface can be quite rugged, and is much cheaper to replace than a Plasma, LCD or DLP rear projection display. In the next chapter we discuss research using interactive displays, advances in front projection technology, and work related to virtual rear projection.

Chapter II

RELATED WORK

The traditional vision of pervasive computing assumes that computer displays are scattered throughout the environment in a variety of sizes [75]. The displays are assumed to react appropriately to a user's actions and needs, either through ubiquitous sensing or by being interactive. Some commercial products such as the LiveBoard [12, 36] and SmartBoard [58] deliver on the promise of Weiser's yard scale interactive displays which have both input and and output capabilities, but these large scale interactive displays have not enjoyed wide deployment to users' homes and offices. When compared to their smaller counterparts, the inch sized displays of cell phones and PDAs and the foot sized displays in laptops and computers, large scale displays are much less pervasive. The economic reality is that current large scale interactive displays are difficult for one person to move or install by themselves, and cost much more than inch or foot scale displays. Additionally, many of the large scale displays that do exist are used to display and interact with applications originally designed for the smaller displays of personal computers. Few applications that are specific to large scale displays exist outside the laboratory. The following sections in this chapter will examine research on large displays and their applications, work on improving projected displays, and research closely related to Virtual Rear Projection that involves the reduction of shadows or elimination of blinding light.

2.1 Large Displays & Applications

Perhaps one of the earliest interactive large display applications was Myron Krueger's projected "Videoplace" artwork [33]. In addition to artistic endeavors, many applications for large displays have been prototyped by researchers. Work on electronic whiteboards [46], digital tape drawing [3], and focus plus context displays [7, 6] have demonstrated potential application areas suited for a single user, wall sized interactive display. Collaborative applications that have been prototyped on tiled display walls include genomic data visualization, iso-surface extraction, and collaborative control rooms[73]. Additionally, remote meetings and video conferencing has been widely investigated as an application area for large displays [70, 28, 27, 54, 8, 79], as have whiteboards for collaborative meetings [12, 50] and design sessions [32].

Recently, researchers have demonstrated some benefits that large displays provide. Tan *et al.* found that displays that filled a larger portion of the user's field of view (around 100 degrees) resulted in users having better performance in 3D navigation tasks [67, 66]. Additionally, physically larger displays, even when viewed at the same angle as a smaller display of identical resolution, improve performance on spacial tasks [68]. They can also be used as an alternative to head mounted displays for virtual reality simulations [49].

Microsoft Research has examined issues that arise from the use of large scale displays by a single desktop computer user. These issues include losing the cursor, problems with information access across large spaces, and window and task management problem [55]. MacIntyre *et al.* and Robertson *et al.* have demonstrated the use of miniature versions of windows displayed on a large peripheral display [37], or on the side of a large display or a secondary monitor [56] for task management purposes. Research on the Stanford Interactive Mural has developed interaction and screen management techniques [18] for wall sized interactive surfaces.

Much of the research work on large displays has used rear projected displays located in the researchers' laboratories, the equivalent of which is still not commonly found in the real world. Researchers who used front projection to prototype large upright displays ran into problems with shadows. For example, focus plus context displays that use a front projector have been "tilted slightly" so the projector can be ceiling mounted to "keep the [sitting] user from casting a shadow on the projection screen"[7]. Many researchers choose to use rear projected displays to avoid the issue of shadows. For example, the builders of the Stanford Interactive Mural decided that, "to avoid self-shadowing that would result from interacting with a front-projection system, we used rear projection" [19].

2.2 Projected Display Technology

Using projection to create a large display has size, weight and cost advantages over using a direct image display. Projectors are the most cost effective way to generate a large image, but even in a rear projection configuration, projected images can suffer from problems not related to shadows and blinding light. One of the primary applications of rear projection in the laboratory is the construction of tiled wall sized displays. There are four main challenges to building display walls from tiled projectors. First, as the number of projectors used becomes larger than can be driven by a single computer, the graphics pipeline must be distributed among multiple synchronized nodes. PixelFlex, a front projected reconfigurable display at UNC Chapel Hill, used a single SGI InfiniteReality 2 system with dual 4 channel graphics pipes to drive eight projectors [79]. Both the Princeton display wall [34, 73] and the Stanford Interactive Mural [23, 19] used clusters of computers supporting distributed rendering. The Stanford work led to software systems for distributed OpenGL rendering such as WireGL/Chromium [22, 24].

Second, the output of the projectors must be precisely calibrated so that the images they project are correctly aligned. Individual projectors can be physically aligned using motorized or manual gimbals, but this requires extensive calibration work. For example, to calibrate the Stanford Interactive Mural, each of the twelve projectors required an hour to align, and this alignment would only be "stable over several weeks" [19]. The majority of tiled displays are now only aligned roughly by hand, and a combination of computer vision and software image warping is used to correct the alignment. Typical techniques involve calculating a projective transform (homography) for each projector using computer vision [11, 79]. Although we do not tile individual projected displays, Virtual Rear Projection uses a similar technique to overlap them, described in Section 3.3.

Third, the outputs of each projector must be combined in such a way that the display is seamless. The two general techniques are to abut the two images exactly (perhaps at pixel boundaries) or to merge overlapping images together smoothly by ramping their intensity. The Stanford Interactive Mural abutted displays using physically taped masks near the screen [19]. Early versions of the Princeton Display Wall used shadow masks located near the projectors to cause penumbral shadows that optically ramped the images together on the screen [35]. The PixelFlex display performed this ramping using alpha blending in software [79]. The Active form of Virtual Rear Projection (AVRP) uses software alpha blending when joining images from separate projectors at a seam as described in Section 5.3.3.

Finally, the brightness and color of the individual tiles must be calibrated to present the illusion of a single uniform display. Majumder found that many projectors exhibit spatially varying levels of intensity, usually delivering a brighter image in the center and suffering from intensity falloffs near the edges [40, 38, 41]. These intensity variations were larger than color (hue) variations and although these variations do not seriously detract from the perceived image quality of a single projector's image, they contribute to the visibility of seams in multi-projector tiled displays. These intensity variations are even more pronounced when using projectors in the off-axis configuration of Virtual Rear Projection, and we use a simplified version of Majumder & Steven's Luminance Attenuation Maps [39, 42] to reduce the visual impact of intensity variations as described in Section 5.3.2.

Although rear projection displays require a relatively specialized transmissive display surface, front projectors can be used to project images on any surface that is sufficiently reflective. Obviously, a specialized surface that is planar and consistently reflective (such as a white wall or specialized reflective screen) provides the best image, but researchers have developed methods for improving images projected onto surfaces with uneven textures or backgrounds. Projecting onto a non planar surface requires geometric correction[5, 53]. A few high end projectors from NEC include the ability to warp their projected image using built in 3D graphics hardware after manual calibration [76]. Projecting onto non-regular, colored, or textured surfaces requires photometric correction, which calculates a pre-corrected image to project that corrects for the pre-existing color or texture on the display surface [14, 17, 47].

Because they are not tied to a specific projection screen, front projectors can be dynamically repositioned, either by users, or via motors under computer control. Projectors which can shift the location of their projected image are called *steerable.* Steerable displays allow a single projector to project images onto many locations through a room. These images can be used as independent displays, or to project graphics that seamlessly integrate with and augment the environment [52]. The Everywhere Displays projector is a steerable projector augmented with a MIDI controlled pan-tilt mirror and computer controlled focus [51]. It can compensate for shadows by detecting when users were blocking its projection path, and move the projected image to an alternative location. PixelFlex used an array of eight steerable projectors to build a tiled display wall that could be dynamically reconfigured to change the aspect ratio and resolution [79]. Although PixelFlex did not detect users and was intended as a non-interactive display, it could be configured to produce a redundantly illuminated display similar to Passive Virtual Rear Projection (PVRP). The calibration of steerable projectors can be simplified if the projectors are physically rotated around their center of projection, instead of using a pan-tilt mirror to steer the reflected image [44]. The combination of geometric correction, photometric correction, and steerable projectors allow front projected displays to be placed on arbitrary surfaces in an environment and give more flexibility about where to position a display than competing display technology.

2.3 Shadow Elimination and Blinding Light Suppression

The use of projector camera systems to improve upon front projected solutions by eliminating shadows and eliminating blinding light is a relatively new area of research. Desney Tan demonstrated how to use IR lights and camera to detect a person and create a black "mask" over the projected graphics [69], which creates a pre-emptive shadow that eliminates the blinding light from a projector. A similar technique is used by a commercial appliance from iMatte, sold as an add-on for existing projectors. These systems suppress the blinding light, but leave a shadow on the display surface. They are useful for some applications (such as giving a presentation) where the user is mobile (i.e. can move the shadow away from the screen if needed) and does not need to interact with the display. However, for other applications where the user must interact with the display (e.g. writing on an electronic whiteboard or selecting links in a web browser) the shadow cast on the screen by the user's body is problematic and coping behaviors become evident [63].

The technology of virtual rear projection, or the use of multiple projectors to provide a robust display in the face of occlusions, has been explored by a small community of researchers. Previous research at Compaq Labs and Just Research by Rahul and Gita Sukthankar and Tat-Jen Cham, in conjunction with Jim Rehg, introduced the idea of using multiple projectors and a camera to correct shadows on a display [60]. Their system used a camera which assumed an unoccluded view of the display, and while correcting for shadows, would project additional light onto occluders, potentially blinding the user if they turned to face the projectors. A later extension to their work "polled" the projectors to determine which projector was being occluded [62, 61]. It would then reduce the light from occluded projectors, eliminating the blinding light on the occluder. This system also assumed an unoccluded view of the display surface and worked at lower than interactive frame rates. A laboratory evaluation of these systems is presented in Section 5.4 and a TVCG journal paper [64]. Both of these systems suffered from two drawbacks. First, they required that the camera have an unoccluded view of the display surface to detect shadows. Second, they could only display pre-selected graphics, which made them unsuitable for interactive displays.

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 [30]. 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 *et al.* enhanced this work to increase the speed to approximately nine frames per second, by updating bounding regions instead of individual pixels [29]. Similar to the Shadow Elimination and Shadow Elimination with Blinding Light Suppression techniques described in Sections 5.1 & 5.2, their system requires numerous frames to converge to a stable display. Their updated system still requires that cameras have an un-occluded view of the screen, and does not eliminate blinding light. Recent work by Audet and Cooperstock demonstrates a system to eliminate blinding light and correct shadows on the display by using a pair of calibrated stereo cameras to detect occluders [2]. Because they are calculating the location of occluders in 3D, their cameras and projectors must be fully calibrated in 3 dimensions, unlike AVRP which only requires a four point projective calibration between projectors and camera. They calculate a rectangular bounding region for each occluder from the viewpoint of each projector and use this to generate shadow masks. Their system works well for occluders moving in a room, but was not demonstrated for users approaching close to an interactive display. All of the previous work described here in the areas of shadow elimination and blinding light suppression has been entirely technical, and involved no user evaluation. The following chapters cover the technical details involved in the implementation of Warped Front Projection, Passive Virtual Rear Projection, and Active Virtual Rear Projection, in addition to user studies that motivated and evaluated the work.

Chapter III

INITIAL DEVELOPMENT OF FRONT PROJECTION FOR INTERACTIVE SURFACES

3.1 Warped Front Projection

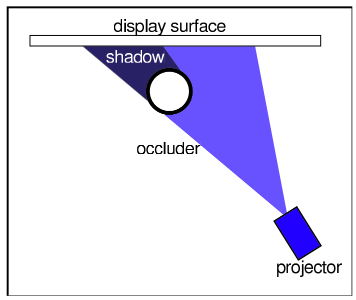

Figure 2: Warped Front Projection

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 where it is less likely to shine light on users. By moving the projector off-axis with respect to the display surface, it can project at a highly acute angle to minimize the area occupied by the projection frustum and hence the likelihood of occlusions (Figure 2). A standard data projector can be mounted at a moderately acute angle (30° to 35° off-axis), and commodity 3-D video cards can be used to pre-warp the projected image to compensate for keystone distortions. Because of the software image warping required to present a distortion free display, we call this technique *Warped Front Projection* (WFP).

The limiting factor for how far a standard projector can be mounted off-axis is it's depth-of-focus, or the range in distance from the projector within which the image remains in focus, which is typically one to two feet. As the angle becomes more acute, portions of the display surface will start to leave the field of focus, and the edges of the display will begin to appear blurry as they move out of optical focus.

A WFP display can be constructed using a standard projector and software tools to pre-warp the image such as the nVidia driver NVKeystone feature, or our WinPVRP software application

(Section 6.2.1). Advanced projectors also have limited built in horizontal keystone correction features, which may work optically (lens-shift) or via geometry processing video chips, but usually only allow for 10° to 15° or less of off-axis placement.

A few commercial projectors such as the 3M Idea Board [74] and the NEC WT-600 projector [77] are designed to be mounted within 1m (3ft) of the display surface and use specialized optics such as aspherical mirrors to warp the projected image. In addition to warping the projected image to compensate for keystone distortions, these optics also have appropriately varying focal lengths for the varying lengths of the beam path. Software based warping can not compete with custom designed optics from a performance or quality standpoint, but these low-volume niche application projectors are typically three to five times more expensive than a commodity video projector.[1] Even with a very acute projection angle provided by expensive optics, these warped front-projection systems suffer from some occlusions whenever the user comes close to or touches the display, making them less than ideal for interactive applications. The areas of occlusion can be filled-in by using a second projector to provide redundant illumination.

3.2 Passive Virtual Rear Projection

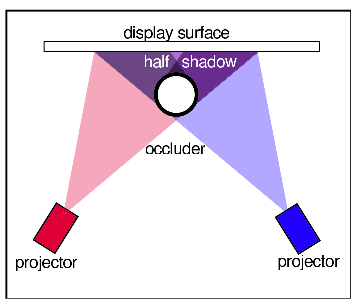

Figure 3: Passive Virtual Rear Projection

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 which use multiple redundant front projectors to approximate the experience of a rear projected surface. A *Passive Virtual Rear Projection (PVRP)* display (Figure 3) uses two (or more) projectors to provide

[1] In 2007, four years after it was introduced, the NEC WT-600 could be purchased from discount online retailers for as low as $2,500, and the updated NEC WT-610 (2000 *vs* 1500 lumens) was similarly priced. A comparable new 2000 lumen XGA projector without the aspheric mirror technology could be purchased for less than $700.

redundant illumination, without actively compensating for occluders.

Most areas that are shadowed in one 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 largest challenge to providing passive redundant illumination is for the system to accurately align the projected images on the display surface. Computer vision and homographies can be used to align the projected images to within sub-pixel accuracy.

3.3 Computer Vision and Homographies for Calibration

In a multi-projector system, several projectors are positioned so that their outputs converge onto a display surface (Figure 3). The goal is to combine light from the projectors to create a single, sharp image on the surface. Clearly, one cannot simply project the same raw image simultaneously through the different projectors; not only does a given point on the surface correspond to very different pixel locations in each projector, but the image produced on the surface from any single projector will suffer from keystone distortion as the individual projectors are mounted off-axis. By using a camera to find a relationship between the projectors, we can calculate how to pre-warp the source image for each projector so that the multiple projected images converge into a single image on the display surface.

We assume that the positions, orientations and optical parameters of the camera and projectors are unknown; the camera and projector optics can be modeled by perspective transforms; and that the projection screen is flat. Therefore, the various transforms between camera, screen and projectors can all be modeled as 2-D planar homographies:

$$\begin{pmatrix} xw \\ yw \\ w \end{pmatrix} = \begin{pmatrix} p_1 & p_2 & p_3 \\ p_4 & p_5 & p_6 \\ p_7 & p_8 & p_9 \end{pmatrix} \begin{pmatrix} X \\ Y \\ 1 \end{pmatrix} \tag{1}$$

where (x, y) and (X, Y) are corresponding points in the camera and projector frames of reference, and $\vec{p} = (p_1 \ldots\ p_9)^T$, constrained by $|\vec{p}| = 1$, are the parameters specifying the homography.

These parameters can be obtained from as few as four point correspondences, using well known camera-projector calibration techniques [59, 20]. One method to determine the homography for each camera-projector pair T_{c,P_i} is to project a rectangle from the projector into the environment. The coordinates of the rectangle's corners in projector coordinates (x_i, y_i) are known *a priori*, and the coordinates of the corners in the camera frame (X_i, Y_i) are located using standard image processing techniques.[2]

The user can interactively specify the display area by manipulating the outline of a projected quadrilateral until it appears as a rectangle of the desired size and position on the display surface. This directly specifies the homography between the selected projector and the screen $T_{p_i,s}^{-1}$; the outline of the selected rectangle can then be detected in the camera image as discussed above to determine the camera to screen homography $T_{c,s}$.

The projector-screen homographies $T_{P_i,s}$ model the geometric distortion (keystone warping) that is induced when an image is projected from an off-center projector P_i. This distortion can be corrected by projecting a *pre-warped* image, generated by applying the inverse transform $T_{P_i,s}^{-1}$to the original image.[3]

Since $T_{\{P_i,s\}}^{-1} T_{\{P_i,s\}} = I$, one can see that the pre-warping also aligns the images from different projectors so that all are precisely projected onto the screen S. Applying the homographies derived from camera images, a multi-projector array can thus be efficiently configured to eliminate keystoning distortions and redundantly illuminate the display surface. In practice, our system is able to achieve alignment within one pixel, meaning that each pixel touches the same pixel projected from other projectors.

This method is used by our WinPVRP application (Section 6.2.1) allowing users to easily calibrate two projectors into a PVRP display using a webcam. As demonstrated in Section 6.4.2, programmers using the PROCAMS toolkit are able to calibrate multiple projectors using this technique with a single function call after allocating projectors and cameras.

[2]Hough-transform line-fitting [4] locates the edges of the quadrilateral, and its corner coordinates are given by intersecting these lines.

[3]In our system, this pre-warp is efficiently implemented using the texture-mapping operations available in standard 3-D graphics hardware.

Chapter IV

PVRP EVALUATION

We decided to investigate just how much of a problem occlusions and shadows posed and how advanced the technology would have to become to be useful. Specifically, we questioned if it was necessary to dynamically compensate for shadows caused by the users. Simply providing redundant illumination (resulting in "half shadows"), without actively attempting to compensate for occlusions or suppress blinding light, might be sufficient for users to operate effectively.

Although it is our intuition that occlusions and shadows pose a problem to users of upright front projected displays (possibly explaining why many large scale interactive displays have been implemented using rear projection) we were unable to locate work that quantified the problem. We present here the first, empirical, end-user study of virtual rear projection. The study described here is designed to: 1) Determine the extent to which shadows on a front projected surface affect user task performance. 2) Investigate user strategies for coping with imperfect display technologies (which allow occlusions). 3) Evaluate two of the new projection technologies **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 [63].

4.1 Projection Technologies Studied

Figure 4 illustrates the projection technologies we studied:

- ***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 using 3D graphics hardware to provide a corrected display on the screen. Commercial and research prototypes demonstrate this on-board warping function, such as used by

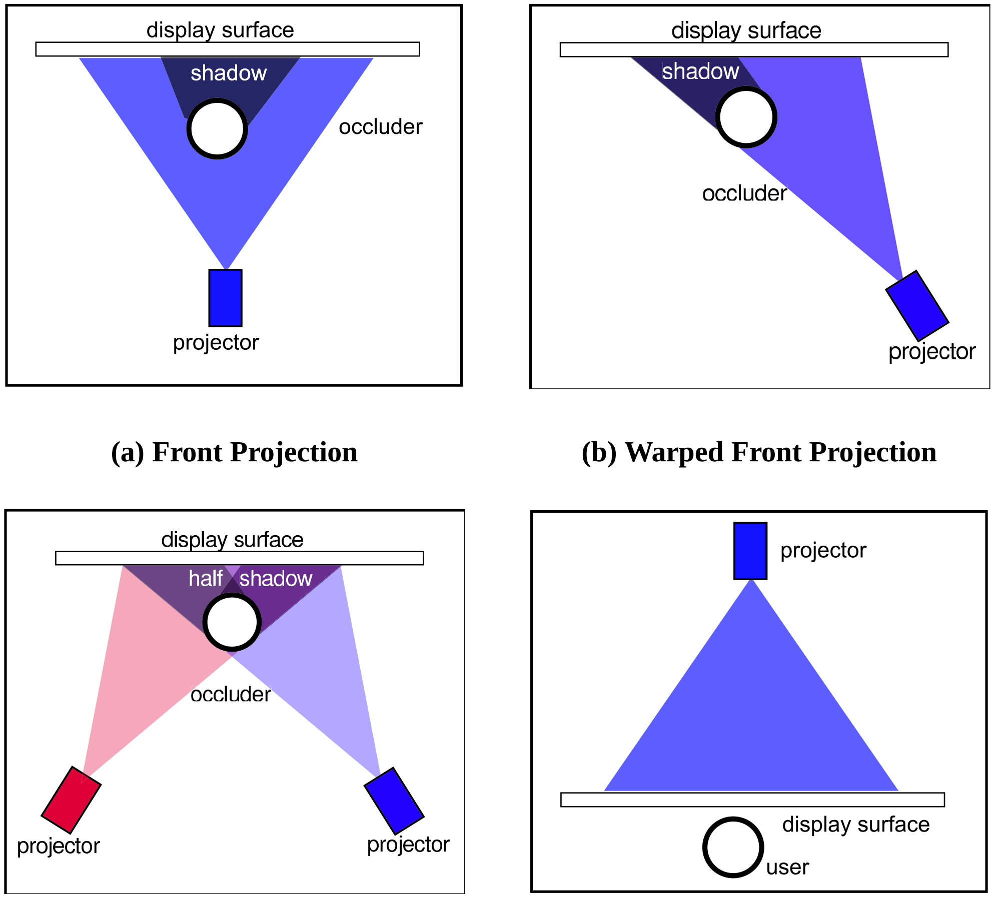

(a) Front Projection **(b) Warped Front Projection**

(c) Passive Virtual Rear Projection **(d) Rear Projection**

Figure 4: Taxonomy of Projection Technologies in our study.

the 3M IdeaBoard [25], NEC WT-600 [77], or the Everywhere Displays Projector [51]. Additionally, the latest version of the nVidia video card drivers includes a "keystoning" function which allows any computer running Microsoft Windows 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. After a calibration step using computer vision technology, output from each projector is independently warped (as with WFP) to correctly overlap on the display screen. This reduces the size and frequency of occlusions. Users standing very close to the screen may still completely occlude portions of the output but usually only occlude the output of one of the projectors, resulting in "half-shadows" where the output is still visible at a lower level of contrast.

- ***Rear Projection (RP)* -** By using a single projector mounted behind the screen, a rear projection solution prevents occlusions and shadows completely, but requires extra dedicated space for the beam path.

We performed this study when we had developed warped front projection and passive virtual rear projection technologies to a point where we felt they were ready to be evaluated by end users. We wanted to determine if this passive version of the technology would be sufficient to replace true rear projection, and if not, use the results to inform development of more active virtual rear projection technologies.

4.2 Study Setup

The study evaluated the effects of four different projection technologies on a *single user* working with a large scale interactive surface. Participants were asked to perform interactive tasks on a rear projection capable SmartBoard which utilized a contact sensitive film (touch screen) on the display surface for input. Our study presented participants with four counterbalanced conditions**:**

- **Front Projection (FP)**
- **Warped Front Projection (WFP)**
- **Passive Virtual Rear Projection (PVRP)**
- **Rear Projection (RP)**

4.2.1 Equipment Setup

Care was taken to adjust all conditions so that the intensity and resolution of the output was equal. Intensity was measured by a Sekonic Twinmate L-208 light meter to equalize light levels for all conditions and the output resolution was adjusted to provide an apparent resolution of 512x512, covering the entire SmartBoard screen, which measures 58” (1.47m) diagonally (See Figure 9). For the front projection conditions (FP, WFP, VRP) 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 participant's interaction with the display surface.

4.2.2 Study Participants

Our study participants were seventeen (17) college students, 9 males and 8 females, mean age of 21.3 (σ=1.77), from the experimental pool of the School of Psychology at our institute. To avoid handedness effects, we selected right-handed participants who exclusively used their right hand for interacting with the screen (without a pen or stylus). All participants had normal eyesight or wore corrective eye-wear to bring their eyesight to normal.

4.2.3 Study Tasks

A photographic image was used to evaluate subjective image quality, and three tasks were presented to the participants. These tasks exercise the basic searching, selecting, dragging and tracing options that a user performs with an interactive surface to perform such UI interactions as button pushing, slider movement, icon dragging, sketching etc. Although they did not directly simulate the use of real applications, we felt that the tasks are relevant for many standard UI interactions and hence, many applications.

Crosses Task (Accurate Selection) - 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 using the SmartBoard touch sensitive surface.

Box Task (Fast Search, Selection, and Dragging) - Boxes with 2" sides appeared pseudo-randomly in one of 8 positions around the perimeter of the screen (Figure 5), with a 4" target placed in the center. The user was instructed to drag each box into the target. Each user moved eighty (80) boxes

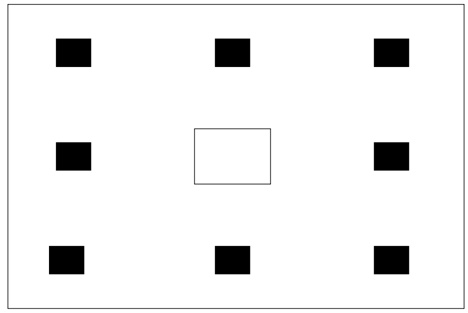

Figure 5: Center target and the eight possible box starting positions.

(ten boxes from each of the eight positions) for each projection technology.

For each box, the search/select (acquire) time, drag time, and total time were recorded, as were the number of drags and touches needed to move the box 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 visible.

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

4.3 Results

Figures 6 and 7 summarize our significant results and present the pairwise T-tests resulting from our statistical analysis. In our within-subjects design, participants experienced each condition in a counter-balanced order. Subjective measures were collected via questionnaire after each condition, while quantitative measures were recorded by the computer administering the tasks. We analyzed the data using a repeated measures ANOVA. To correct for a potential violation of the sphericity assumption we applied a Greenhouse-Geisser correction in all cases. The independent variable was treatment condition (FP,WFP,PVRP,RP).

4.3.1 Subjective Results

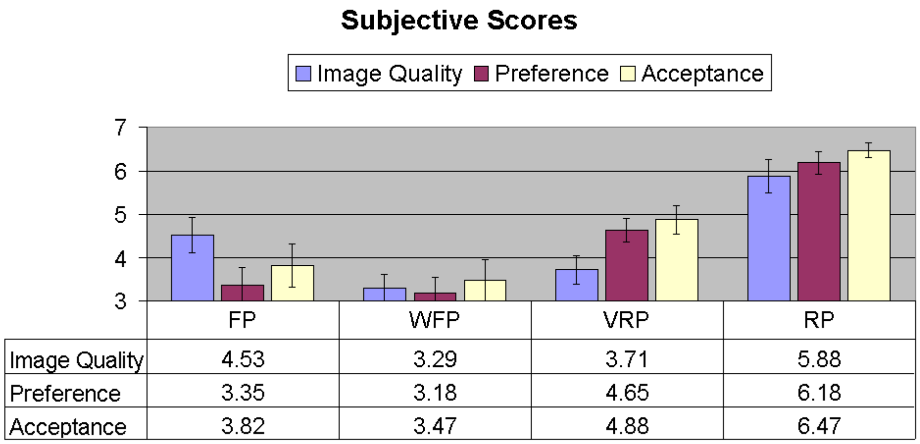

Condition	Compared With:	Image Quality	Preference	Acceptance
FP	WFP	**1.235**	0.176	0.353
	PVRP	0.824	**-1.294**	**-1.059**
	RP	**-1.353**	**-2.824**	**-2.647**
WFP	PVRP	-0.412	**-1.471**	**-1.412**
	RP	**-2.588**	**-3.000**	**-3.000**
PVRP	RP	**-2.176**	**-1.529**	**-1.588**

Figure 6: (Top) Subjective scores from participant questionnaires. (Bottom) Pairwise comparisons of Image Quality, Preference, and Acceptance scores based upon treatment condition. Positive numbers indicate the condition scored higher than the “compared with” condition. Statistically significant differences ($p<0.05$) are presented **in bold**.

A main effect was found for all subjective measures. *[Image Quality: $F(2.224, 35.589) = 9.755$, $p < 0.001$; Preference: $F(2.359, 37.745) = 20.812$, $p < 0.001$; Acceptance: $F(2.156, 34.5) = 17.366$, $p < 0.001$].*

Image Quality - Because we were projecting onto a display surface optimized for rear-projection, the rear projection condition was strongly biased and had the highest reported image quality.[1] In the post session interview of the primary study we found that the factor leading to the image quality score was primarily the sharpness (or blurriness) of the image (100% of the participants) with some

[1] *“How would you rate the image quality of the display technology? [Poor Quality = 1 2 3 4 5 6 7 = Excellent Quality]”*

of the participants citing intensity or color saturation (29%) and shadows (6%) as additional factors.

We attribute the poor showing of PVRP and WFP (leftmost bars in the graph of Figure 6) to using the SmartBoard's display (designed for on-axis projection) for all conditions, which was needed to control for extraneous variables. To control extraneous 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. The use of the rear projection display surface in all conditions resulted in biased subjective image quality scores, and these numbers should not be trusted as they will not generalize to other types of display surfaces.

We performed a small followup study with ten participants running an image quality survey on a front projection screen with the front projected conditions (FP, WFP, and PVRP) (See Section 4.4). One goal of this study was to determine the effects of our primary studies' projection surface which was optimized for rear projection, on the image quality scores for the front projection cases. Participants in this secondary study did not perform the performance measurement tasks (Crosses, Box, Spiral). The same photographic image, intensity, resolution, and questionnaire were used to measure subjective image quality. Although the image quality scores in Table 1 cannot be directly compared to the primary study, the trends in image quality scores indicate that warped front projection can produce an image quality that rivals that of a front projector, while suggesting that the slight differences in image alignment for virtual rear projection produce a slightly lower quality image, even on a front projection surface.

Preference - Rear projection was preferred over the other projection technologies on the preference question[2] with passive virtual rear projection being preferred over the single projector conditions (FP & WFP). When asked to volunteer what factors they considered when making their preference

[2] *"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]"*

judgments, about half of the participants mentioned image quality (65%) and an equal number mentioned shadows (65%) or lack thereof. Users ranked the image quality of PVRP lower than that of FP and WFP, yet their preference rankings for PVRP were significantly higher than that of FP & WFP. This, combined with the large number of participants who volunteered that shadows were a factor in their preference rankings indicates that PVRP was preferred because of its ability to eliminate virtually all occlusions.

Acceptance - The user acceptance question[3] 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%), and shadows (53%). Ease of performing the task (12%), touch-screen problems (12%), unspecified reasons (6%) and "just kind'a a gut reaction" (6%) made up the remainder of responses.

4.3.2 Quantitative Measures: Speed & Accuracy

Box Task (Fast Search, Selection, and Dragging) - 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 acquisition time between occluded and unoccluded boxes and recorded the behaviors participants adopted to compensate for shadows (see Section 4.3.3). Figure 8a 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 because our projectors were mounted near the ceiling and the users were right-handed. Additionally, WFP and VRP reduced the time it took users to acquire an occluded box. This was due to the fact that less of the user's shadow would cover the screen, allowing them to uncover and detect the box with less motion.

[3] *"Please rate your willingness to use this display technology on the following scale: [Absolutely unacceptable = 1 2 3 4 5 6 7 = Completely acceptable]"*

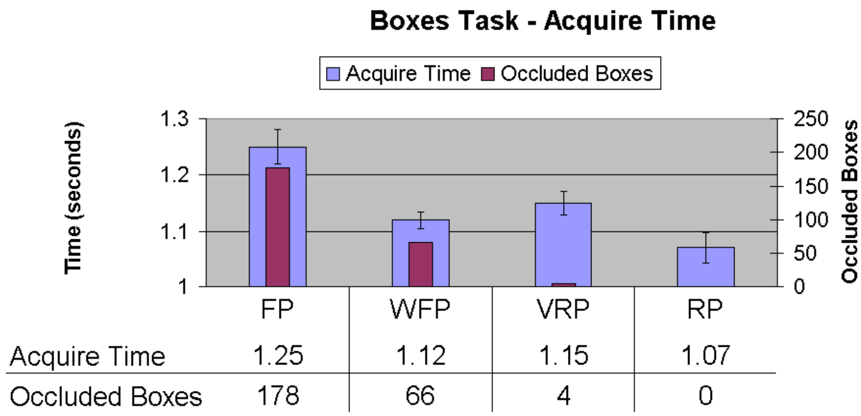

Condition	Compared With:	Mean Diff. (ms)	Std. Error (ms)	Sig.
FP	**WFP**	**128**	**25.1**	**0.000**
	PVRP	**102**	**24.9**	**0.001**
	RP	**185**	**29.2**	**0.000**
WFP	PVRP	-25	13.0	0.072
	RP	**57**	**20.8**	**0.014**
PVRP	**RP**	**82**	**17.4**	**0.000**

Figure 7: (Top) Acquire times in the Box task with number of occluded boxes in each condition. (Bottom) Pairwise comparisons of Box Acquire Time (in milliseconds) based upon treatment condition. Positive numbers indicate how much slower the "condition" is than the "compared with" condition. All statistically significant differences ($p<0.05$) are presented **in bold**.

In the Box task the dependent variables, measured in milliseconds, were (box) Acquire Time and Total Time. A main effect was found based upon the treatment condition for Acquire Time *[Acquire Time: $F(2.127,34.036) = 23.940$, $p <= 0.001$]*; no significant difference was found between conditions for the total task completion time, although the data trends were similar to that shown by the acquire time dependent variable. The lack of statistical significance with N=17 is attributable to a larger variance in the task completion time data.

Crosses & Spiral (Accurate Selection & Fast Tracing) - These tasks differed from the Box Task in that the whole task was visually presented at once (a full spiral or all crosses) allowing the participants to plan their motion. In the Crosses Task, participants would generally work from one side of the screen to the other, keeping their shadow away from crosses they were working on. We

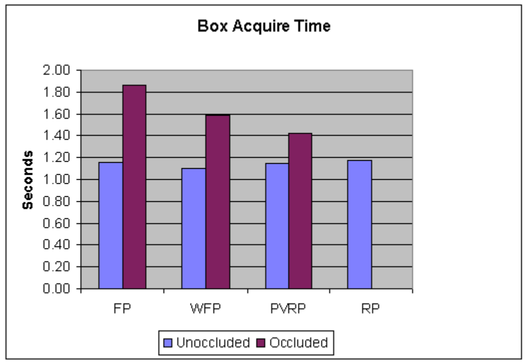

Figure 8: Acquire time for occluded and unoccluded boxes.

found no significant difference between the four conditions for accurate selection.

The Spiral Task measured the user's ability to trace a curve quickly, exercising muscle motions similar to free form drawing or writing in a more controlled setting. Users would sway to avoid casting a shadow on the portion of the spiral they were currently tracing. Conditions which eliminated or reduced shadows (RP & PVRP) had faster mean completion times than conditions which did not (FP & WFP), but these trends are not statistically significant.

4.3.3 Coping Strategies

Behavior in the PVRP and RP cases (minimal to no occlusions) were identical for all of the tasks, 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.

In the FP and WFP conditions, the participants adopted coping strategies to work around their shadows. 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 participant 3, see the "Dead Reckoning" strategy below) 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.

Strategies developed for the Box Task, which included randomly appearing targets, were much more involved. Participants generally used one of the following four strategies. Almost all participants settled into a single strategy fairly quickly (within 10 boxes). Participant 9 changed from the

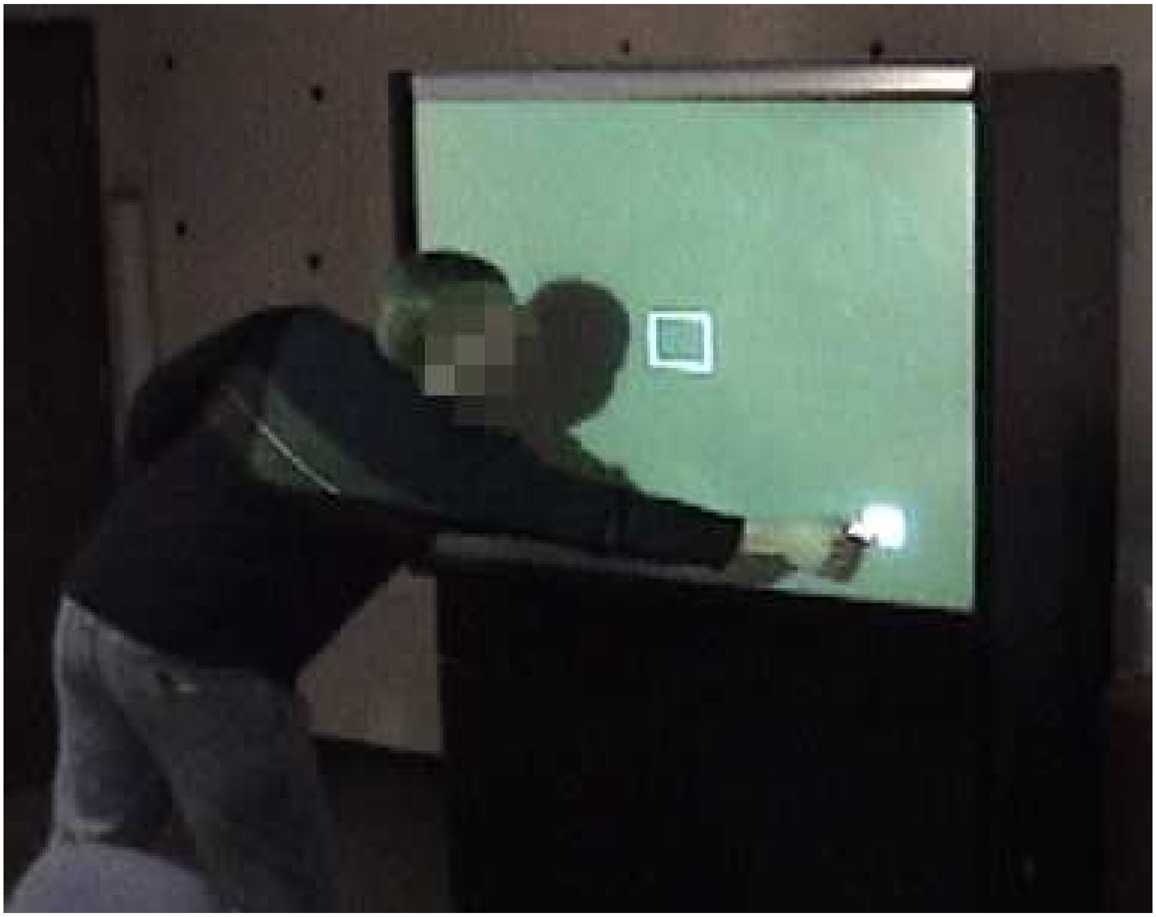

Figure 9: Participant exhibiting the edge-of-screen coping strategy while working the Box Task in the Front Projection condition.

Edge of Screen to the Move on Occlusion strategy half way through the run, and is counted in both.

- *Edge of Screen* (7 of 17 participants) - Participants stood at the edge of the screen. Four participants would lean inward to move boxes, immediately returning to their home position to insure that they were not occluding the next box. (See Figure 9.) Three participants 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 in line with the target). Three participants were short enough to occlude few boxes, while four participants would occlude boxes and would "sway" their entire upper body twenty to forty degrees to find occluded boxes.

- *Move on Occlusion* (3 of 17 participants) - Participants 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) - This participant 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, this participant would "drag through" his shadow along the curve, also with impressive accuracy.)

4.3.4 Participant Awareness of Shadow Coping Strategies

About half of the participants (47%) volunteered that they developed strategies to cope with occlusions, (*"Were there any specific strategies you used to perform the tasks?"*) while others (47%) 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 relatively diminutive participant (6%) who had only occluded 3 boxes (the average participant occluded 14.6 boxes) declared that she had no problems with the shadows.

Interestingly, of the eight participants who volunteered that they had developed strategies to deal with the shadows, seven (41%) stated that shadows were a factor in their preference ratings, while one (6%) only reported having considered image quality. Of the eight who only recognized their shadow coping behavior after being prompted by the interviewer, three (18%) cited shadows as a factor in their preference ratings, while five (29%) reported using image quality exclusively.

4.4 Followup Blinding Light Comfort Level Study

While investigating image quality on a front projection surface (followup study described in our Image Quality section) we also evaluated the necessity of a VRP system to provide blinding light suppression. To investigate this issue we added the task of reading two cards displayed at the back of

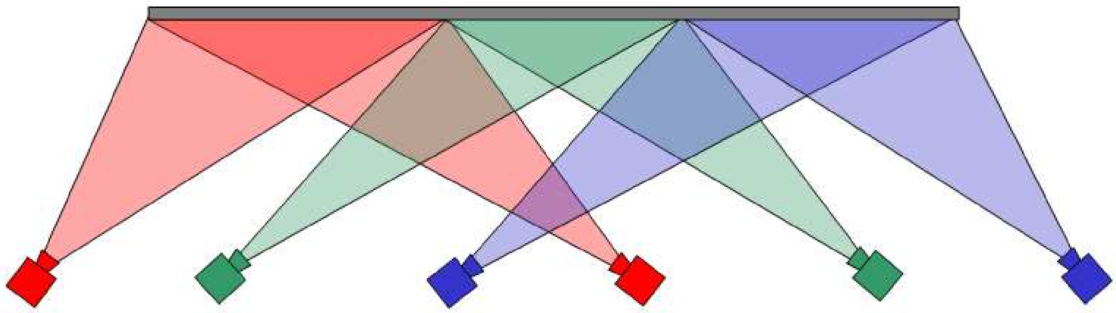

Figure 10: Projector locations and beam-paths for a 17.5ft (5.3m) wide electronic whiteboard using passive virtual rear projection. Users find it extremely difficult to avoid standing within projection beams.

the room which forced the participants to face the projectors as if giving a presentation. Participants were then asked to rate the "Annoyance" level of each condition.[4]

Table 1: Mean (Standard Deviation) subjective measures on a 7 point scale, on image quality and annoyance of projected light on a front projection screen. **Bold** data indicates statistical significance.

Condition	**Image Quality**	**Annoyance**
Front Projection (FP)	4 (1.15)	6.5 (0.53)
Warped Front Proj. (WFP)	4.1 (0.99)	5.9 (1.37)
Passive Virtual Rear Proj. (PVRP)	3.2 (1.62)	**4.5 (2.07)**

As with the primary study discussed previously, the user was placed in a specific location when performing the image quality task (three feet from the screen, two feet to the left of center). This placement was chosen so that they were *not* blocking the beam path for the front projection (FP) and warped front projection (WFP) conditions, and *were* blocking the beam path of the left projector for the passive virtual rear projection (PVRP) condition. This location was chosen based upon our observations of projector users, who almost exclusively choose to stand outside of the beam path when possible. We deliberately placed participants in the beam path for the VRP condition, as it is much harder to avoid a pair of projectors, and the actual deployment of virtual rear projection technologies will likely make it even more difficult to avoid beam paths. Figure 10 shows that as you add projectors for a wall sized PVRP system, the locations where a user is "safe" from being projected upon is drastically reduced, especially as they approach the display surface for interaction.

[4]*"Did you find the light from the projector(s) to be annoying? [Annoying = 1 2 3 4 5 6 7 = Unnoticeable]"*

The result of this decision was that neither the FP or WFP conditions beamed light directly into the participant's faces. The comfort scores in Table 1 for FP and WFP are understandably higher than for VRP, and even with such a limited participant pool the difference between PVRP and the other conditions was significant ($p <= 0.05$).

Essentially, the blinding light aspect of this followup study only had two conditions (user in beam, user out of beam), although because we were running it in conjunction with the image quality questionnaire we had to run all three (FP, WFP, PVRP) conditions. It is unsurprising that the differences in the comfort scores of FP and WFP are not significantly different. However, we have shown that the effect of being in the path of a projection beam (the case with the PVRP condition) is large enough to make a detectable difference with even a small sample size (N=10), leading evidence that the projected light is noticeable and annoying.

4.5 Discussion

In our studies, we found that humans are able to adapt to occlusions and shadows from front projection systems via coping behaviors to maintain their level of task performance. We observed four different types of coping behavior which users developed early and quickly in the front projection (FP) sessions. This indicates that at least for simple tasks, and only considering efficiency, a single front projector is sufficient.

However, there are two important qualifications. First, our tasks were quite basic, and we did not measure the amount of cognitive load executing the coping strategies placed on the users. More cognitively challenging tasks may suffer from the use of front projection coping strategies. Secondly, and more importantly, even though performance was comparable, our participants strongly disliked front projection when comparing it to rear projection (a significant subjective preference rating difference between 3.35 and 6.18). There are very few applications where the user's preference does not play a strong role in acceptance and adoption, and these preference scores cannot be discounted.

Assuming that a system already has an accelerated 3D graphics card, a warped front projection (WFP) system adds nothing to the hardware cost of a traditional front projection (FP) system, although system software must be designed to use the graphics card to correctly warp the output.

Our primary study indicates that such a system reduces occlusions by an average of 62% when compared to a straight front projection system. We believe the low preference score for WFP in our primary study was due to the unfair disadvantage presented by the off-axis projection onto the rear-projection surface. Our followup study on a front-projection surface showed that WFP image quality was virtually identical to a standard front projection system when used on a front projection surface. We recommend warped front projection in situations where only a single projector is available and the application software allows the easy addition of warping code.

Passive virtual rear projection (PVRP) had the highest user preference scores out of the front projection technologies, eliminated user's coping behavior and virtually eliminated occlusions. For these reasons, we recommend PVRP when the user desires a rear projection (RP) solution, but is constrained by the available space. If the space and resources are available, a rear projection system continues to provide the best user experience.

However, the twin facts that 1) users preferred rear projection to our passive virtual rear projection (PVRP), and 2) that they found blinding light annoying, motivate further development of VRP technologies. Although seemingly obvious, we have empirically confirmed that users notice when they are in the beam path of a projector and find it moderately annoying, motivating the addition of shadow elimination and blinding light suppression to active virtual rear projection technologies. For this reason, we must expand virtual rear projection our taxonomy of projection technologies discussed previously as follows:

- ***Active Virtual Rear Projection (AVRP)*** **-** Similar to PVRP, AVRP adds a camera or other sensor which determines when one of the projectors is occluded. The system then attempts to compensate for this occlusion by boosting output power from the other projector(s) to increase contrast in the "half-shadow" area(s), effectively eliminating them [30, 61].

- ***AVRP with Blinding Light Suppression (AVRP-BLS)*** **-** Similar to AVRP, AVRP-BLS adds the ability to detect and turn off projector output that is shining on an object other than the screen, such as an intervening user. This blinding light suppression allows users to comfortably face the projectors without blinding light or distracting graphics being projected into their eyes or onto their bodies [61].

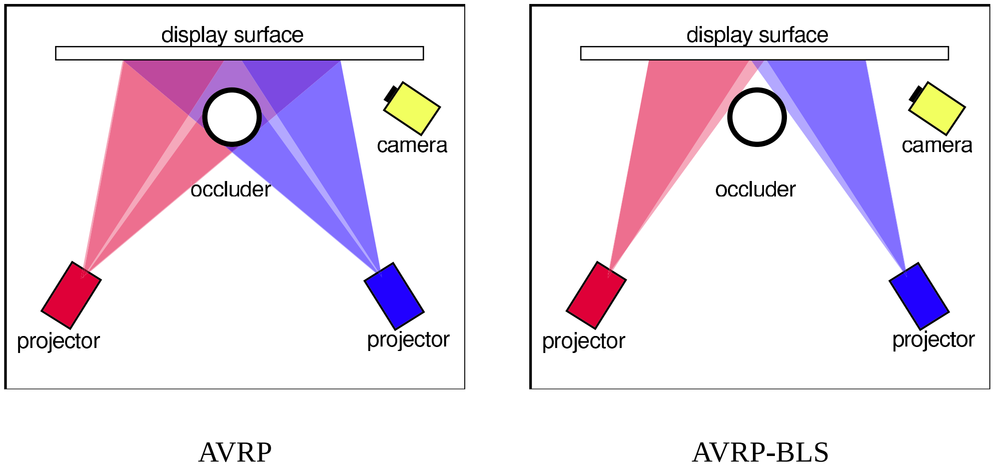

Figure 11: Additions to projection technologies taxonomy.

Technically, active virtual rear projection (AVRP) is more complicated than passive virtual rear projection (PVRP). To implement a PVRP system, the two projectors must be calibrated once upon installation (and whenever they are moved), a step which can be done in under a minute with computer vision techniques. AVRP and AVRP-BLS requires continuous processing after the calibration step, to automatically locate occluders and modify the projector's output to compensate for occlusions and shadows, remove blinding light in the case of AVRP-BLS, and blend the output of the projectors to present a seamless display. Ideally, all of this must be accomplished fast enough to be imperceptible to users.

Chapter V

ACTIVE VIRTUAL REAR PROJECTION

This chapter describes three algorithms (Shadow Elimination, Shadow Elimination + Blinding Light Suppression, and Switching) that actively compensate for shadows and occlusions. The first two algorithms were originally developed by a group of researchers at Compaq Research Labs, including Rahul Sukthankar, Tat-Jen Cham, Gita Sukthankar and my adviser James Rehg[62, 60]. In the course of my thesis, I re-implemented these algorithms (equation 5 reported in Section 5.2.4 is corrected from the original paper) developed the switching algorithm (Section 5.3) in conjunction with Masters student Ramswaroop Somani, and performed the comparative evaluation reported in Section 5.4. During the course of implementation, development, and evaluation, the switching form of AVRP was the clear winner and in subsequent chapters the switching form of AVRP was used for deployment in the PROCAMS toolkit and user evaluations.

5.1 Shadow Elimination

By adding a camera or other sensor (Figure 12) that is able to detect the shadows on the display surface it is possible to dynamically correct penumbral shadows by projecting additional light into the region from one of the non-occluded projectors. This shadow elimination system must precisely

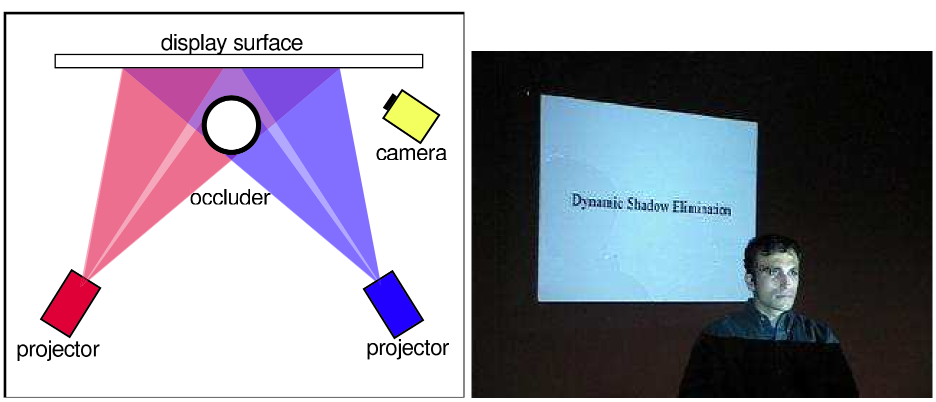

Figure 12: *Left:*Shadow Elimination. *Right:* Penumbral shadows are eliminated but the blinding light remains.

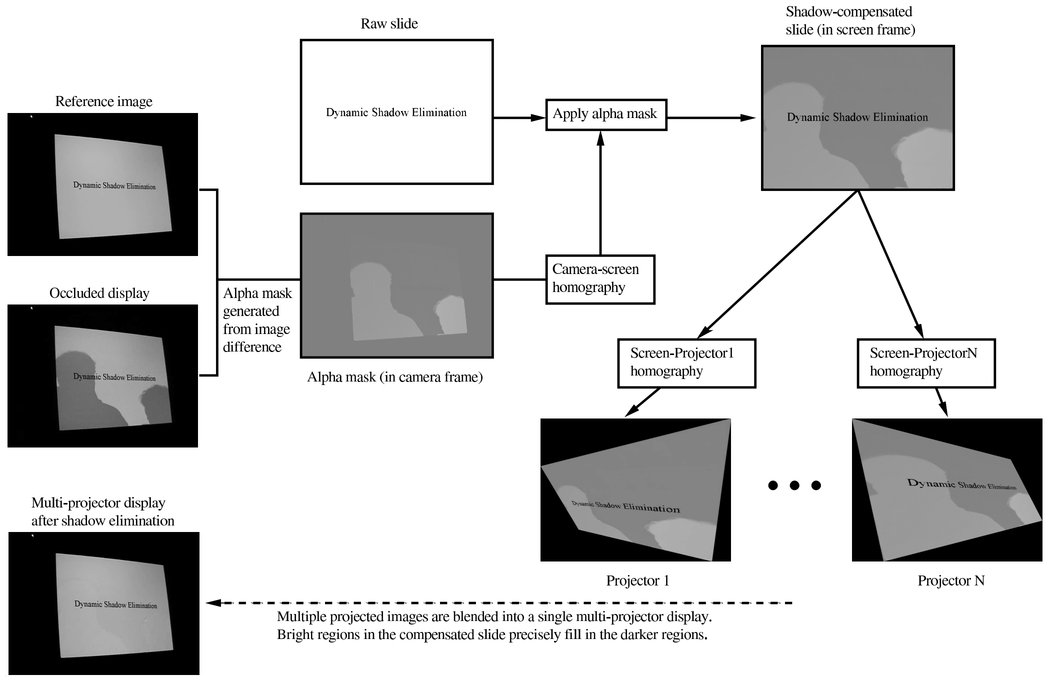

Figure 13: This diagram summarizes the occlusion detection and shadow elimination algorithms. The images in the left column were taken by the system camera during operation. The two penumbral occlusions caused by the person blocking both projectors are identified and corrected to create a shadow-free display (bottom left). See text for details.

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, over-illumination artifacts will be created. The shadow boundaries must be treated carefully since humans are very sensitive to edge artifacts.

5.1.1 Occlusion detection

The shadow elimination system focuses exclusively on detecting artifacts on the display surface. These can occur for either of two reasons. First, uncorrected penumbral occlusions appear as darker regions in a camera image that can be corrected by projecting additional light into the region. Second, artifacts may be caused by over-illumination of the display area, and occur most often when an occluding object (whose shadows had been eliminated) moves away suddenly. These bright spots are corrected by reducing the light intensity in the region. Our shadow elimination algorithm makes no assumptions about the locations, sizes or shapes of occluders.

Figure 13 illustrates the algorithm. During its initialization phase (when the scene is occluder-free) the system projects each image it wishes to display and captures several camera images of the projected display. These images are pixel-wise averaged to create a reference image for that slide, and this image represents the desired state of the display (Figure 13, top left). The goal of occlusion detection is to identify regions in the current image that deviate from this ideal state. During operation, the system camera acquires a current image of the projected display which may contain uncorrected shadows. For example, the image shown in Figure 13 (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).

Since the display surface remains static, a pixel-wise image difference between current and reference camera images can be used to locate shadows and over-compensation artifacts. To reduce the effects of camera noise and minor calibration errors, we apply a $5x5$ spatial median filter to the difference image. 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 (α_t), which when applied to the current camera image, should bring it closer to the reference image. Alpha values range from 0 (dark) to 255 (bright), and the mask is initialized to 128 at $t = 0$. The alpha mask is updated at every time-step using the following simple feedback system:

$$\alpha_t(x, y) = \alpha_{t-1}(x, y) - \gamma \left(I_t(x, y) - I_0(x, y) \right),$$

where $I_t$is the camera image at time t, I_0 is the reference image, and is a system parameter (set to 0.3 in our implementation). For a static scene, the alpha mask converges to a stable fixed point in a very short period of time. A noteworthy point about our shadow elimination system is that all of the projectors in the multi-projector system use the *same* alpha mask for shadow removal. This reduces the amount of processing required, but results in additional light being projected onto occluders as described below.

5.1.2 Eliminating Shadows

The alpha mask (described above) integrates the previous state of the shadow correction, and information from the current difference image. However, since it was computed in the camera frame of reference, it must be transformed into the screen frame of reference before it can be applied; this is done using the camera-screen homography $T_{c,s}$, discussed in Section 3.3.

It is surprising that using the *same* alpha mask for all projectors correctly eliminates *all* of the penumbral shadows! This can be explained by the following argument. Consider the two-penumbra shadow configuration generated by the two-projector, one-occluder system shown in Figures 12 (right) and 13. From P1's perspective, the left high-alpha region falls precisely on the left penumbra (Shadow2) while the right high-alpha region simply over-illuminates the occluder. From P2's perspective, the left high-alpha region falls on the occluder (without effect) and the right one corrects for the right penumbra (Shadow1). Thus, both projectors are able to use the same image to eliminate shadows.

Since this algorithm does not use photometric models of the environment, projectors or camera, it cannot predict precisely how much light is needed to remove a shadow. However, the iterative feedback loop used to update the alpha mask allows us to avoid this problem: the system will continue adding light to shadowed regions until the region appears as it did in the reference image. This approach has additional benefits. For instance, the system is able to correct for the fuzzy occlusions caused by area light sources (e.g., the diffuse shadow created by a hand moving near the projector) without requiring an explicit model of the shadow formation process. One drawback to such an iterative technique is that the alpha mask can require several iterations to converge; in practice, shadows are eliminated in approximately 3 iterations. The second drawback of this form of active virtual rear projection with shadow elimination is that it indiscriminately projects additional light onto the occluder (user) as well as the areas of shadow on the display surface. If the user turns to face the projectors this blinding light is distracting [63].

5.2 Shadow Elimination + Blinding Light Suppression

To combat this blinding light being cast upon users, we must be able to determine which pixels in each projector are falling upon occluders. After the projectors have been geometrically aligned, we can easily determine which source pixels from the projectors contribute to the intensity of an arbitrary screen pixel. In the following analysis, we assume that the contributions are at some level additive. Given N projectors, the observed intensity Z_t of a particular screen pixel at time t may be expressed by:

$$Z_t = C\Big(k_{1t}S_1(I_{1t}) + \cdots + k_{Nt}S_N(I_{Nt}) + A\Big), \tag{2}$$

where I_{jt} is the corresponding source pixel intensity set in projector j at time t, $S_{j(\cdot)}$ is the projector to screen intensity transfer function, A is the ambient light contribution, assumed to be time invariant, $C(\cdot)$ is the screen to camera intensity transfer function and k_{jt} is the *visibility ratio* of the source pixel in projector j at time t. Note that all the variables and functions also depend on the

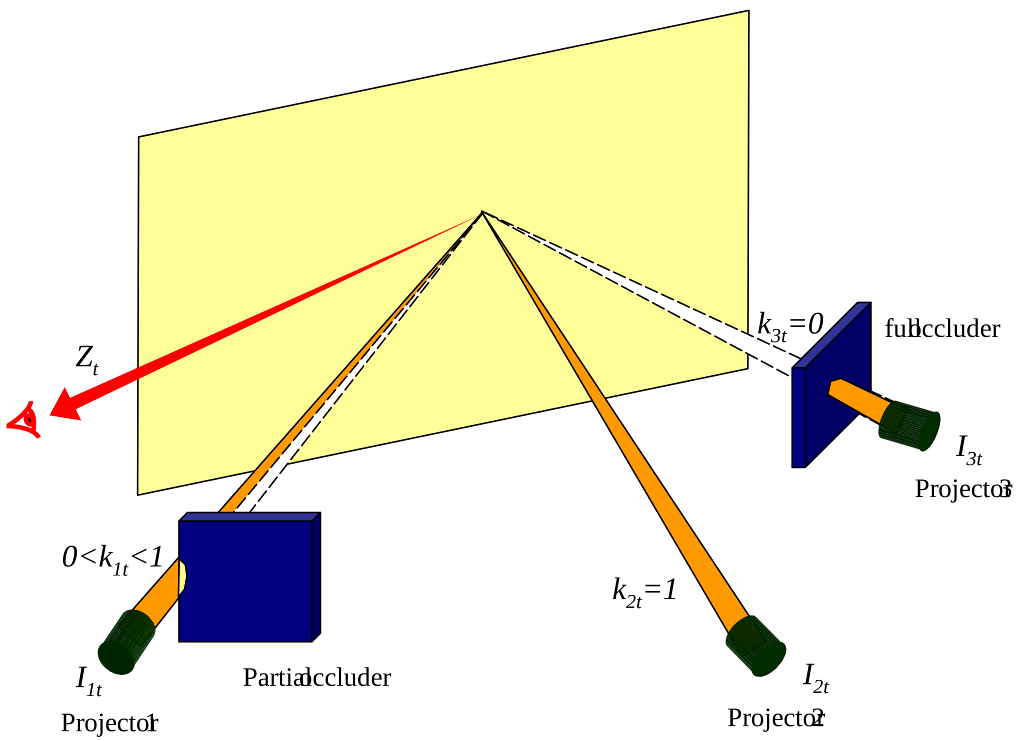

Figure 14: Photometric framework. This diagram illustrates equation (2), in which the observed display intensity Z_t is related to the combination of projector source pixels I_{jt} and the corresponding visibility ratios k_{jt}. The visibility ratios vary accordingly with non-occlusion, partial and full occlusion.

When occluders obstruct the paths of the light rays from some of the projectors to the screen, Z_t diminishes and shadows occur. This situation is quantitatively modeled via the visibility ratios, which represent the proportion of light rays from corresponding source pixels in the projectors that remain unobstructed. If the projectors were modeled as point-light sources, occluders would block either none or all of the light falling on a given pixel from any particular projector; therefore, k_{jt}

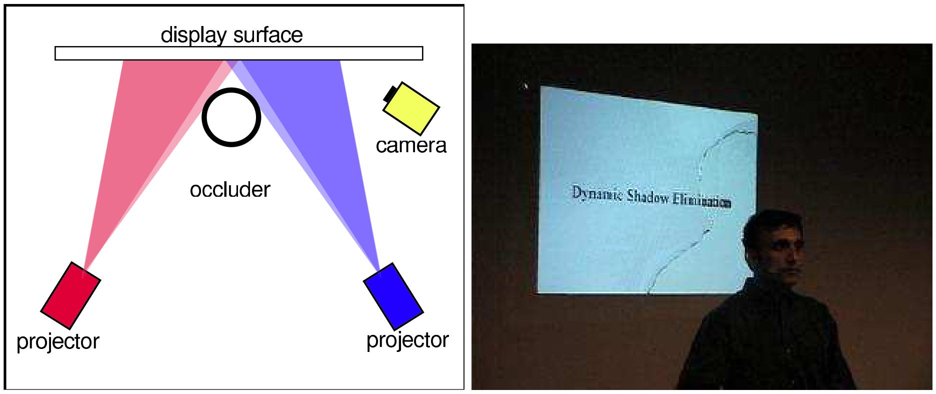

Figure 15: *Left:* Shadow Elimination with Blinding Light Suppression. *Right:* Light is kept off of the occluders face.

would be a binary variable. However, this assumption is not valid in real-world conditions. Our system must cope with partial occluders (created by objects near the projector) that cast fuzzy-edged shadows on the screen. In these cases k_{jt} denotes the degree of occlusion of projector j for the given pixel.

5.2.1 Occlusion Detection

The Blinding Light Suppression system focuses exclusively on detecting deviation of the observed intensities on the screen from the desired intensities when occluders are not present. The major cause of deviation is occlusion, although deviation can also occur because of changes in ambient lighting, projector failure, etc. Our system can handle all of these problems (as discussed in the next section). No assumptions are made about the locations, sizes or shapes of occluders.

Mathematically, the desired intensity of a particular screen pixel may be represented by Z_0. This may be obtained in the initialization phase when the system projects each presentation slide and captures several camera images of the projected display while occluders are absent. As an occluder is introduced in front of projector k to create penumbral shadows, the visibility ratio k_{jt} decreases, such that $k_{jt} < 1$. Hence $Z_t < Z_0$. These deviations in the screen can be detected via a pixel-wise image difference between current and reference camera images to locate shadow artifacts.

5.2.2 Iterative Photometric Compensation

Our system handles occluders by

1. compensating for shadows on the screen by boosting the intensities of unoccluded source pixels; and

2. removing projector light falling on the occluder by blanking the intensities of occluded source pixels.

The degrees-of-freedom available to us are the source pixel intensities I_{jt}, which may be changed. Hence for a shadowed screen pixel where $Z_t < Z_0$, we ideally want to compensate for the shadow (i.e. setting $Z_{t+1} = Z_0$) by (i) increasing $I_{j(t+1)}$ to be larger than I_{jt} if $k_{jt} = 1$, and (ii) reducing $i_{j(t+1)}$ to zero if $k_{jt} < 1$.

However, it is very difficult to accurately model $C(\cdot)$ and $S_j(\cdot)$. Even if we know the exact values for the ambient lighting and visibility ratios, it is almost impossible to update the source pixels such that in one time step the shadows are eliminated. Fortunately, we expect $C(\cdot)$ and $S_j(\cdot)$ to be positive monotonic, and an iterative negative feedback loop can be used to compute $I_{1t} \ldots, I_{Nt}$ required to minimize $Z_t - Z_0$.

The advantages of such a system are:

- it does not require explicit modeling of $C(\cdot)$ and $S_j(\cdot)$,
- it does not require explicit measurement of the visibility ratios k_{jt},
- it is able to handle slowly varying ambient light.

As in Section 5.1, the change in the intensity of each source pixel in each projector is controlled by the alpha value associated with the pixel:

$$I_{jt} = \alpha_{jt} I_0, \tag{3}$$

where I_0 is the original value of the source pixel (i.e. pixel value in the presentation slide) and is the same across all projectors, while α_{jt}, which can vary between 0 and 1, is the time-varying, projector-dependent alpha value. The alpha values for the source pixels in one projector are collectively termed the alpha mask for the projector.

The earlier shadow elimination system described in Section 5.1 can compensate for shadows but is incapable of suppressing projected light falling on the occluder. In particular, that simpler method cannot distinguish between the contributions of individual projectors. Instead, all projectors boost their pixel intensities for each occluded region. This has two undesirable consequences: (1) bright "halos" may appear around eliminated shadows, particularly when occluders are in motion; and (2) the amount of distracting light projected on users is *increased* rather than reduced by the system. This motivates the need for a more complex solution where the alpha masks are different for different projectors.

The approach adopted here is to design components which separately handle the problems of shadow elimination and occluder light suppression, and integrate them into a complete system. These are discussed in the following sections.

5.2.3 Shadow Elimination

Eliminating shadows involves increasing values for corresponding source pixels. The shadow elimination (SE) component of the system is based on

$$(\Delta\alpha_{jt})_{\text{SE}} = -\gamma(Z_t - Z_0), \tag{4}$$

where $\Delta\alpha_{jt} = \alpha_{j(t+1)} - \alpha_{jt}$ is change of α_{jt} in the next time-frame, and γ is a proportional constant (γ is 0.7 in our implementation). This component is a simple, linear feedback system.

5.2.4 Blinding Light Suppression

Suppressing projector light falling on occluders involves diminishing the source pixels corresponding to the occluded light rays. We determine whether a source pixel is occluded by determining if any changes in the source pixel result in changes in the screen pixel. However, since there are N possible changes of source pixel intensities from N projectors but only one observable screen intensity, we need to probe by varying the source pixels in different projectors separately. This cyclical probing results in a serial variation of the projector intensities.

The light suppression (LS) component of the system is based on

$$(\Delta\alpha_{jt})_{\text{LS}} = -\beta \frac{\Delta\alpha^2_{j(t-N)}}{\Delta Z_t^2 + \epsilon}, \tag{5}$$

where $\Delta Z_t = Z_t - Z_{t-N}$ is the change in the screen pixel intensity caused by the change of alpha value $\Delta\alpha_{j(t-N)}$ in the previous time frame when projector j is active, β is a small proportional constant and ϵ is a small positive constant to prevent a null denominator (β and ϵ are 0.1 in our implementation).

The rationale for (5) is that if the change in α_{jt} results in a corresponding-sized change in Z_t, the subsequent change in α_{jt} will be relatively minor (based on a small β). However if a change in α_{jt} does not result in a change in Z_t, this implies that the source pixel is occluded. The denominator of (5) approaches zero and α_{jt} is strongly reduced in the next time frame. Hence occluded source pixels are forced to black.

Note that the probe technique must be employed during shadow elimination as well. In particular, the system must be able to discover when a pixel which was turned off due to the presence of an occluder is available again, due to the occluders disappearance. This constraint is smoothly incorporated into our algorithm.

5.2.5 Integrated System for Shadow Elimination and Blinding Light Suppression

The integrated iterative feedback system combines (4) and (5) to get

$$\Delta\alpha_{jt} = (\Delta\alpha_{jt})_{\text{SE}} + (\Delta\alpha_{jt})_{\text{LS}}\,. \tag{6}$$

The alpha values are updated within limits such that

$$\alpha_{jt} = \begin{cases} 1, & \text{if } \alpha_{jt} + \Delta\alpha_{jt} > 1, \\ 0, & \text{if } \alpha_{jt} + \Delta\alpha_{jt} < 0, \\ \alpha_{jt} + \Delta\alpha_{jt}, & \text{otherwise}. \end{cases} \tag{7}$$

The following synthetic example (See Figure 16) illustrates the system. Suppose that each projector has an initial alpha value of 0.5 (both projectors illuminating equally at half brightness, $\alpha_{1t} = 0.5$ and $\alpha_{2t} = 0.5$. If source pixel 1 is suddenly occluded then $Z_t < Z_0$ because half of the light is blocked. Both projectors initially increase brightness, However, $\Delta\alpha_{2t}$ becomes dominated by $(\Delta\alpha_{2t})_{\text{SE}}$ which forces source pixel 2 to be bright. On the other hand, $\Delta\alpha_{1t}$ becomes dominated by $(\Delta\alpha_{1t})_{\text{LS}}$ since the screen pixel does not change when α_{jt} is changed. This forces source pixel 1 to be dark. Note that even when source pixel 1 becomes unoccluded, nothing changes if source

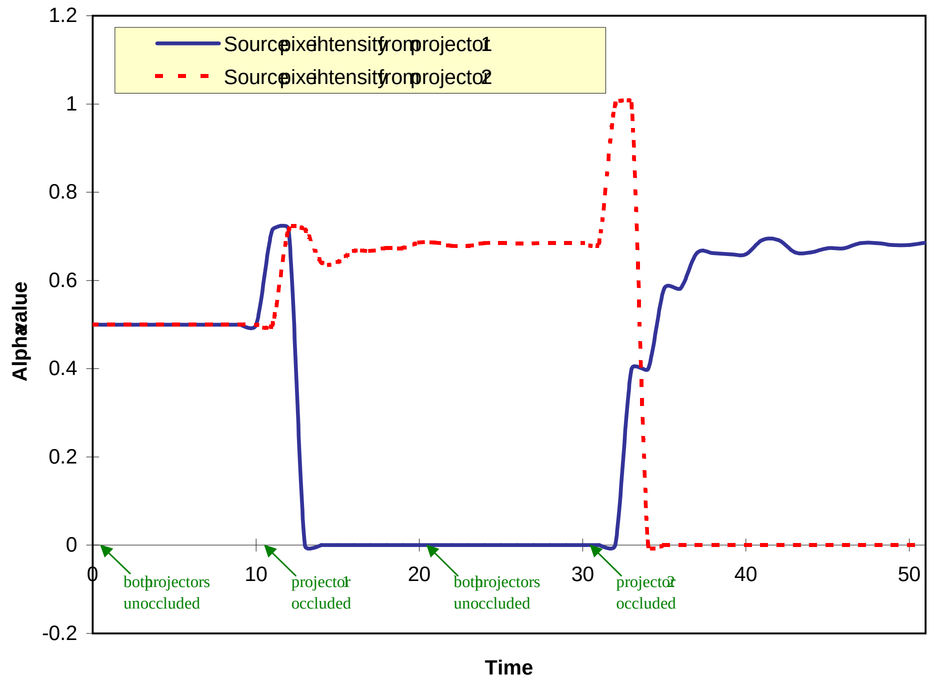

Figure 16: Synthetic example of transitions in projector source pixel intensities. This graph shows the intensity transition of two corresponding projector source pixels over time, subject to four events of occlusions and deocclusions. Note the hysteresis effect in which the source pixels are not boosted or blanked until new occlusion events occur.

pixel 2 remains unoccluded since the shadows have already been satisfactorily eliminated. This particularly illustrates the *hysteresis effect* in which source pixels are not boosted or blanked until new shadows are created – the system does not automatically return to an original state, nor change as a result of deocclusion.

Since we do not have good photometric models of the environment, projectors or camera, we cannot predict precisely how much light is needed to remove a shadow. However, the iterative feedback loop used to update the alpha mask allows us to avoid this problem: the system will continue adding light to shadowed regions until the region appears as it did in the reference image. Similarly, the system will blank projector source pixels which are occluded and do not affect the observed images. This approach has additional benefits. For instance, the system does not require an accurate photometric model of the shadow formation process to correct for occlusions with non-binary visibility ratios, e.g. the diffuse shadow created by a hand moving near the projector. The drawback to such an iterative technique is that the alpha mask can require several iterations to converge; in practice, shadows are eliminated in approximately 5–7 iterations.

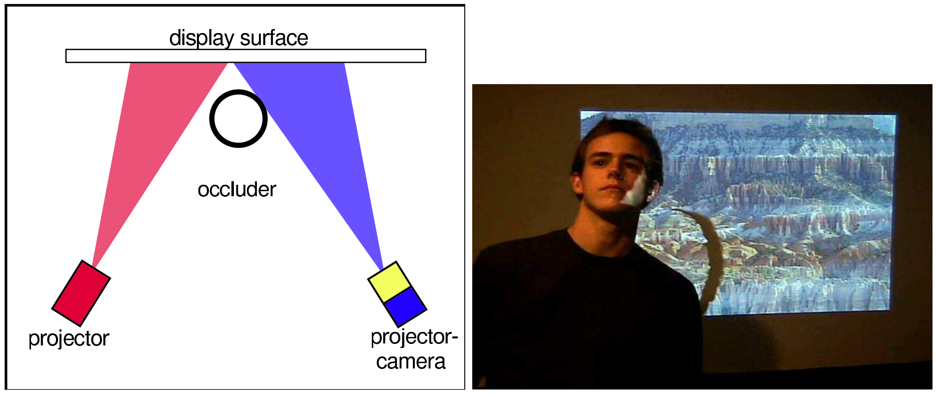

Figure 17: *Left:* Switching VRP. *Right:* Shadows are eliminated and blinding light is suppressed with a moving user. The gap in the display caused as the user moves into the scene will be corrected in the next iteration.

In our software only implementation, the AVRP-BLS system is able to calculate 1.6 iterations per second (See Table 2) Even assuming advances in processing power, when using commodity projectors, which are limited to 60 or 85fps, a series of 5–7 iterations would produce a visual artifact for up to $1/10^{th}$ of a second [1]. 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 which operate at 120fps 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.

5.3 Switching

The previous systems provide redundant illumination to each pixel 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 un-obstructed view of the display, we will have to correct the projector's output blindly, without feedback. To do this successfully, each pixel on the display surface is illuminated by only one projector at a time. As the projector illuminating a pixel

[1]As with the active shadow elimination system, the largest intensity changes happen in the first or second iteration. As the iterative feedback loop converges, subsequent iterations are much less noticeable.

is occluded, responsibility for illuminating that pixel is shifted to another (unoccluded) projector. 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 which 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 sub-images 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.

5.3.1 Occlusion Detection

In our approach, we chose the projector that was less likely to be occluded and designated it as the *primary* projector, responsible for the entire display by default. We positioned a camera close to the projector lens of this 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 [47]. 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 IR camera images is not affected by light projected from the projectors and, as shown in Figure 21(b), the back-lit silhouette of occluders creates a strong contrast between foreground and background.

Because we are detecting occluders (instead of shadows) we do not need to pre-shoot background plates for each expected frame [62] or predict the expected appearance of each image when projected onto the display surface [30]. **This is a significant advantage when projecting arbitrary interactive graphics.**

For each compensation step, the IR camera image must be processed to meet the challenge

Figure 18: Boundary between regions of varying projector ownership. *Left:* before seam blending. *Right:* after seam blending.

of preserving high image quality in the face of varying pixel-projector ownership. These steps are illustrated in Figure 21. First, the acquired image must be warped to align with the display surface using a camera-surface homography. Second, the image is segmented into occluder and non-occluder regions. Our implementation uses background subtraction. In some cases, median filtering is needed for noise removal, but in our experiments the back-lit 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. Figure 18 illustrates the necessity for blending to avoid distracting seams.

5.3.2 Photometric Uniformity

The projected display from one projector must appear photometrically uniform to another projector to insure the VRP displays consistently. Calibration for photometric uniformity is necessary to make the hand-off of a pixel from one projector to another unnoticeable.

Majumder and Stevens have found that the major source of apparent color variation across multiple projectors is primarily due to luminance variation, and that the chrominance of projectors (of the same brand) are very similar [40, 38]. Their work has focused on tiled multi-projector displays where the projectors are oriented perpendicular with the display surface.

In a virtual rear projection 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 over-illuminated, while the far side is under-illuminated. This angle-induced ramp 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 being used to illuminate the pixel.

The method we use to generate the attenuation maps is similar to those used by Majumder and Stevens for their Luminance Attenuation Maps (LAM) [39] 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 one step at a time (to account for non-linear projector and camera responses) until the target intensity is uniform across the display. Figure 19 shows two example LAMs and the following pseudo-code describes our simple algorithm for their 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 l = UPDATE_LIMIT
  project black for all other projectors
  while l > 0
    project LAM(*,p) and capture image
    for each pixel i
      if (intensity(i) > d*)
        LAM(i,p)--
      end if
    end for
    l--
  end while
  low-pass filter LAM(*,p)
end for
```

5.3.3 Edge 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

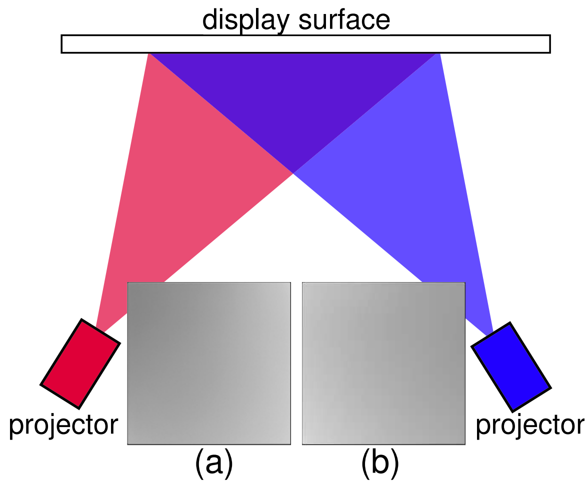

Figure 19: *Luminance Attenuation Maps (LAMs):* (a) LAM for projector positioned to the left of projection surface (b) LAM for projector positioned to the right of the projection surface. Note that the dark regions of each LAM correspond with the shortest projection distance to the display surface.

using edge blended alpha masks to limit the output of each projector, generated as follows:

1. Order your projectors from $P_0 \ldots P_N$. 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 projector P_0 and P_1 will be handled by projector P_2 and so on through P_n. Associate an initially zero alpha mask with each projector $\alpha_0 \ldots \alpha_N$ which will be used to control the active output pixels.

2. Generate an occlusion mask $O_0 \ldots O_N$ for each projector, indicating which projector pixels are occluded.

3. For the alpha mask of the ith projector $\alpha_{0<i<=N}$ turn on all pixels which are not occluded in the occlusion mask O_i and have not already been turned on in any previous alpha masks $\alpha_{0 \ldots i-1}$. This results in a set of mutually exclusive alpha masks which favor projectors based on their ordering. A pixel must be occluded in all projectors before it will be lost.

4. We then perform the following operations on each alpha mask to add a feathered edge which hides the seam:

   (a) Filter each alpha mask $\alpha_0 \ldots \alpha_N$ with a $3x3$ median filter to remove noise.

(b) Dilate each alpha mask three times to expand their extent.

(c) Blur the expanded alpha masks with a Gaussian filter to feather their edges.

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.

5.3.4 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 switched VRP may be optimized by exploiting today's programmable graphics cards (GPUs). Masters student Matthew Flagg moved the switching algorithm onto the GPU, translating OpenCV operations into programmable vertex and texture shaders. I subsequently integrated this code into the PROCAMS toolkit (Chapter 6). Image processing on the GPU shifts the speed limit of switched VRP away from computation on the CPU to capture and display rates of the camera and projector. Figure 20 illustrates our image processing pipeline using the GPU and Figure 21 gives example textures at each stage.

There are three capabilities of GPUs and DirectX 9.0 that we exercise in order to eliminate the bottleneck of image processing: (a) multiple render targets, (b) pixel shaders and (c) multi-head resource sharing. First, the Multiple Render Targets (MRT) 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.

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 [13]. 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 pixels 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 sub-sampled 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 used by them is illustrated in Figure 20.

Finally, multi-head 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 multi-head resource sharing, one display device performed $2n$ texture samples and the other sampled $2(n + k)$ pixels ($4n + 2k$ total samples). After multi-head sharing, a dilation using $2n$ texture samples is shared among both display heads and a 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 performed for each display head separately due to differences between the occluder masks.

5.4 Quantitative Evaluation of Virtual Rear Projection Methods

To evaluate their 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 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

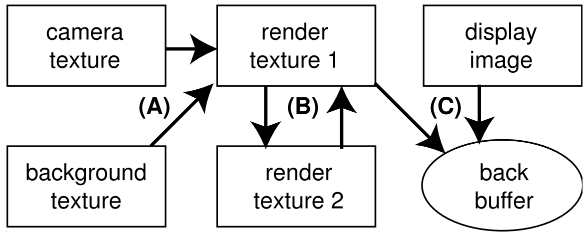

Figure 20: *Pixel Shader Pipeline:* Boxes represent textures and arrows denote texture sampling operations used in pixel shaders. (a) Background subtraction shader stores result in render texture 1 (b) Render textures 1 and 2 are used as sampling buffers for dilation and blurring operations, each of which require 2 independent shaders (c) the final occluder mask is composited with a display texture and rendered into the back buffer for display.

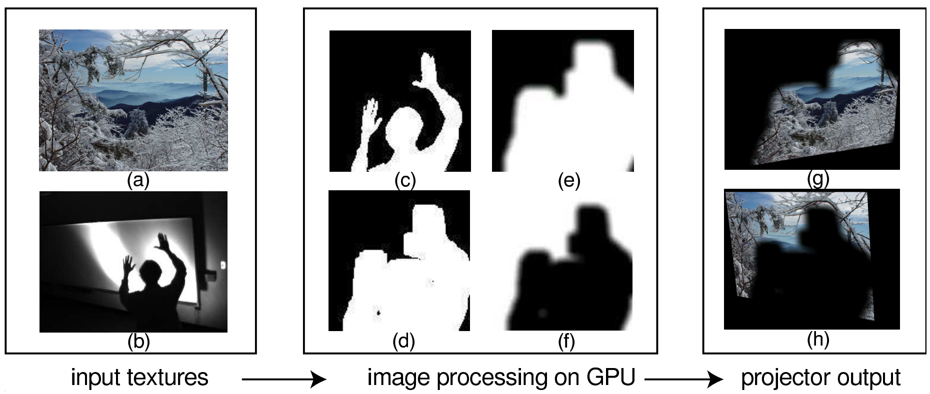

Figure 21: *GPU-centric architecture:* (a) display texture (b) IR camera frame (c) occluder mask texture (d) dilated mask to tolerate inter-frame occluder movement (e) blurred mask for projector 1 blending (f) blurred mask for projector 2 blending (g) keystone-corrected projector 1 output (h) keystone-corrected projector 2 output.

an occluder into the beam path of one projector and re-started the algorithm.

We used a static occluder which 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* (SSD) 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. An overview of the results are presented in Figure 22 and Table 2.

Table 2: Algorithm Performance Measures

Condition	Frames to Converge	SSD Er ror	Occluder Light	F.P.S.
WFP	n/a	3379	166	23.8†
PVRP	n/a	2509	167	23.7†
AVRP-SE	7	1052	221	23.3†
AVRP-BLS	7	1165	34	1.6
Switching	1	1466	12	9.5‡

† Because WFP and PVRP do not actively compensate for shadows, their frame-rate scores represent the sensing limitation of our 30fps camera and evaluation code. AVRP is only slightly slower than the passive solutions.

‡ 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 algorithms. The GPU version of the switching algorithm runs at 85fps, limited by the refresh rate of our projectors.

5.4.1 Experimental Setup

Each algorithm was run on a dual processor Pentium4 Xeon 2.2Ghz Dell Precision workstation with 2 GB 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 430cm apart on a bar 360cm from the display surface, 240cm above the floor. The display surface was 181cm wide and 130cm high, mounted so that it's bottom was 63cm from the floor. Each projector was 34° off of the projection surface's normal, giving a total angular separation of 68° between the projectors.

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

After the occluder was introduced, the system was re-started. To the algorithms, this gave the appearance of an instantly appearing occluder which blocked approximately 30 percent of one projector. In the graphs, the occluder appears in frame five.

At this point, the algorithms were allowed to run normally until they had stabilized.

In the simple cases of warped front projection and passive virtual rear projection, the system performed no compensation, and the light on the occluder and errors in the displayed image are immediately stable. As you can see from Table 2 (SSD Error) and the graphs in Figure 22 passive

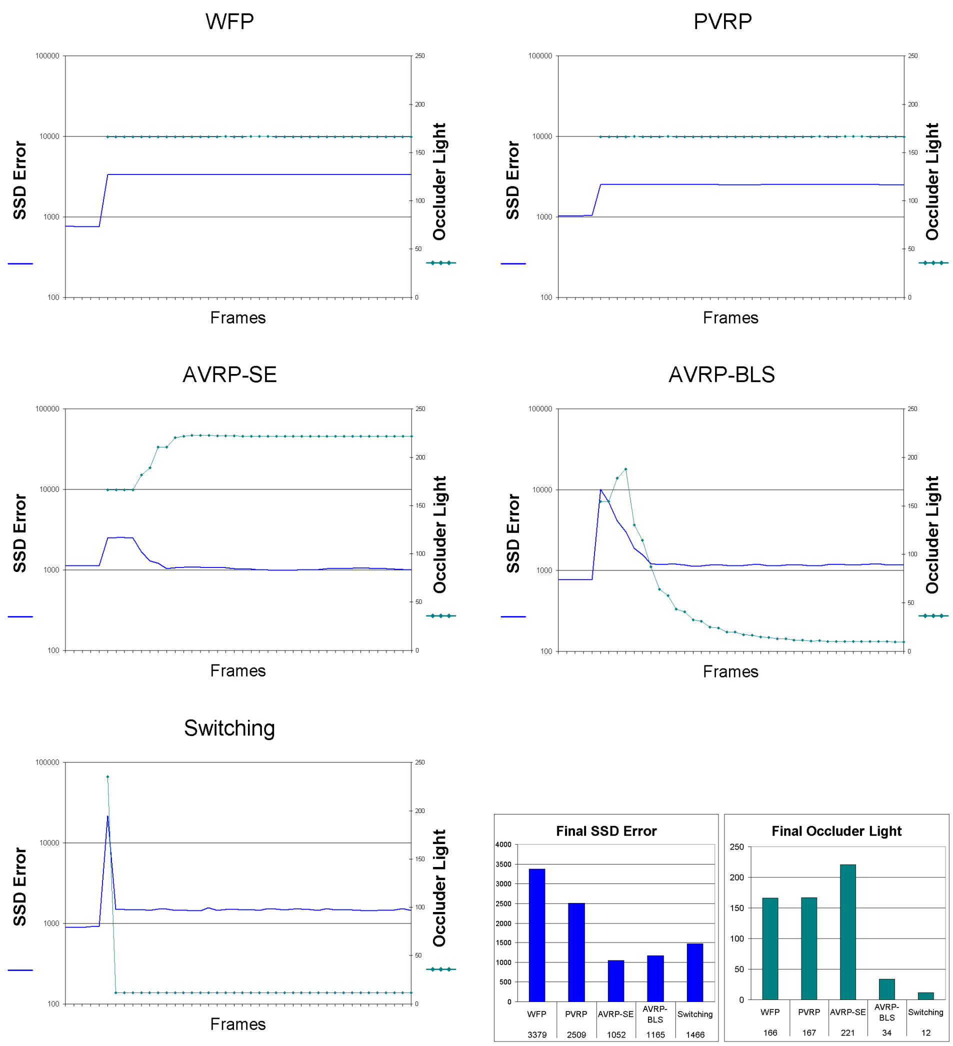

Figure 22: Top Left: Warped Front Projection **Top Right:** Passive Virtual Rear Projection **Middle Left:** Active Virtual Rear Projection - Shadow Elimination **Middle Right:** Active Virtual Rear Projection - Blinding Light Suppression **Bottom Left:** Switching Virtual Rear Projection **Bottom Right:** Final SSD and Occluder Light Measures

virtual rear projection improved the image quality over that achieved by a single projector solution (Warped Front Projection) despite taking no implicit compensatory action.

Shadow elimination, 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 non-active 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.

The shadow elimination with blinding light suppression system (AVRP-BLS), also took seven iterations to converge, but due to the increased processing required by this algorithm, 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 as 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 shadow elimination or blinding light suppression cases, 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.

Chapter VI

PROCAMS TOOLKIT

The PROCAMS (Projector/Camera) toolkit is a collection of software modules that ease the development of applications using projectors and cameras together [65]. It consists of hardware interface components, computer vision components, and utility classes that ease the development of multi-projector applications. The PROCAMS toolkit has been developed in conjunction with the implementation work needed to deploy and evaluate the projection technologies described in Chapters 3 and 5.

In addition to a toolkit that can be used via the programming API, a utility program (WinPVRP) was constructed using the toolkit and has been released as a ready to install windows application. A computer with the appropriate hardware (minimum 2 video outputs with one of them connected to a projector) can use this utility program to create a Warped Front Projection and (with two projectors) Passive Virtual Rear Projection display. In addition to this general purpose application, various sample applications using the PROCAMS programming API are bundled with the PROCAMS toolkit download. These sample applications can be studied by programmers to see how the toolkit is used in actual applications, or used as a base from which to build similar applications.

The hardware interface components are divided into input (cameras) and output (projectors). The input components standardize camera input from different API's such as VideoForWindows, CVCam, and the Matrox camera interface into a generic camera input object. This allows any camera that supports one of the above mentioned interfaces to be used by an application developed using the PROCAMS toolkit. Although the VideoForWindows interface is specific to Microsoft Windows, the CVCam and Matrox camera interfaces are supported on Linux.

The output components take advantage of the DirectX API to use hardware acceleration to quickly warp images with a projective transform, allowing "square" images to be projected onto arbitrary planar surfaces in the environment. The reliance on the DirectX API currently limits the toolkit to computers running a Microsoft Windows operating system, but porting these output components to use OpenGL would allow the toolkit to run on many POSIX/Unix based operating systems that support OpenGL.

The computer vision components implement basic algorithms useful for calibration of projectors and cameras, as well as detecting users. For example, they are used to calibrate multiple projectors via a camera to project warped images so that the output from each projector overlaps with the other projector's image to form a single image on the display surface. These computer vision components are built on top of the open-source OpenCV library, and wrap the low level computer vision algorithms, abstracting them to a much higher level operation for the programmer. They are only dependent upon the OpenCV library, and would work on any platform for which the OpenCV library has been ported (currently, Microsoft Windows and Linux).

The utility classes bundle together functionality, using the input and output classes together with the computer vision components to ease the creation of multi-projector displays. In addition to the work presented in this document, the PROCAMS toolkit has been used to prototype a capture resistant environment [71], and multi-planar display system [1].

6.1 PROCAMS Abstractions

PROCAMS supports three main features: enhanced keystone correction via warping, the calibration needed to align multiple redundant projectors into a redundantly illuminated display, and algorithms to detect occluders and project compensated images. It abstracts the 3D programming, camera access API's, and computer vision techniques needed by programmers to deploy novel projected applications quickly. These abstractions allow a programmer to concentrate on the application functionality, not the graphics and computer vision programming needed to display images from multiple, arbitrarily-positioned projectors.

In the simplest case, PROCAMS allows a programmer to warp the output of a single projector onto an arbitrary planar surface using a projective transform performed by the accelerated 3D video card (See section 6.2.2, and Figure 23). This *warped front projection* (WFP) allows a projector to be placed in an arbitrary location with respect to the display surface.

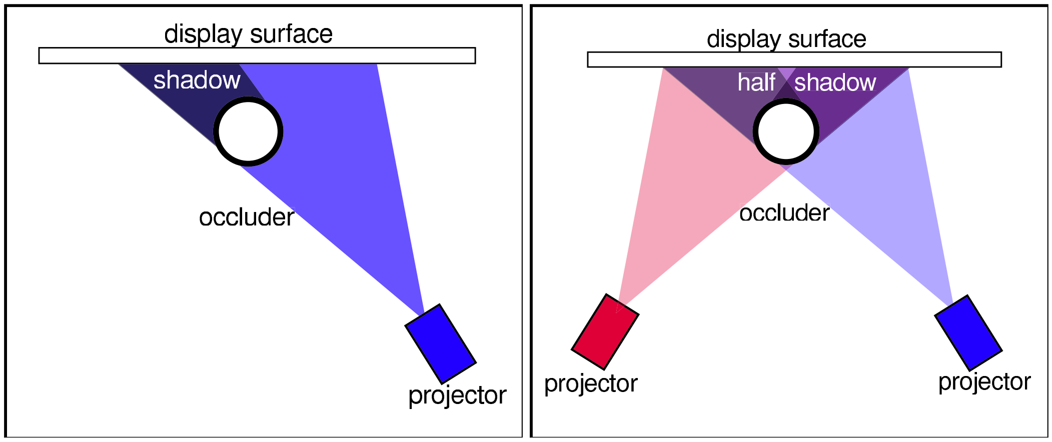

Figure 23: left: A Warped Front Projection (WFP) display. The enhanced keystone correction allows more freedom in projector placement. **right:** A redundantly illuminated display (Passive Virtual Rear Projection) uses two or more projectors to increase brightness and provide robustness in the face of occlusions and shadows.

Although warped front projection can be a useful tool to easily position projectors, redundant illumination is the key feature provided by PROCAMS that other software does not offer. Redundant illumination allows users to approach the display surface without completely occluding the display with their own shadows, providing a user experience similar to rear projection. Figures 24 and 27(left) illustrate users interacting with redundantly illuminated displays which are robust to shadows. These displays are created by adding a camera and second projector to the system. PROCAMS handles the computer vision needed to calculate the homography between each projector and the camera. By using the camera's view as a frame of reference, multiple projectors can be calibrated so that their output overlaps on the display screen (Figure 23) forming a PVRP display.

6.2 PROCAMS Applications

We have used the PROCAMS toolkit to build dedicated applications (such as the interactive game in Figure 24, and the banner display in section 6.2.2) as well as the WinPVRP application. The WinPVRP program is a solution for users attempting to implement a warped front projection or passive virtual rear projection display. Programmers can download and use the underlying C++ based PROCAMS toolkit to experiment with multi-projector systems and build custom applications.

Figure 24: An interactive game using redundant illumination provided by PROCAMS. The redundant illumination prevents shadows from hampering the game-play.

6.2.1 Redundant Illumination - WinPVRP

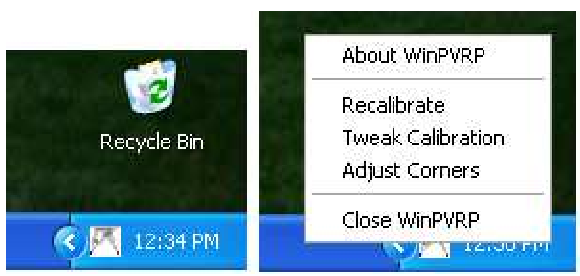

Figure 25: WinPVRP tray icon and menu.

At Georgia Institute of Technology, the School of Aerospace Engineering has retrofitted a classroom into a Collaborative Design environment (CODE) (Figure 26). The CODE provides student design teams experience solving design problems in collaborative team rooms, which are becoming more common in the workplace. The design of the CODE includes several interactive, upright large-format computer displays. However, because of space and cost constraints, rear projection screens

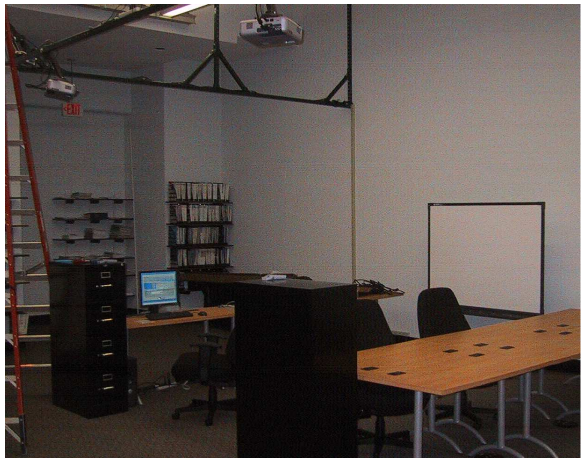

Figure 26: Breakout Area 1 in the Collaborative Design Environment (CODE) at the School of Aerospace Engineering.

could not be installed. We used PROCAMS to build a Windows tray application that allows a standard Windows desktop to be projected using passive virtual rear projection. The two projectors were mounted on the left and right sides of the touch sensitive surface (as in Figure 23). This positioning, combined with the redundant illumination, provides robustness to occlusions and almost eliminates shadows. Figure 27 shows displays created using dual projectors and the WinPVRP application.

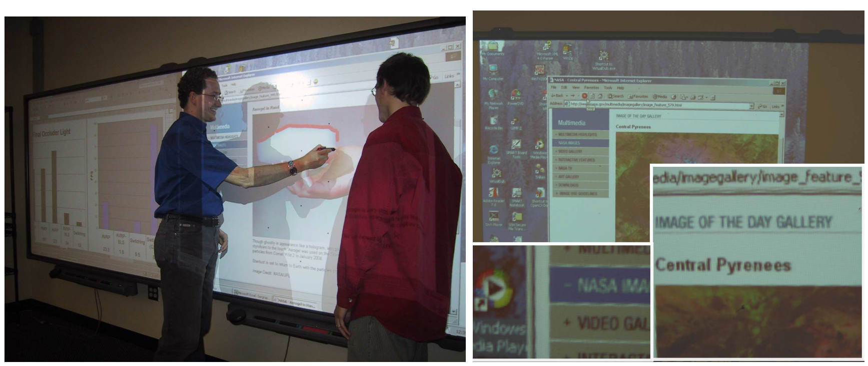

Figure 27: right: The WinPVRP application provides camera based calibration of dual projectors to provide a passive virtual rear projected (PVRP) display surface. The redundant illumination provided by dual projectors allows users to approach, and interact with, the surface without completely occluding it. Although users cast “half-shadows”, graphics are still visible within the semi-occluded regions. **left:** The calibration accuracy can be seen in the two enlarged call outs at the bottom of this figure illustrate.

The WinPVRP application (Figure 25) allows users with a Windows desktop and two projectors (3 total video ports) to create a passive virtual rear projected display using any Video for Windows device (such as a USB webcam) to calibrate the two projectors.[1] If the WinPVRP application only detects a single projector, it will automatically fall back into warped front projection mode. WinPVRP provides an easy way to take an existing Windows application (or even a windows manager such as Scalable Fabric [56]) and project it onto a touch-sensitive interactive surface using passive virtual rear projection so that user's shadows do not occlude the display.

6.2.2 Warped Front Projection - Banner Display

The Banner program reads lines of text from a file and renders the text onto a sign. We used it to implement a Trolley Timer (Figure 28), which displays the predicted wait time for the next few trolleys at the stop outside of our building (using GPS data). The best place to locate the Trolley Timer sign was on a hallway wall at a "T" intersection. This location was chosen due to the location of windows and doors that precluded other locations, as well as the normal traffic flow patterns in the building. Unfortunately, the hallway at right angles to the chosen wall had no good locations to place a projector. To mount a projector in the correct location to project the sign, a projector mount would have had to be installed by facilities workers. This would have increased the project cost, and significantly delayed deployment.

The banner application, created using PROCAMS (the code in Sections 6.4.1 & 6.4.3) allows the user to position the display at the desired location, while placing the projector at an extreme off-axis angle. By adding one line of code (*display->userMouseOutCorners();*), the programmer allows the user to interactively specify where each corner of the display should be placed. Mouse input, calculation of the correct projective transforms, implementation of the projective transforms on the 3D graphics hardware, and feedback to the user are all handled by the PROCAMS toolkit. The projector was placed in an existing cabinet, and the warped front projection allowed the sign to be projected correctly in the desired location.

[1] Manual calibration of two projectors is also possible, but use of a camera greatly speeds the process.

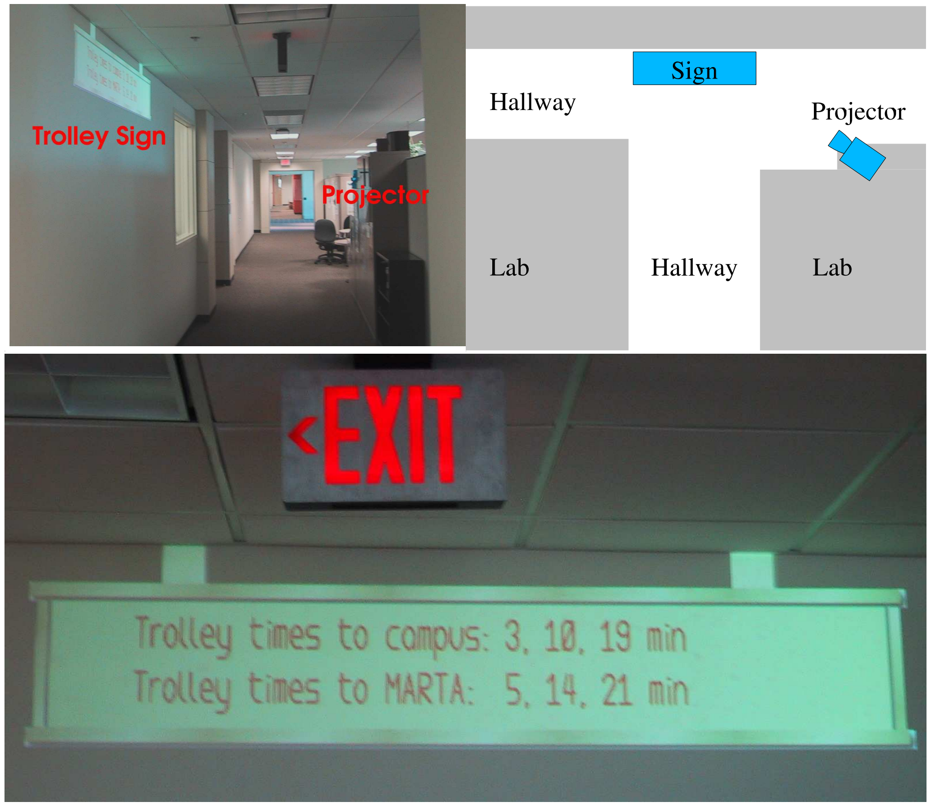

Figure 28: Trolley Timer sign environment and floor-plan.

6.3 PROCAMS Architecture

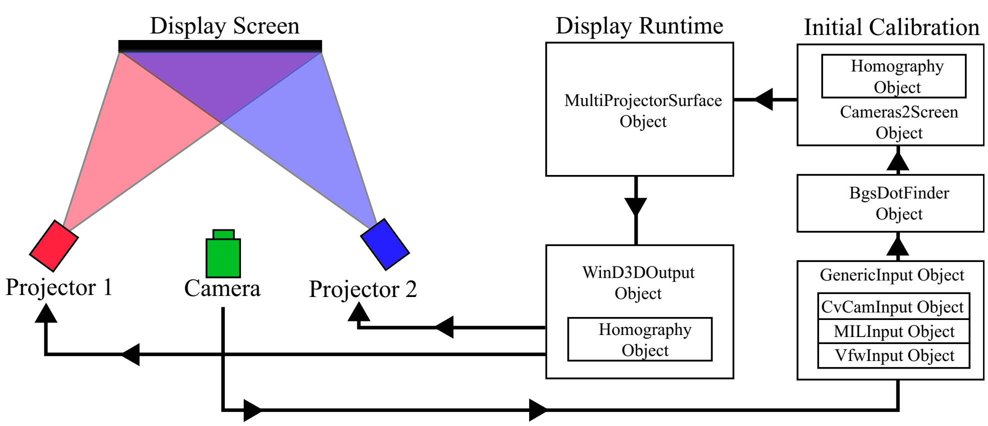

Figure 29: Architecture diagram of the PROCAMS toolkit showing data flow for calibration and use.

PROCAMS has three main functional components with which a programmer interacts:

1. MultiProjectorSurface - This object represents a single display “surface” which can be made up of one or more projected outputs. The user adds cameras and projected outputs to this object, and it handles the computer vision needed for calibrating multiple projectors. The MultiProjectorSurface also provides user interface mechanisms for an end user to position the display interactively using the mouse.

2. GenericInput - PROCAMS supports three different camera API’s: Video For Windows, Matrox Imaging Library (MIL), and the CVCam interface provided by OpenCV. This allows various USB webcams and more professional IEEE 1394 (Firewire) cameras to be used. Each camera interface is a subclass of GenericInput. A user creates an object to interface with the specific camera they have, and passes it to the MultiProjectorSurface via the *addCamera* method after casting it as a GenericInput.

3. WinD3DOutput - This object handles full-screen window creation and image warping using the 3D graphics card. Programmers use the WinD3DOutput object to “grab” one or more video ports (connected to projectors) in full-screen mode. The WinD3DOutput object is then given to the MultiProjectorSurface, which uses the projector(s) in creating the display.

Figure 29 shows the data-flow through these three components. In addition to these three programmer visible objects, the math and vision routines needed to calibrate multiple projectors and calculate the appropriate projective transform to warp their outputs are encapsulated within three objects that are used internally by PROCAMS. The following three objects are hidden from the casual programmer:

1. Homography - These objects encapsulate the math needed to calculate a homography between two planes. It is used by the Cameras2Screen object to calculate the relationship between projectors and cameras, as well as by the WinD3DOutput object to calculate the appropriate warping for a projected image. The Homography object will also be useful to advanced programmers who wish to calibrate any two planes, such as an input surface and a projected display.

2. BgsDotFinder - This object uses GenericInput objects to access a camera feed and encapsulates a background subtraction and "Dot Finder" computer vision algorithm. It is used by the Cameras2Screen object to detect projected calibration patterns. Advanced programmers can use the background subtraction routines from this object, useful as the first step in detecting human activity.

3. Cameras2Screen - This object handles the projection of calibration patterns, their detection via a camera, and the calibration and alignment of multiple projectors into a redundantly illuminated display.

As shown by the code samples in Section 6.4, the default interface to PROCAMS is relatively easy to use. Programmers allocate one or more projectors (via the WinD3DOutput object), an optional camera (via one of the Input objects, cast to a GenericInput) and give these objects to a MultiProjectorSurface, which handles the calibration and user interface for display placement. From that point forward, the programmer is free to create the desired graphics which are handed off to the MultiProjectorSurface via the *drawImage* method. One feature not demonstrated by the code samples is that PROCAMS allows programmers to save calibration state between program executions to a file. This allows projector calibration and/or display placement to be done only on initial setup or when projectors are moved.

6.4 PROCAMS code samples

The PROCAMS toolkit provides hardware abstractions for camera input (used for computer vision) and warped output (using accelerated 3D hardware to provide massive keystone correction quickly), and tools for easily calibrating multiple projectors via computer vision. It also handles interactive display alignment and position specification by the user.

6.4.1 Allocating and Positioning a Display

The following example code grabs the 1st monitor (which is attached to the projector), adds it to a "MultiProjectorSurface" (which in this case has a single projector), asks the user to position the corners of the display interactively, and projects a welcome image:

```
#define VIDEO_OUT 0
```

```
WinD3DOutput* graphics;
MultiProjectorSurface * display;

// Get the screen attached to the projector
graphics = new WinD3DOutput();
graphics->grabScreen(VIDEO_OUT);

// Add the projector to the display surface
display = new MultiProjectorSurface( graphics );
display->addProjector(VIDEO_OUT);

// User positions display with the mouse
display->userMouseOutCorners();

// Initialization and calibration complete
// display can now be used for output:
display->drawImage( cvLoadImage("Hello.jpg") );
```

The code above is all that is required to set up a single projector display (Warped Front Projection) as shown in Figure 23(left). Once the user specifies where the display should be located, subsequent *display.drawImage()* calls will update the display. Although the above code could be used to set up a traditional front projected display, the main advantage offered by PROCAMS is the ability to warp the display so that it can be positioned at arbitrary locations with respect to the projector.

6.4.2 Calibrating Redundant Projectors using Computer Vision

The following code demonstrates the use of a Video for Windows camera (USB Webcam) to calibrate two projectors into a redundant display (Figure 23(right)). The cast of the vfwInput object to the *genericInput* type allows for the use of other types of cameras (PROCAMS currently supports the Matrox Imaging Library, Video For Windows, and the CVCam interfaces).

```
// We use 2 projectors
#define PROJECTOR1 0
#define PROJECTOR2 1
```

```
WinD3DOutput * graphics;
MultiProjectorSurface * display;
// We use the first camera.
#define CAMERA 0
vfwInput * camera;

// Grab the projectors...
graphics = new WinD3DOutput();
graphics->grabScreen(PROJECTOR1);
graphics->grabScreen(PROJECTOR2);

// Grab the camera
camera = new vfwInput(CAMERA);

// Add the projectors & cameras
// to the display surface
display = new MultiProjectorSurface(graphics);
display->addProjector(PROJECTOR1);
display->addProjector(PROJECTOR2);
display->addCamera( (genericInput) camera);

// Calibrate the projectors!
display->findHomographys();

//User positions display with mouse
display->userMouseOutCorners();

// redundant display ready
display->drawImage( cvLoadImage("Hello.jpg"));
```

The *display->findHomogrpahys()* function call is abstracting a large amount of calibration work. When this function is called, a calibration pattern is projected from each projector, detected by the camera, and the projectors are calibrated so that their displays are overlapped. The redundancy these multiple front projectors provide greatly increases the displays robustness to occlusions and shadows. Instead of casting full shadows on the display, users only cast “half shadows” within which the computer output remains visible. This provides a virtual rear projected display allowing users to approach and interact with it.

6.4.3 Native Image Format

PROCAMS uses the Intel Image Processing Library (combined with the OpenCV library) iplImage as its native image format. The OpenCV library provides methods for loading and saving iplImages to/from standard file formats such as JPG, GIF, TIF, etc. In addition, the OpenCV and IPL libraries provide basic drawing functions (lines, circles, arcs, polygons, text) for iplImages. As an example of generating images to display via PROCAMS, the following snippet of code (From the banner display example application of Section 6.2.2) loads a background image from a file and renders text onto it from a text file before displaying the final image.

```
#define RED CV_RGB(255,0,0);
// Load text from a file
char buff1[80];
char buff2[80];
FILE* f = fopen(“message.txt”,”r”);
fscanf(f,”%[^\n]\n%[^\n]\n”,&buf1,&buf2);
IplImage* sign = cvLoadImage(“sign.png”);

// Render text over the sign...
CvFont cvfont; CvPoint p1,p2;
p1.x = 150; p1.y = 85;
p2.x = 150; p2.y = 130;
cvInitFont(&cvfont,CV_FONT_VECTOR0,1,1,0,2);
cvPutText(sign,buff1,p1,&cvfont,RED);
cvPutText(sign,buff2,p2,&cvfont,RED);

// Display the image, free it.
display->drawImage(sign);
cvReleaseImage(&sign);
```

Displaying a video simply requires a looping construct to iterate through the frames:

```
CvCapture* movie = cvCaptureFromAVI( “c:\\movie.avi” );
IplImage* frame;

while (cvGetCaptureProperty(movie, CV_CAP_PROP_POS_AVI_RATIO) < 0.99))
{
  aviFrame = cvQueryFrame(movie);
  display->drawImage(aviFrame);
}
```

Chapter VII

EVALUATION STUDIES

The previous study explored a single-user task, where the user did not have to face the projectors. In this situation, users were able to complete the tasks and rated the Passive VRP condition higher than the single front projector conditions. We confirmed that users found blinding light from projectors to be annoying when they were facing them (Section 4.4). This motivated our development work on an active version of VRP (AVRP), that more fully removes shadow artifacts and eliminates blinding light. As reported in Chapter 5, we choose the Switching algorithm implemented on the GPU as the Active Virtual Rear Projection used in this work. The active compensation of AVRP prevents light from shining on users, and fills in shadows, but introduces some visible artifacts on the display surface. The studies described in this chapter were conducted in order to test AVRP with respect to the other projection technologies developed in this thesis[1] and to explore situations with more realistic tasks and collaborative groups of users.

7.1 Research Questions

Recall the overall thesis of this work:

By using a projector-camera system to mitigate shadows and blinding light, a virtual rear projected (VRP) display improves upon the user experience with respect to a traditional front projected display.

In Chapter 4 we compared Warped Front Projection, and a passive form of VRP, that mitigated shadows, with more traditional front and rear projected displays. The studies in Chapter 4 were conducted with a single user repeating simple tasks (moving boxes, hitting targets, following spirals) designed to emulate low-level GUI operations. The results show that individual users working on simplified tasks prefer PVRP to a traditional front projection display due to its ability to mitigate shadows. Additionally, users working with front projected displays adopted observable coping

[1]To simplify this study, Front Projection was excluded as a straw-man because previous chapters have already shown that WFP and PVRP are preferred by users.

behaviors not observed when they used PVRP. We also confirmed that users were annoyed by projector light that struck them on the face, and we worried that this “blinding light” from the PVRP condition would become a problem, especially as the size of interactive surfaces, and the number of projectors to support them increased.

This motivated the work in Chapter 5 to develop an active form of PVRP, AVRP, that would simultaneously eliminate shadows and blinding light. As shown in section 5.4, AVRP compensates for shadows and reduces blinding light better than previous work, but it is not imperceptible to users [64]. When compensating for occlusions, the seam between the two projectors is detectable despite the photometric uniformity and edge blending techniques (Sections 5.3.2 and 5.3.3) used to minimize the visual artifacts. Although not imperceptible to users, AVRP does operate at sufficiently high frame rates (75 Hz) to be interactive and we felt that it was ready for user evaluation. The studies in this chapter are designed to supplement our earlier user studies on Virtual Rear Projection by examining the following new features:

1. Active VRP (AVRP), which mitigates blinding light.

2. The use of more realistic tasks.

3. Multiple collocated users, both in collaborative-groups and in presenter/audience configurations.

Our general research questions for these studies are:

1. Do users prefer WFP, PVRP, or AVRP? What factors about the technology do users consider when forming their preferences?

2. Does the robustness to occlusions provided by redundant illumination (of PVRP & AVRP) *vs* a single projector (WFP) condition cause:

    (a) an observable effect on the user’s behavior?

    (b) a significant difference in the user’s preferences?

3. Do users find the active compensation of AVRP to be:

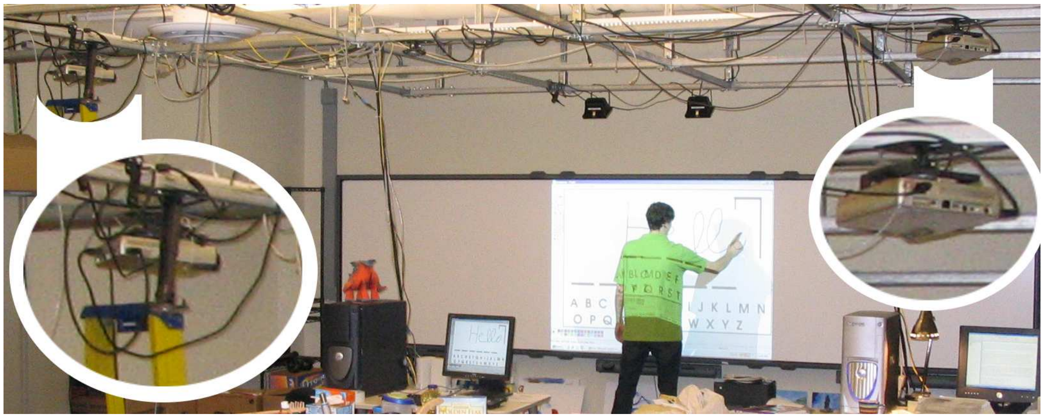

Figure 30: A view showing the two projectors (far left above ladder and far right), two IR lights (black, above the user's head), and the SmartBoard. The system is using PVRP in this photograph, and graphics are projected on the users back.

(a) noticeable?

(b) annoying?

(c) worthwhile enough to outweigh any drawbacks?

4. Does the blinding light elimination of AVRP cause

(a) an observable effect on the user's behavior?

(b) a significant difference in the user's preferences?

We are interested in identifying changes in user preference as well as observable differences in individual and group behavior, based upon the projection technology (condition) used.

7.2 Study Format

The study environment (TSRB Room 224, see Figure 30) has no exterior windows and is illuminated with standard office lighting (fluorescent lights). Two projectors are mounted on a unistrut beam 12' from the SmartBoard, with approximately 62 degrees of angular separation between the projectors. The projectors are 8' above the floor and separated by 14.5'.

The groups were introduced to the study, and asked to work on a collaborative problem (task) on a large interactive display for fifteen minutes, split into three five-minute sessions. The projection technology used (WFP, PVRP, AVRP) was changed for each of the five minute sessions in a counterbalanced order. At the end of the fifteen minutes, the group members were asked to fill out individual questionnaires, and then engaged in a focus group interview concerning the projection technologies used.

7.2.1 Tasks

Because we were interested in both collaborative problem solving tasks and situations where one person (the presenter, or driver) interacts with the display while observed by the rest of the group, we divided the study into two sub-studies with different tasks. Each sub-study had twenty-four participants, drawn from slightly different participant pools, and a different task.

The first sub-study used Aerospace Engineering graduate students who were presented with a representative task from their curriculum, while the second sub-study used general college students who played a game of Hangman. The Aerospace Engineering task is ecologically valid in that it closely mimics an actual task the students have received training on and would be expected to perform in their typical jobs. The Hangman game, although not a task you would find in a typical workplace, is designed to represent a collaborative group discussion and problem solving session around an interactive surface containing pertinent information. We used the well known and easily learned game so that it could be easily mastered by a general audience.

7.2.1.1 Aerospace Engineering Task - Quality Functional Deployment

The screen-shot in Figure 31 shows the design tool which was used by the participants in the Aerospace study. Their task was to solve design problems in the aerospace domain by using the design tool to select a single set of design options from a pre-specified design space that includes billions of possible combinations. The problems and pre-specified design space were originally prepared as an exercise in an Aerospace Engineering class. This is one stage in an Aerospace Engineering design process called Quality Functional Deployment that the participants were familiar with due to their educational program. In the task, users selected (and possibly un-selected) design options with check-boxes, as well as manipulated sliders at the bottom and right of the display as it was projected on the SmartBoard while discussing the resulting design alternatives.

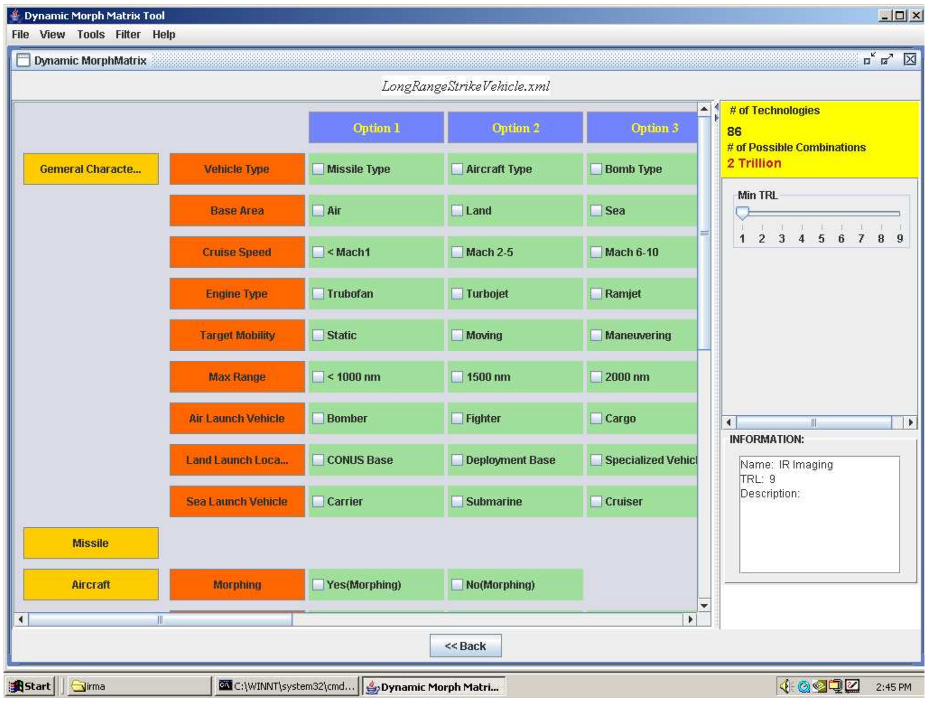

Figure 31: Missile analysis tool used for the task.

7.2.1.2 Hangman Task

The general college students were asked to play a game of Hangman. This task involves one participant (the "driver") drawing a card with a "secret" word on it and the audience made up of the rest of the group attempts to guess the word letter by letter. The driver marks down letters that are correct, and crosses off letters that have already been guessed, as well as keeping track of mistakes by drawing a figure. After each word, an audience member replaces the driver, who returns to the audience, allowing each group member to experience the task from both viewpoints. This task represents any activity where one person is driving an interactive application while interacting with an audience. Figure 32 shows the provided game-board for the Hangman game. Drivers could use their fingers to draw letters above the blanks, cross out letters from the alphabet at the bottom of the display, and keep score.

7.2.2 Rationale for Task Selection

The two tasks for this study involve multiple users working collaboratively. The Aerospace task includes group discussion and may result in some users turning away from the board into projector

Figure 32: The Hangman game-board, before game play has begun.

light, but we did not anticipate that users would spend a large portion of the time turned away from the board. Because the Aerospace users would primarily face the board, it was unlikely that they would notice graphics projected onto their backs.

The Hangman game task was designed so that the person driving the board (marking letters) would turn towards the projectors when they face the audience. Another important difference between these two tasks is that in the Aerospace task, all participants are collaboratively interacting with the SmartBoards simultaneously, while the Hangman task has a specific driver who interacts with the board, and the remainder of the group acts as an audience with whom the driver must interact. Thus, in the Hangman task, the driver is more likely to be looking back towards the audience and projectors. Also, the audience is more likely to notice any graphics that may be projected on the driver because they are located behind the driver.

7.3 Participants

Due to the varied nature of the tasks, two different participant populations were used. This complicates comparisons between studies, but allowed us to use a more ecologically valid task for the Aerospace study.

7.3.1 Aerospace Engineering Students (Aerospace Task)

Six groups of participants were used, each consisting of three to five members. The groups were made up of current or former graduate students from the School of Aerospace Engineering's Advanced Design Methods class. We selected this participant pool because the task was a decision support tool used in their class and in industry. Overall, twenty-four participants, made up of 18 males and 6 females with a mean age of 27.9 years ($\sigma = 6.85$ years) took part.

7.3.2 College Students (Hangman task)

Six groups of four participants were used. Individuals were recruited from the School of Psychology subject pool, via word-of-mouth recruitment, and via newsgroup posts to git.ads (a Georgia Institute of Technology advertising newsgroup) and assigned to groups. Although the members of one group were recruited together, and knew each other, the majority of groups were made up of strangers. Overall, twenty-four participants, made up of 16 males and 8 females with a mean age of 22.4 years ($\sigma = 4.32$ years) took part.

7.4 Experimental Procedure

The treatments (projection technologies) were varied in a within-group, counterbalanced manner with each group using each of the three projection technologies (WFP, PVRP, AVRP) for one of their five minute task sessions. The independent variable was the projection technology used. Data collected and analyzed as dependent variables include:

- Individual participant responses to questionnaires administered after they had used all three projection conditions.
- Individual responses to questions posed in a focus group interview.
- Time-lapse (1 fps) overhead camera view of the participant's occupancy of the space in front of the SmartBoard.
- Video footage of the participants interaction with the SmartBoard, captured from behind and to the right of the groups.

7.4.1 Research Procedure

The following linear procedure was followed when conducting the studies:

1. The researcher greeted the participants and completed the consent procedure. He then gave the participants a brief demonstration of the SmartBoard and associated software before introducing them to the task using the projection technology that they would be exposed to as condition number 1. The participants were then allowed to experiment with the SmartBoard until satisfied that they could perform the task.

2. Participants began to work on the task for fifteen minutes. Every five minutes the researcher announced the end of that condition, asked the participants to step away from the SmartBoard, and changed the projection technology used.[2] The researcher then announced "We are now starting condition number #" before resuming. WFP, PVRP, AVRP were used by each group in a counterbalanced order.

3. *After the third and final condition:* The researcher stopped the group and asked each participant to fill out a questionnaire individually.

4. After all participants completed the questionnaire, the researcher led the group in a focus group discussion about the three different projection technologies they experienced.

5. Near the end of the focus group (after question 7 from section 7.4.2) the research demonstrated the operation of the three different projection technologies to the group, and then proceeded with the final questions in the focus group interview.

7.4.2 Researcher Focus Group Questions

The following questions were presented (in order) to guide the focus group discussion. If some topics had already been covered based upon a previous question they could be skipped at the researcher's discretion.

[2] The changeover procedure took between 3 and 30 seconds depending upon which conditions were being switched between. The most overt gesture that the researcher had to make was pointing a remote towards the overhead projector when switching to or from WFP. The switch between PVRP and AVRP is accomplished with a few mouse clicks that are not visible on the main SmartBoard.

1. How did you like the SmartBoard? Do you feel that it helped you work on the task?
2. During the task we used three different types of projection technology. Did you notice any differences between the three different conditions?
3. Which condition did you like the most? The least? Why?
4. Did your opinion about the earlier conditions change after you saw the later conditions? Why?
5. What did you think about image quality on the different conditions? Which was best? Worst?
6. Did you have any problems with shadows in any of the conditions? If so, how did you deal with them?
7. What did you think of the light coming from the projectors?
8. *Before this question, the researcher brought the group back to the SmartBoard and demonstrated the three different conditions so that the group members could observe them again and ask questions.*

   Now, I'd like you to imagine that you have joined a small engineering firm as a new employee, and one of the first jobs your manager gives you is to spend $5000 upgrading their conference room with a SmartBoard and new furniture. Lets suppose that you have spent $3000 to buy the SmartBoard and two projectors, setting them up in the "two projector simultaneous" setup which you had for part number (X). You can also use the single projector warped mode if you choose. This leaves you $2000 to buy new chairs and furniture. The SmartBoard salesman says, "You know, we could upgrade your setup to a <AVRP> system for a $500 more." Would you be willing to reduce your furniture furniture budget to $1500 for the upgrade? If not, how much would you be willing to pay?
9. Is there anything that you think I'm forgetting to ask, or that you'd like to add?

7.5 Analysis & Results: Aerospace Task

When planning this study, we hypothesized that:

1. Users would prefer AVRP to PVRP because of (a) the reduction of blinding light and (b) lack of visible half-shadows, when standing between the projector(s) and the board.

2. Users would continue to prefer PVRP to WFP for the reason identified in our previous study (Section 4.3.1): reduction of full shadows on the display.

3. Users would report more annoyance with projected light in the PVRP and WFP conditions when compared to the AVRP condition, because AVRP reduces blinding light.

4. When using the PVRP and AVRP conditions, users would gather closer to the screen than when using WFP (due to the shadow elimination).

5. The tendency to gather closer to the display in the dual projector conditions (AVRP/PVRP) would increase collaboration.

7.5.1 Research Metrics & Analysis

To test hypotheses 1 & 2 we used the preference answers in our questionnaires for raw preference scores and used the focus group interviews to learn what reasons users gave as the basis for their preferences. We asked the users to rate each condition on a 7-point Likert scale.[3] The questionnaire also asked each user to rank order the conditions by preference, and had a free response area for them to write reasons for their choice. During the focus group interview the users were also asked to comment on why they liked or disliked specific conditions.

To test hypothesis 3, a second question on our questionnaire attempted to investigate how annoying light from the projectors was in each of the three conditions.[4] We also asked about the light coming from the projectors in the focus group interview.

To test hypothesis 4, we mounted a time-lapse video camera (capturing 1 frame per second during the studies) overhead to collect data on user's movement patterns. This overhead video data was programmatically analyzed using adjacent frame differencing with analysis of aggregate motion to determine the average distribution of the group.

---

[3] "Overall, how do you rate the display technology for the task performed...
Definite Dislike = 1 2 3 4 5 6 7 = Liked very much"

[4] "Did you find the light from the projectors to be...
Annoying = 1 2 3 4 5 6 7 = Unnoticeable"

When we designed the Aerospace study, we hypothesized that the two-projector conditions (which offer redundant illumination) would result in greater collaboration between group members (hypothesis 5). We hypothesized that the redundant illumination would allow group members to gather more evenly around the board, enhancing opportunities for collaboration. As a second objective metric for "collaboration," we decided to code the group member's interactions with the SmartBoard. A second video camera mounted facing the board behind the group collected video that was coded by two independent researchers recording each interaction with the board. Any interaction with the application that was completed before the user's hand returned to an idle position was counted as a single interaction. For example, checking a check-box, sliding a slider, or repositioning a slider by clicking repeatedly on a scroll button are all coded as single interactions with the SmartBoard. For example, the interaction log from the AVRP condition of the first Aerospace study looks like this: 3,3,3,1,1,1,2,1,1,4,1,2,4,4,4,3,1,3,1. Each number in the sequence represents a single user interaction with the board. In this session, participant three interacted with the board three times, followed by participant one, who also touched the board three times. Participant two then touched the board, followed by participant one, and so on.

The video was coded independently by two researchers and then the results were compared. The sequences from the two coders usually differed in only one or two places, usually due to a disagreement about how many times a particular individual interacted with the board (insertion or deletion errors), and not the sequence of changes between participants. The two coders would then review the video and agree on the correct sequence.

From this raw data we calculated the total number of interactions with the board, the number of interactions by participant, and the total number of changes between participants. For example, in the example sequence above, participant three interacted with the board five times, and there were eleven changes between participants. We hypothesized that the number of changes would be larger in the dual-projector conditions because more people would be able to stand closer to the board and take direct control of the application.

All between condition measures were analyzed using a repeated-measures ANOVA, using an $\alpha = 0.05$ criteria to check for statistical significance.[5]

[5]Where the sphericity assumption was violated, a Greenhouse-Geisser correction was applied.

7.5.2 User Preference

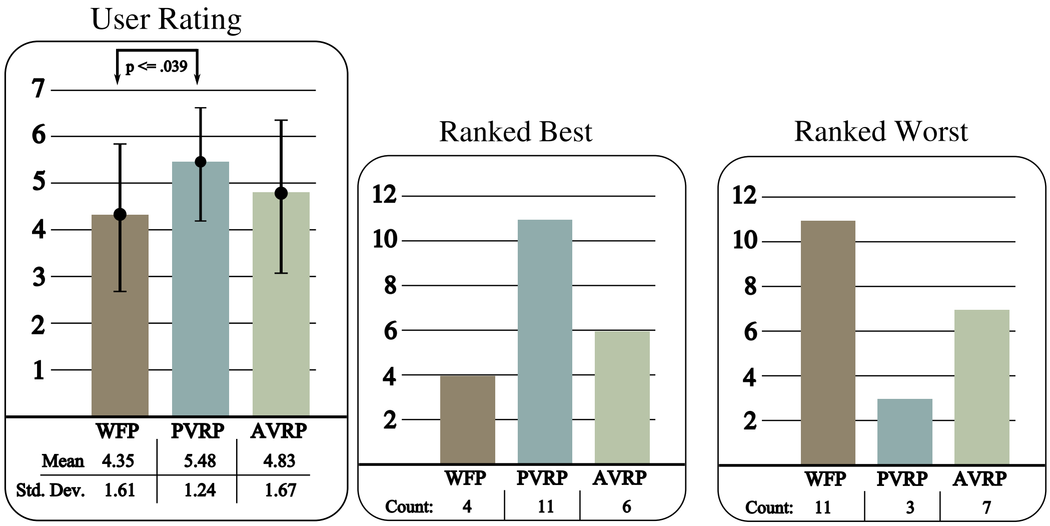

Figure 33: User rating scores, and forced ranking for the Aerospace task.

The user rating graph of Figure 33 shows the mean values (and Standard Deviation) for user responses to the rating question. Arrows (with p-values) at the top of the graph indicate when the values for two conditions show a statistically significant difference. Users preferred PVRP (mean rating 5.5) to WFP (4.3), *p <= 0.041, F(1.642,23)=3.737,* $\eta^2 = 0.592$.

AVRP was rated 4.8, but analysis revealed no significant difference between it and the other two conditions.

The results from the Likert scale rating question was consistent with the results from the exclusive choice questions. When asked to identify the condition they liked the most, 11 preferred PVRP, 6 preferred AVRP, and 4 preferred WFP. When asked for the condition they disliked the most, 11 chose WFP, 7 choose AVRP, and 3 choose PVRP.[6] A χ^2 analysis comparing these to a normal distribution (7,7,7) indicates that the differences are not significant. The dual-projector conditions (AVRP and PVRP) were prefered by 17 people, compared to 4 people who prefered the single projector condition (WFP), which is also not significant when compared to a normal (14,7) distribution.

Individual Questionnaires

On the individual questionnaire free response areas, users gave various reasons for liking and

[6]Of the 24 participants, 3 left these two questions blank, resulting in only 21 total responses.

disliking the three projection conditions. A single researcher coded questionnaire results, forming more general categories when responses were similar. These categories supported by more than two participants are reported below:

Liked WFP:

- *Not as blurry / better image quality [than two projector conditions]* (two participants)

Disliked WFP :

- Shadows (thirteen participants) including:
  - *Shadows* (eight participants)
  - *Shadows interfering with the use of the board* (three participants)
  - *Having to take actions to cope with shadows* (two participants)

Liked PVRP:

- *Less/Reduced Shadows* (ten participants)
- *Lack of visual artifacts (as opposed to AVRP)* (two participants).

Disliked PVRP:

No two participants agreed on a reason they disliked PVRP.

Liked AVRP:

- *Reduced full shadows* (four participants)

Disliked AVRP:

- Visual artifacts (nine participants) including:
  - *flicker/blinking* (three participants)
  - *projectors filling in areas differently* (two participants)

  - *weird overlapping shadows* (two participants)
  - *fuzzy halo shadows*, (one participant)
  - *lag* (one participant)

- *Intermittent failures to detect/correct occlusions* (three participants),
- *Couldn't predict where shadows would appear* (two participants)
- *Shadows overwhelm board when multiple people approach* (two participants).

Minority Responses

Some users gave reasons for liking or disliking various conditions that were not echoed by other participants, and were not generalized into categories. These views represent a minority opinion that was not volunteered by other users in their individual questionnaire results or are likely a mistake (e.g. "Less shadows" in the WFP condition) due to misremembering the order of conditions. However, some of these minority opinions re-appear in the focus group interviews and would sometimes gain more support. The minority opinions held by single participants are reported here:

Reasons for liking WFP: *Familiarity with single projector system*, *People able to stay out of the light*, and *Less shadows*. A reason for disliking WFP: *Having to figure out who is casting the shadow*.

Reasons for disliking PVRP: *Blurriness* (of the image), *Inability to determine source of shadows* (from dual projectors), *Didn't like half see-through shadows*, *Projector light more noticeable*, and *more shadows.*

Reasons for liking AVRP: *Easier to use/manipulate the board*, *Didn't have to worry about where others in the group stood {to avoid their shadows}*, *Projected light was less annoying*, and *Enjoyment of the novel visual effect* (This participant called it "The predator effect", likening it to the alien's cloaking mechanism from the popular 1987 movie). Reasons for disliking AVRP: *Harder to see, More shadows, Have to guess where things are on the board, More difficult to stay out of light.*

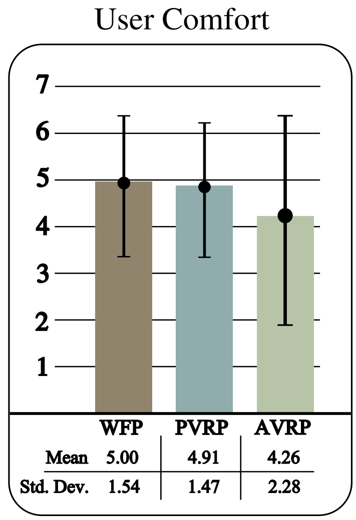

Figure 34: Self reported user comfort for the Aerospace task.

7.5.3 Annoyance of Blinding Light

During the study we observed no behavior (e.g. shielding of eyes with hands, squinting) to indicate that users felt the "blinding light" from the projectors was a problem. When the participants were asked about the light in focus group interviews the majority replied that light coming from the projectors was not an issue.

Many participants said that they had not noticed light from the projectors because they had been focused on the task on the SmartBoard:

> "I didn't look back."
>
> "I was facing the screen the whole time so I didn't notice anything."
>
> "I think it was more that I didn't look around much."
>
> "But in terms of what she said about looking back and being annoyed by the light, we never had an opportunity to turn around and be annoyed"
>
> "We didn't really turn around"
>
> " We didn't turn around, yeah."
>
> "When you were standing waiting between tasks you'd occasionally turn into the light, but when I was working on the task I never faced back."

Even participants who had turned away from the SmartBoard occasionally usually did not find the projectors to be annoying:

> “It wasn’t blinding, ’cauz it’s not, I did turn around I think once or twice, and I did notice that one of the other projectors had turned on. It’s at an angle from which, I don’t think, most people are tall enough, I didn’t turn around while I was at the board, but if you were standing back a few feet you’re not going to get a light in your eye.”
>
> “Yeah, maybe if you are elevated up a little more higher and you can see the projectors more, but since they are coming down it’s not really that easy, because you will be looking out at the audience, not up at the projectors, hopefully.”
>
> “Like shining on your eyes or something? No, not really.”

Only two people (in the same session) mentioned that the projector light was annoying, and another person hypothesized that it would be annoying if they had to turn towards an audience:

> “just that if I turned more towards this direction of the room, it was sorta...*Q: in your face?* yes”

> “I think if there had been an audience that wasn’t involved in the task and we had to turn around and talk with them, then maybe it would have bothered me”

One person disliked the projected light falling on the paper given to the group at the beginning of each condition that described the problem, but had not noticed blinding light from the projector on his face:

> "The only time that I noticed it was when I was trying to read the task, and so I was trying to find a place where I could not have light on the paper." "White paper shows the image very well"

Two other participants commented on the thermal output from the light beams and noise from the projectors:

> "I don’t know if it’s important for 5 minutes, but for a long study the light from the projectors might be annoying in terms of heating up, the a, room, or you know,

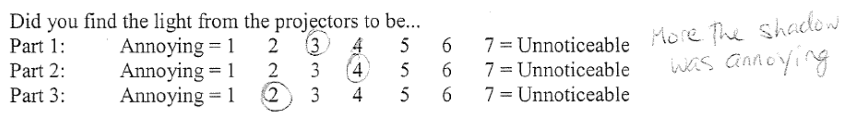

Did you find the light from the projectors to be...

Part 1:	Annoying = 1	2	(3)	(4)	5	6	7 = Unnoticeable
Part 2:	Annoying = 1	2	3	(4)	5	6	7 = Unnoticeable
Part 3:	Annoying = 1	(2)	3	4	5	6	7 = Unnoticeable

Figure 35: Margin note added by user to the comfort question.

> you have something like a warm beam..." *Q: So, were you feeling heat?* "I was feeling yeah...That was one of my first reactions to the whole setting, yeah, if I stay here for half an hour maybe ."
>
> "I would have been more annoyed about the sound instead of the heat...that was the first thing, when I walked into the room, Wow, these things making more noise than everything else."

Because the majority of our users did not report any annoyance from blinding light in the interview, we believe that they individually re-interpreted our comfort question and answered it based upon other factors. Because the only "light from the projectors" they had seen was on the SmartBoard, (and not shining in their eyes) they answered the "Did you find the light from the projectors to be...Annoying ...Unnoticeable" question based upon the image projected on the SmartBoard. The shadows cast by WFP, the half-shadows cast by PVRP, and the visual artifacts caused by the AVRP system were likely to all have affected answers to this question. Figure 35 illustrates that one participant even added a note to the margin of our questionnaire expressing how he had interpreted the question.

The only conclusion we can safely draw from this question combined with the focus group interview results is that the majority of users did not notice or suffer discomfort from the projected light.

7.5.4 Image Quality

In an attempt to determine if any of the three conditions (WFP, PVRP, AVRP) had a noticeably better image quality, one of the Likert scale questions asked the users about the perceived image quality of the display.[7] No statistical significant difference was detected between the three conditions, and

[7] "How would you rate the image quality of the projected display...

the average scores were 4.63 (AVRP), 4.71 (WFP), and 5.04 (PVRP), indicating that the perceived quality of all three technologies was good, but had room for improvement. As one participant put it, the display's image quality was "Good enough for what we were doing." (348) PVRP had the nominally highest score, but also had the largest σ of 1.52 (compared to 1.33 and 1.35 for WFP and AVRP respectively).

When participants were asked about image quality in the interview, their replies were equally mixed. We found that the factors they predominantly mentioned when asked about image quality were the brightness and clarity levels of the display. Some users felt that the added brightness of PVRP (from the simultaneous operation of two projectors) gave it a better quality image, while others felt that the slight blurring caused by the two projectors overlapping resulted in a lower quality image. Artifacts in the AVRP condition were rarely mentioned as reflecting negatively on image quality, and were predominantly mentioned when explaining why users had disliked AVRP. A few participants felt that the WFP display was crisper and less blurry than the dual-projector conditions, and others had not noticed any differences in image quality because they were engaged in the task:

> "No, I was trying to design an aircraft."
>
> "I think with the 2 projector system *(PVRP)* the screen is brighter and I like that, but I didn't like the 2 shadows that were cast, but with the one projector switching *(AVRP)* and the one projector *(WFP)* it wasn't as bright, and I didn't like that part about it."
>
> "I thought the image quality was best in the 1st one *(WFP)* the 2nd one *(PVRP)* was blurry.....so straight lines weren't straight and things like that. The third one *(AVRP)* you had like I guess like contrast or contrast issues where different areas were different levels of brightness and those were noticeable and you know like any artifacts from shadows and whatever..."

7.5.5 Mean Group Activity

The work in this section was performed in conjunction with Mario Romero, who analyzed the data using adjacent frame differencing as part of his research, we collaborated on the interpretation and

---

Poor Quality = 1 2 3 4 5 6 7 = Excellent Quality" analysis of the processed data. The overhead time-lapse camera video was processed by detecting differences in adjacent frames over time to record aggregate motion (See Figure 36. The graphs in Figure 37 show overhead activity maps displaying the average location of user motion averaged by condition. Each chart represents the motion of users averaged over five sessions (overhead video was not captured for session 1 in the Aerospace task).

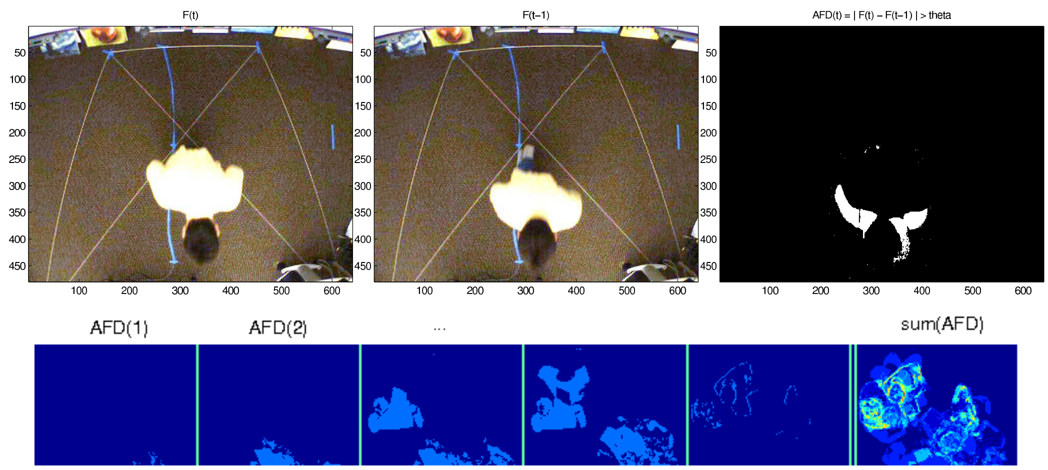

Figure 36: Visual explanation of the adjacent frame differencing method. The difference between temporally adjacent frames (*top right)* is summed over time to aggregate user activity.

In the WFP condition, users are clearly split by the projected light (entering diagonally from the bottom right towards the SmartBoard located at the top center) which results in the large (blue) area showing minimal activity near the middle of the room. The people to the right of the projector beam are standing forward, towards the wall and away from the projected light. The PVRP and AVRP conditions also show a bi-modal distribution, but those groups are much closer together, and when compared to the WFP condition, the right group is not pushed as far forward.

Ideal Group Position

To numerically compare these three activity maps, we have defined an "ideal" group layout based upon all users equally spaced around the SmartBoard in a semicircular area (Figure 38b). We have chosen this shape because 1) the hole in the center allows all users a view and physical access to the board, and 2) the circular shape also allows social access to other participants. Note that the camera is positioned slightly to the right of the center of the SmartBoard. To compensate, we positioned the idealized space usage image slightly to the left so that it was centered on the

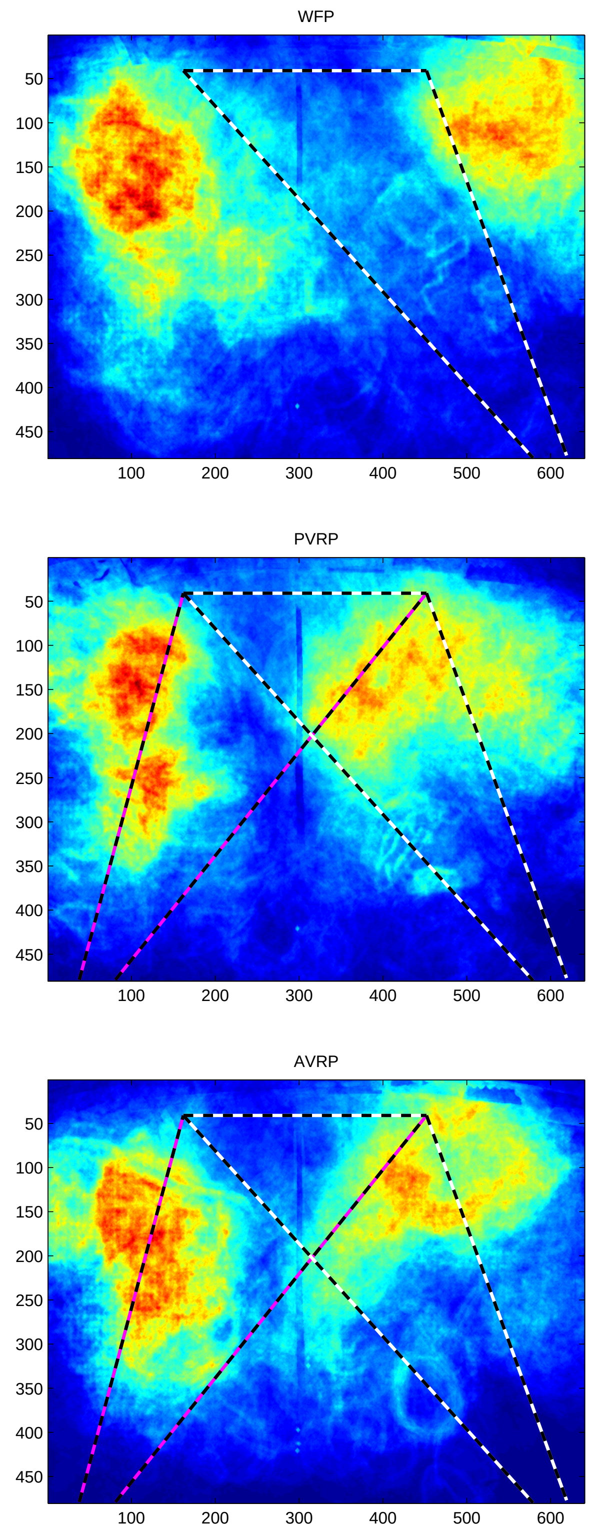

Figure 37: User motion by condition, with overlaid projector beam paths, in the Aerospace study. Horizontal and vertical axis are numbered by camera pixels.

SmartBoard, and not the image.

It should be made clear that we build this model only to get a numerical representation to quantify our observations. This "ideal" model is simply a representation of our subjective analysis, used for quantification of data, and not meant to be theoretically correct or ideal in a global sense.

We programmatically compare the average activity map for particular conditions by subtracting our ideal image from a thresholded version of our activity map and squaring the differences (making them all positive). The sum of these squared differences (SSD) is a metric of the difference between the average activity in each condition and the ideal model. This calculation is shown graphically in Figure 39. As the conditions progress from WFP (74.6%) to PVRP (76.1%) and AVRP (79.6%) the location of activity approaches the abstract ideal. To demonstrate that this calculation is stable with respect to the parameters specifying the model, we calculated the match with models of varying sizes (Figure 40) and demonstrated that while the absolute percentages may change slightly, the relative ordering of the conditions remain constant. For our "ideal" model, we chose the alternative (number 2) that gave the largest overall match.

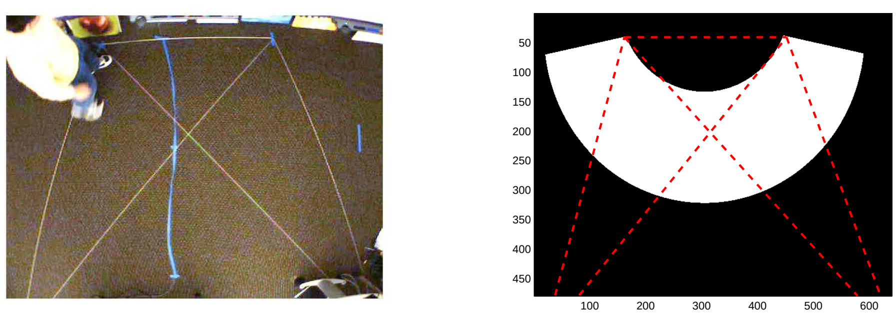

(**a**) (**b**)

Figure 38: (**a**) Overhead camera view of the experimental space. The SmartBoard is located just above the top of the image. The strings representing the projector beam paths were not shown to participants. (**b**) Idealized space usage superimposed over the overhead camera field of view.

7.5.6 Interaction Patterns with the Board

When looking at collaboration, we hypothesized that the number of times the person interacting with the board changed would be higher in the dual-projector conditions. We assumed that because more people would be able to stand closer to the board, they would share direct control of the application and more people would be involved in manipulating the Aerospace decision support application.

	WFP	PVRP	AVRP
Occupancy	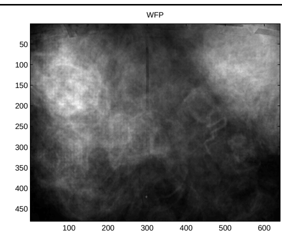	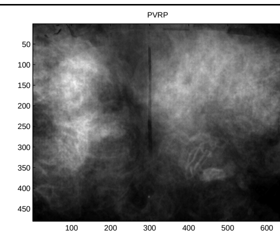	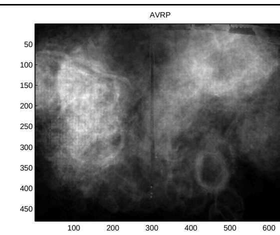
Thresholded	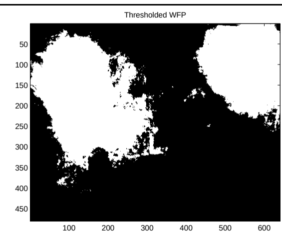	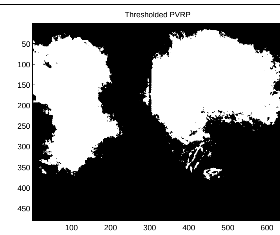	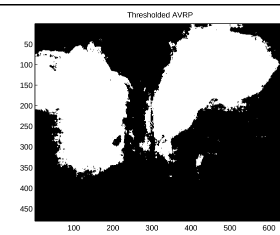
Ideal Match	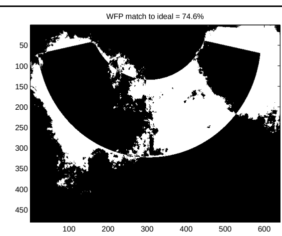	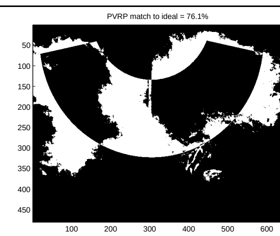	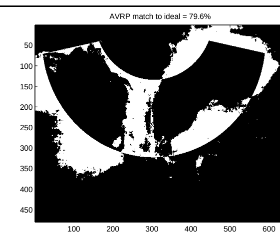
Percentage	74.6%	76.1%	79.6%

Figure 39: Match between each condition and an idealized group layout.

	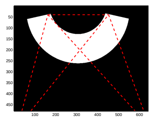	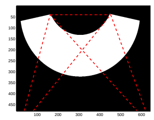	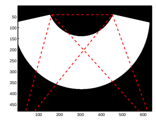	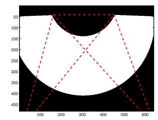		
	Alternative 1	Ideal	Alternative 3	Alternative 4	Average	δ
WFP	70.2%	74.6%	72.7%	69.0%	71.6%	
PVRP	71.7%	76.1%	74.6%	70.8%	73.3%	+1.7%
AVRP	75.8%	79.6%	77.1%	71.8%	76.1%	+2.8%

Figure 40: Matches with alternative ideal models with varying parameters are consistent. Alternative 2 was chosen as our ideal because it provided the closest match with the data.

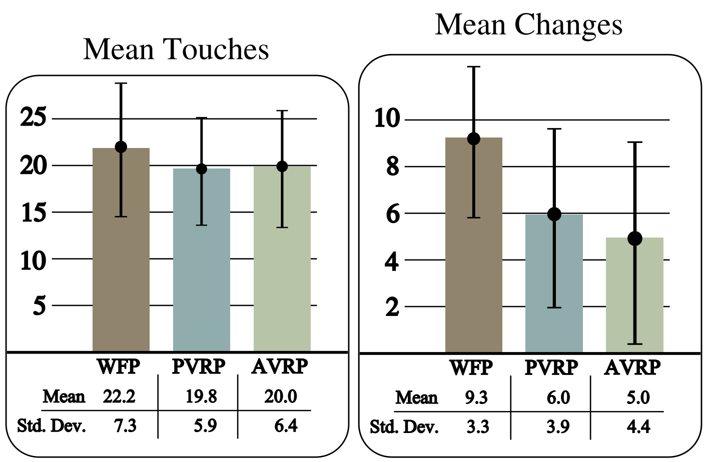

Figure 41: Mean Touches and Changes in the Aerospace Task

As shown in Figure 41, the mean number of touches for all three conditions was very close to 20, with a wide standard deviation. It appears that none of the conditions affected how many interactions it took to complete the tasks.

The number of times the person interacting with the board changed do show a trend. The WFP conditions have an average of 9.3 changes, while the dual-projector conditions have 5 and 6 changes per session. Because these measures are collected on a per-group basis, N=6, and the statistical power of the ANOVA is reduced. The results for the data on mean changes does not meet a strict $\alpha <= 0.05$ test (*F(2,6)=4.017, p <= 0.052,* $\eta^2 = 0.577$). The effect size $\eta^2 = 0.577$ indicates that with a larger N a statistically significance difference between WFP and the dual-projector conditions may be obtained.

However, note that if we naïvely accept that more changes between users is equivalent to better collaboration, this provides evidence against hypothesis four, and indicates that WFP promotes more collaboration than the dual projector conditions. In actuality, the reason for the elevated number of changes is that users were standing on either side of the screen to prevent their shadows from obscuring the image, and this limited their reach. Instead of moving across the screen, users would allow others located on the far side of the screen to interact with UI elements that were out of their reach. In one case, we observed a user begin moving a horizontal slider, only to “hand-off” the thumb to another user standing on the other side of the screen when it reached the half-way point.

In the dual-projector conditions (AVRP/PVRP), it was more common for a driver to emerge and stand directly in the center of the board. Because they stood in the center of the board, instead of on the side (as was typical for the WFP conditions) they were able to reach most of the board without having to move and other users were less likely to interact with the board until the driver stepped back into the group.

7.5.7 Perceived Value of AVRP

As part of the focus group interview, the groups were presented with a scenario where they were given a budget of $5,000 to outfit one of their company’s conference rooms with a SmartBoard system similar to the one they used in the study, and furniture. They were told that the SmartBoard system with two projectors (capable of using the WFP and PVRP modes) would cost $3,000, leaving them $2,000 to purchase furniture. They were then told than the SmartBoard salesman could upgrade their system (to allow it to use all 3 modes, including AVRP) for an additional $500 (leaving $1,500 to purchase slightly less expensive furniture). The $3,000 and $500 prices were chosen to be representative of the actual hardware costs. The participants were then asked if they felt that AVRP was worth the additional $500.

Ten (of 24) participants choose to pay an extra $500 to enable the option of choosing AVRP as a display mode. Four participants did not give a specific reason for this choice. Three participants felt that the AVRP mode would be most useful for presentations: “If you are going to be doing presentations to people, where you are facing the audience, absolutely.” (369) One participant felt that the additional “500 dollars on a pay scale kind of perspective is so tiny” (810) that it was a mistake not to make the investment, another felt that having the ability to switch to the AVRP option when needed was worth the $500, and another “just liked that particular version.” (372)

Six (of 24) participants were not sure if it would be worth $500. Some of them felt that AVRP would be worth the $500 if it were used by only a single user presenting to an audience, as it would might work better than it had with the 4-5 users in their experiment:

> "I don't know that I'd really say that without just having *(seeing)* one person using it cauz I don't know that it would be not, I mean it might not be flickering as much and stuff if you only had one person up there. *Participant 2:* I think it would really depend upon what kind of things you were trying to do, I mean if you were just doing this, then probably not. But if you were going to do something where like I said before, you were turning around and talking to the room a lot which you probably would be in a conference room, then maybe. It would depend, like she said, how it did with one person, how many people you expect to be up working on the board at once."

Eight (of 24) participants did not think AVRP was worth $500. A few mentioned that they felt that other modes (PVRP/WFP) were good enough for their purposes:

> "I don't think so. I think if a person were just to stand on one side, and only had one or two people up there, I think the one projector warped, just tell them stand on this side, and that would be the way to go"
>
> "I just don't think it's worth it, you still get a good image with the 2 projector simultaneously... I don't think it's $500 worth."

Four of the eight specifically mentioned that they did not want to spend $500 because of the visual artifacts, but would purchase it if the artifacts were imperceptible. As one participant said "I'd hold on to those $500 dollars and wait until their is a version 2.0 of the technology that currently doesn't have this artifact." (585)

7.6 Analysis & Results: Hangman Task

The Hangman task was chosen to represent a task where a single user interacts with an application on the board with a collaborating audience. Although only one user is directly driving the board, his or her actions are partially directed by, and influence the audience. The Hangman task gave our participants the opportunity to both work with the board while interacting with an audience, as well as be a member of the audience observing the board and the person directly interacting with it. Two main differences between the Hangman and Aerospace tasks are that only one person directly interacts with the board at a time (the other participants make up an audience) and the task does not require specialized domain knowledge outside of knowledge of English vocabulary. Because the audience is located behind the driver, they may notice graphics projected on the driver's body, and the driver may turn towards the audience (and the projectors) and be affected by incident light from the projectors. The audience members also have a wider view of the entire situation when compared to participants in the Aerospace task.

Our first three hypotheses were shared with the Aerospace study:

1. Users would prefer AVRP to PVRP because of (a) the reduction of blinding light and (b) lack of visible half-shadows.

2. Users would continue to prefer PVRP to WFP for the reason identified in our previous study: reduction of full shadows on the display.

3. Users would report more annoyance with projected light in the PVRP and WFP conditions when compared to the AVRP condition, because AVRP reduces blinding light.

The results for the Hangman study are similar to those in the Aerospace study. Users reported stronger opinions than in the Aerospace study about the three conditions, with the range between the highest and lowest ratings larger than in the Aerospace study. In most cases, the data trends mirrored those in the Aerospace task but with stronger significance. The main differing metric was the comfort (annoying light) question (Section 7.6.2.1).

7.6.1 User Preference

In the Hangman study, users preferred PVRP (with a mean rating of 6.3). AVRP (5.0) came in second, and WFP (3.3) was liked the least, $p <= 0.002$, $F(2,24)=21.55$, $\eta^2 = 1.000$. See Figure 42 for pairwise p-values.

The results from the Likert scale rating question was consistent with the results from the exclusive choice questions. When asked to identify the condition they liked the most, 17 preferred PVRP,

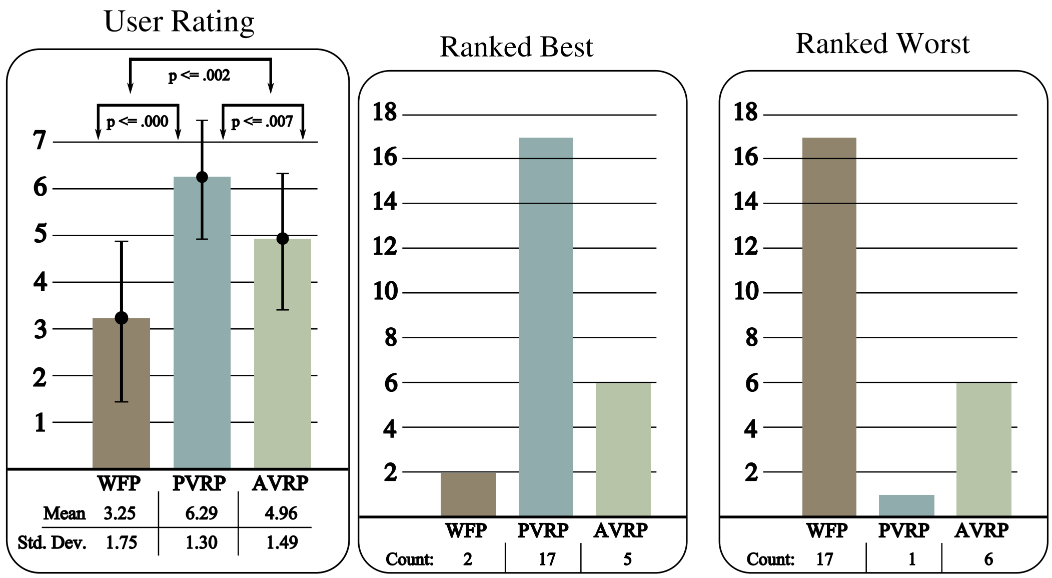

Figure 42: Rating question result for the Hangman study.

5 preferred AVRP, and 2 preferred WFP. When asked for the condition they disliked the most, 17 chose WFP, 6 choose AVRP, and 1 choose PVRP. A χ^2 analysis comparing these to a normal distribution (7,7,7) indicates that the distributions are significant *Ranked-Best:* $\chi^2 = 7.1$ *p <=0.05* *Ranked-Worst:* $\chi^2 = 8.6$ *p <= 0.025*.

Individual Questionnaires

On the individual questionnaire free response areas, users gave various reasons for liking and disliking the three projection conditions. A single researcher coded questionnaire results, forming more general categories when responses were similar. These categories supported by more than two participants are reported below:

Liked WFP:

No two users gave the same reason for liking WFP.

Disliked WFP:

- *Shadows* (fifteen participants) including:
  - *Shadows blocked view of board, or got in the way* (nine participants)

  - *Existence of shadows* (six participants)

- *Drawing or writing on the board was more difficult {for unspecified reasons}* (four participants)

- *Display was dim* (two participants)

Liked PVRP:

- *Less shadows* (six participants)

- *Good image quality*[8] (four participants)

- *Board was easy to see* (three participants)

- *Lighting* (two participants).

Disliked PVRP:

No two participants agreed on a reason they disliked PVRP.

Liked AVRP:

- *Lack of shadows* (three participants)

- *Novel visual effect* (two participants)

Disliked AVRP:

- *Visual artifacts* [9] (four participants)

- *Aesthetically unpleasant* (two participants)

- *Too dark* (two participants)

- *Intermittent failures to detect/correct occlusions* (two participants)

[8] The users reported that the display was “bright”, “clear”, “bright and clear” and that “the colors were shiny”.

[9] Described as (a) *a blob following your hand,* (b) *white shadows*, (c) *smudge effect*, and (d) *distracting half shadow*.

Minority Responses

Some users gave reasons for liking or disliking various conditions that were not echoed by other participants, and were not generalized into categories. These views represent a minority opinion that was not volunteered by other users in their individual questionnaire results or are likely a mistake (e.g. "No shadows to distract us" in WFP condition) due to misremembering the order of conditions. However, some of these minority opinions re-appear in the focus group interviews and would sometimes gain more support. The minority opinions held by single participants are reported here:

Reasons for liking WFP: *No shadows to distract us, Clear to see,* and *Didn't hurt eyes*. Reasons for disliking WFP include: *No Transparent,* and *Projector too bright.*

Reasons for disliking PVRP included: *Too many shadows,* and *Too bright, hurt eyes.* One participant liked PVRP because it was *Transparent*, possibly a reference to the half-shadows created by the redundant illumination.

Two people liked AVRP because it was *Easy to work with the board*, and *better than the other two options.* Two participants disliked AVRP because: *Other people's shadows affected the writing,* and *difficult to use because of errors with touching/selection*.

7.6.2 Annoyance of Blinding Light

7.6.2.1

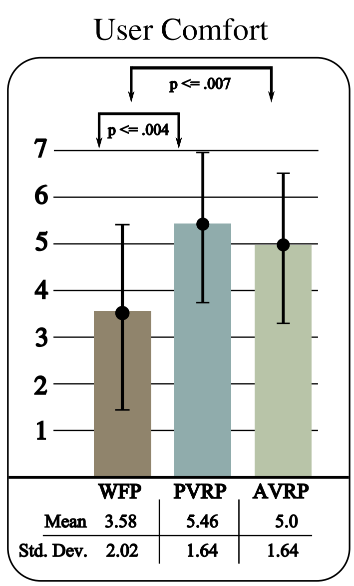

Figure 43: Comfort question result for the Hangman study.

As with the Aerospace study, we observed no behavior (e.g. squinting, raising a hand for shade) that would indicate users were having problems with the projected light. Unlike the Aerospace study, in the Hangman study users reported a statistically significant difference between WFP and the other two conditions on the comfort question $F(1.53,24)=8.41$, $p <= 0.002$. See Figure 43. The primary difference between the results in the Aerospace study and the Hangman study is the drop in the score of WFP, although the scores for PVRP and AVRP do rise slightly in the Hangman study. Looking at the individual scores, 13 (of 24) participants gave the WFP condition a score of 3 or lower (a score of 4 is neutral). PVRP and AVRP each had only four participants give them a score of 3 or lower.

Approximately a third of the participants reported that they had not noticed light coming from the projectors during the experiment, while another third said that any light they noticed hadn't bothered them. When participants were asked if they had had a problem with blinding light, most negative responses were very short, many times consisting of a single "No." The following are illustrative examples:

> "*Q: What did you think about this light from the projectors, did it ever bother you?*
>
> *1st participant:* No.
>
> *2nd participant:* No.
>
> *3rd participant:* I didn't notice.
>
> *4th participant:* No."[10]
>
> "I didn't notice it."
>
> "I didn't notice any of the light ever."
>
> "*1st participant:* I didn't really notice.
>
> *2nd participant:* Yeah, I wasn't really looking at *(the projectors)* all these letters were coming at me that I had to like..."
>
> "I never noticed it."
>
> "*1st participant:* I really didn't notice them.
>
> *2nd participant:*Yeah, I didn't."

[10]This quote represents near simultaneous answers from all four participants.

"I didn't have a problem, well, in all of them."

"Didn't really matter"

"*1st participant:* It didn't bother me at all...

*3rd participant:*Yeah, it's not, sometimes you have a projector on a table and the light shines right in your eyes, but here it's up top and I didn't, it didn't bother me at all.

4th participant: Sure, it wasn't eye level, it's fine."

Less than a third of the participants reported having even minor complaints about the light from the projectors shining in their eyes or faces, and no participant said that it was a serious problem. The following quotes are all of the issues participants expressed with blinding light in the interviews:

"I think that the first case *(AVRP)* I was actually a little surprised as how mild the lighting coming from the projector was when I first walked in here which was scenario number 1. And the 2nd one *(WFP)* was like everything else, because I've presented in other situations before, other situations where you have one projector and you have a screen and you are presenting, and the third one *(PVRP)* I didn't really notice that much disruption from the lighting either."

"I didn't notice it particularly for the first two *(PVRP, AVRP),* third one *(WFP)* I could feel where the light was coming from subconsciously, I didn't really look at it but I could out of the corner of my eye. Whatever it is it's coming from here."

"*Q: If you turned around you might have had light coming from the projectors hitting you in the face, did you ever notice that?*

1st Participant: Yeah, in the 3rd *(WFP)* one.

2nd participant: In the 3rd one, yeah."

"But I thought the second one, that I was standing up there, I thought the 2nd case scenario *(WFP)* was almost a little bit too bright for my taste. *Q: what do you mean by too bright?* Just the shining, the lights, coming from the projector, coming at the screen. *Q:Did anybody else feel that? (others shake heads negatively)"*

"*Q: Did any of you ever have problems with light shinning in your face when you*

> *were looking out at the other people?* It hit me once or twice, but...I work with projectors at OIT, so I'm kinda used to looking at projectors, so that's why, I think I just looked there."

The Hangman study was designed so that it was more likely that users would have problems with blinding light, as the driver would turn towards the projectors when making eye-contact with the audience. Still, two-thirds of the participants claimed to have either not noticed the projected light, or to not have been bothered by it. Even the comments (above) from participants who mentioned the projected light were not terribly negative.

If the user comfort question were an accurate representation of user discomfort caused by blinding light, we would have expected the score of PVRP (which projects twice as much light, from two different directions/projectors) to have been at least as low as WFP, but this is not the case. It is likely that users are considering other factors (such as shadows viewed on the screen) in addition to blinding light from the projectors when answering this question. It is also possible that users are more aware of a single projector than dial projectors due to the location of their shadow on the screen, as hypothesized by these two users:

> "In terms of direct light? I didn't notice it particularly for the first two (PVRP, AVRP), third one (WFP) I could feel where the light was coming from subconsciously, I didn't really look at it but I could out of the corner of my eye. Whatever it is it's coming from here. *2nd participant:* Exactly. Well yeah, it's because of the shadow again that's why, different because of the shadow, you know, it's coming, I mean I remember myself acting, to move like you know out of the way so that the other guys could see." *(This quote contains a sub-quote used earlier.)*

A result of the user comfort question and interview replies is that two-thirds of users did not report experiencing any discomfort caused by blinding light from the projectors on this task in the experimental setting. Of the one-third who mentioned that blinding light had been a problem, the majority were associated with the WFP condition, and even those comments were not extremely negative.

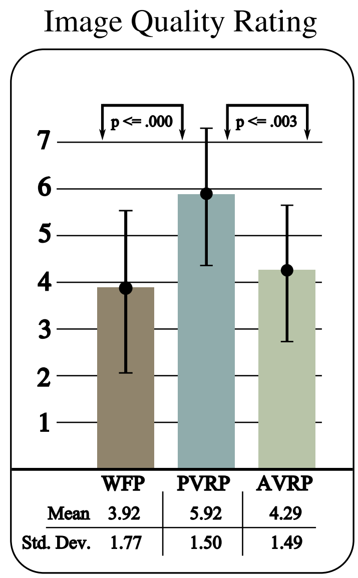

Figure 44: Image quality question result for the Hangman study.

7.6.3 Image Quality

Users felt that PVRP (mean score 5.92) delivered better quality images than WFP (3.92) and AVRP (4.29) $F(2,24)=9.57, p <= 0.001$. See Figure 44. The factors users mentioned the most when asked about image quality were the brightness and clarity of the display. Some users felt that the added brightness of PVRP (from the simultaneous operation of two projectors) gave it a better quality image:

> "I think the first one *(PVRP)* was the best, in terms of that, it was you know, shiny for me. The second *(AVRP)* and third one *(WFP)* was dimmer."
>
> "I think the first one *(PVRP)* was probably the best. *2nd participant:* The first one *(PVRP)* was best, it was the clearest and the brightest, I mean the third one *(AVRP)*, you don't want to have that, that white blob, I mean you could tell the outline of the blob and everything."
>
> "The first one *(AVRP)* was, blurrier than the 2nd *(WFP)* and third *(PVRP)*. To me the 3rd *(PVRP)* one was a lot brighter than the first one."

7.6.4 Mean User Activity

The work in this section was performed in conjunction with Mario Romero, who analyzed the data using adjacent frame differencing as part of his research, we collaborated on the interpretation and analysis of the processed data. The overhead time-lapse camera video was processed by detecting differences in adjacent frames over time to record aggregate motion. The graphs in Figure 45 show overhead activity maps displaying the average location of user motion averaged by condition. Each chart represents the location of user motion averaged over the six hangman sessions.

In the Hangman task, participants were split into two distinct groups, the audience (3 participants) and the driver (a single participant). During a session, the four participants would rotate through the driver role, who was in charge of keeping track of correct and incorrect guesses (from the audience), marking guessed letters on the board, and keeping score by drawing the hangman figure. In Figure 45, the driver is generally responsible for motion from zero to 150 pixels (on the vertical axis) and the audience is generally responsible for motion in the 250 to 480 range. After each word was completed, a member from the audience would pick up a card with a new word on it and move forward to the board (on the right), while the old driver would discard his card and move back into the audience (on the left).

Because the driver would approach the board on the right side, and the hangman drawing area was on the right, they would typically stand to the right of the board (so as to not block the audience's view), and this "home" position is represented in all three conditions with a large blob of activity in the top right.

In the WFP condition, the drivers spent most of their time in this "home" position and only made sorties across the board when absolutely necessary. The audience activity is noticeably shifted to the left, but we do not believe this is caused by them avoiding the projection beam from the single projector on the right. At the position they were asked to stand, five feet from the board (below 250 pixels on the vertical axis) it was very unlikely that a typical audience member would block the projector unless they were taller than 6'2". Any offset in the average audience location is most likely due to the shadow that would be cast to the drivers left. By moving to the left, we hypothesize that the audience members could see the board directly behind the driver, a view that would be blocked

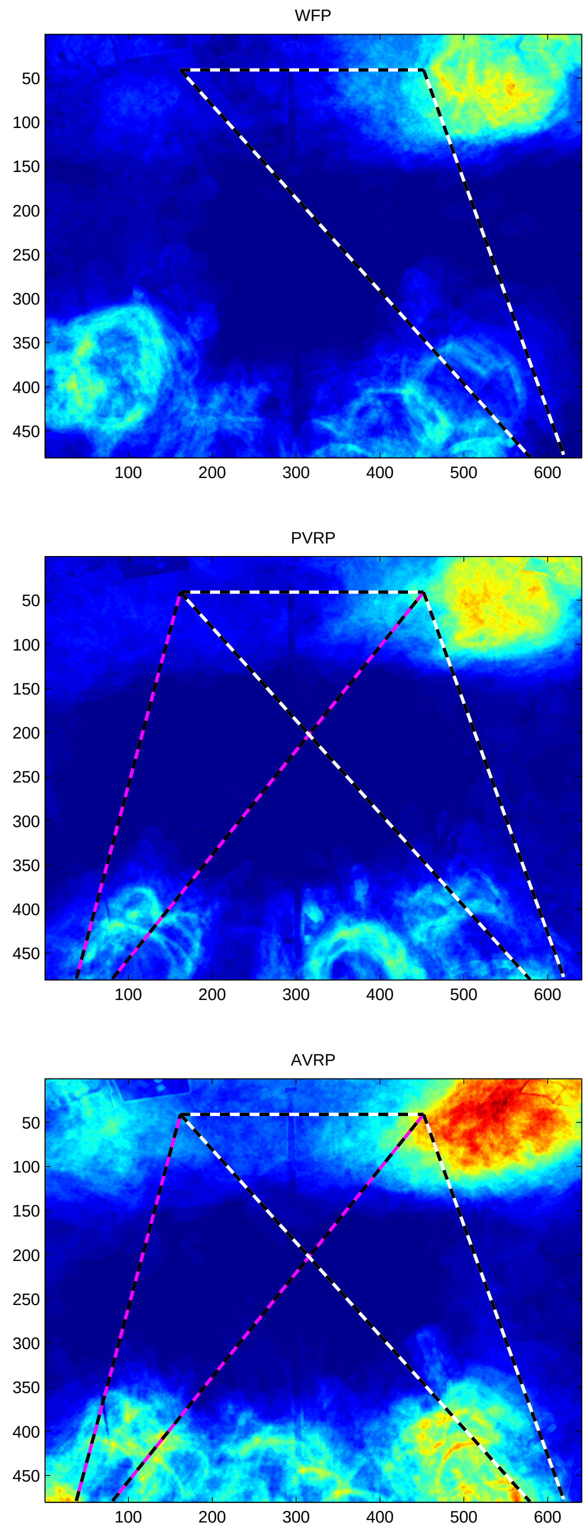

Figure 45: User motion by condition with overlaid projector beam paths in the Hangman study.

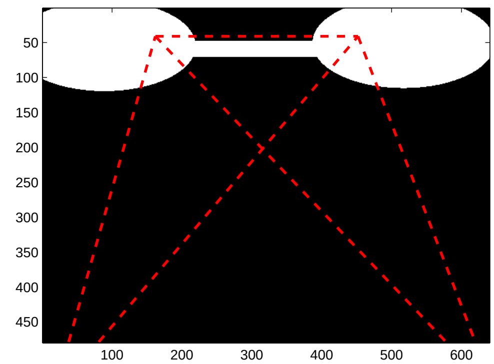

Figure 46: Hangman ideal model for driver activity.

if they stood centered on the board.

In the PVRP condition, the drivers still spent the majority of their time in the "home" position, but they moved in front of the board more than in the WFP condition, as evidenced by the more pronounced cloud extending to the left of the "home" position. The audience members were basically symmetrically about the center of the board, most likely caused by the symmetry of *penumbral* shadows cast by the dual projectors.

In the AVRP condition, both the driver and audience exhibited more motion in general. The driver spent more time on the left side of the board, as exhibited by the bright cloud on the left side, and obviously crossed in front of the board more frequently, as evidenced by the visible connection between the "home" position and the cloud on the left side of the board.

Ideally, the driver would be free to cross in front of the display, although they would spend the majority of their time on one side or the other, so they do not block the audience's view. Figure 46 shows a visual representation of this ideal model. The lobes on either side of the board are larger than the crossing region in the front, because the driver is more likely to turn towards (and approach) the audience to interact with them when not in front of the board. It should be made clear that we build this model only to get a numerical representation to quantify our observations. This "ideal" model is simply a representation of our subjective analysis, used for quantification of data, and not meant to be theoretically correct or ideal in a global sense.

As can be seen in Figure 47, the match percentage of WFP (89.9%) and PVRP (90.5%) are near equal, while AVRP (93.1%) is larger. This difference is primarily caused by more activity in front and to the left side of the board in the AVRP condition, indicating that the drivers crossed the board more frequently. To demonstrate that this calculation is stable with respect to the parameters specifying the model, we calculated the match with models of varying sizes (Figure 48) and demonstrated that while the absolute percentages may change slightly, the relative ordering of the conditions remain constant, with WFP and PVRP being near equal, and AVRP having a higher value. For our "ideal" model, we chose the alternative (number 2) that gave the largest overall match.

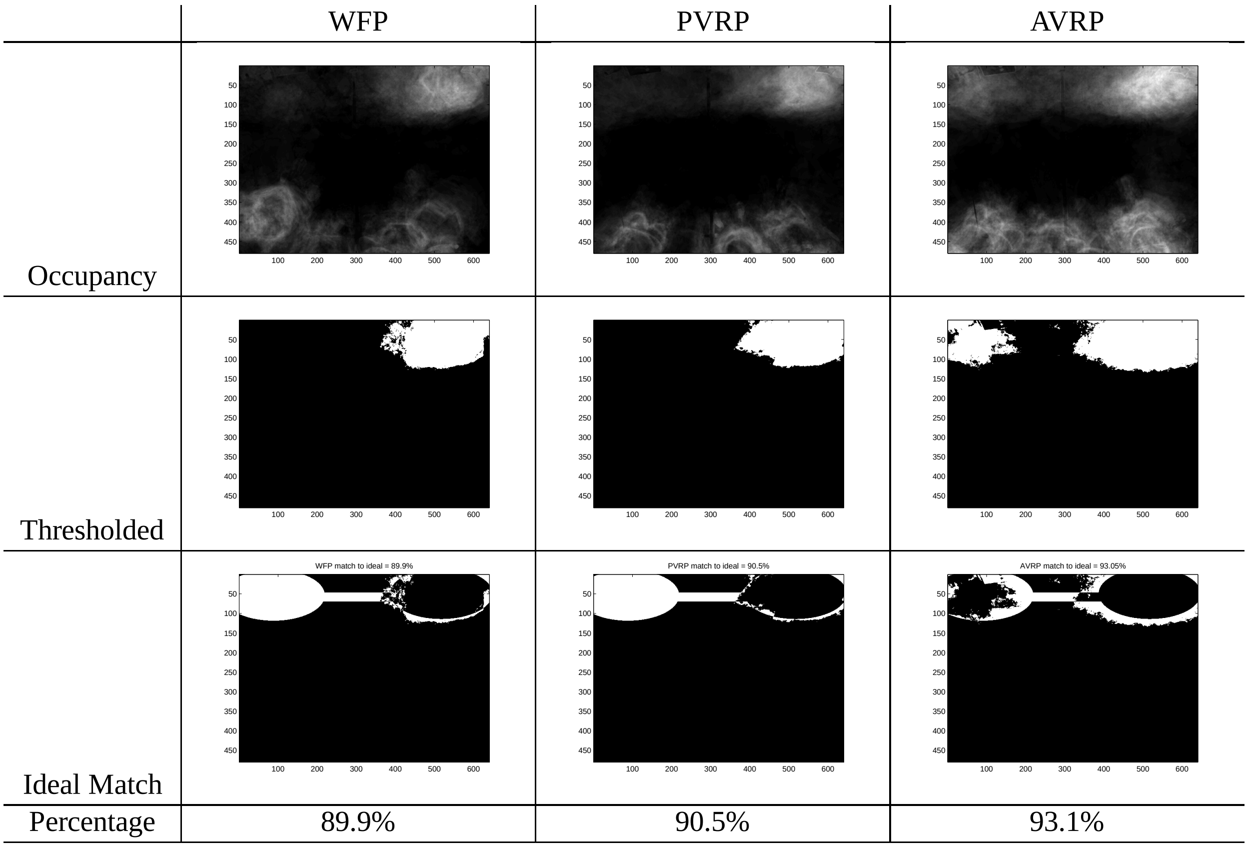

Figure 47: Match between the driver's activity in each condition and an ideal model in the Hangman study.

7.6.5 Perceived Value of AVRP

As in the Aerospace study, participants in the Hangman study were given a scenario (see Section 7.5.7 for details) and asked if they would pay $500 extra for the ability to use the AVRP projection condition. After participating in the Hangman task, only three (of the 24) participants would pay $500 extra for AVRP. One participant felt that "$500 isn't that much more, when you're dealing with already $3000," and another participant felt that "the distortion didn't really bother me, and if there

	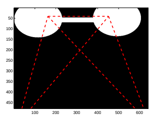	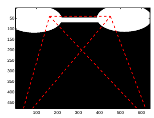	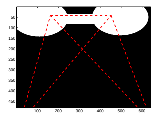	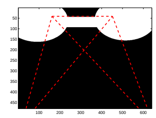		
	Alternative 1	Ideal	Alternative 3	Alternative 4	Average	δ
WFP	87.8%	89.9%	86.2%	83.2%	86.8%	
PVRP	87.2%	90.5%	85.6%	83.2%	86.6%	-0.1%
AVRP	91.3%	93.1%	91.5%	89.8%	91.4%	+4.8%

Figure 48: Hangman matches with alternative ideal models with varying parameters are consistent. Alternative 2 was chosen as our ideal because it provided the closest match with the data.

was a fair amount of interaction involved then I think a person going back and forth throughout the screen would cast that double shadow *(from PVRP)* enough that it would be a hindrance, so I would pick the, I would invest the $500, and office furniture is office furniture."

Three (of the 24) participants were undecided. Two of these participants felt that the technology might be right for a high-tech company, but wasn't yet ready for a regular company. "If I was doing this thing at a tech company like Google or Amazon I would go for the all out because my boss would think that was really cool. But if I was doing it with a normal company I don't think the technology is mature enough to use on a professional level."

The remaining eighteen (of 24) participants would not pay $500 for the ability to use AVRP. Six participants declined for unspecified reasons. Six of the participants disliked the visual artifacts, and some felt that they would be distracting or unprofessional:

> "The two projector switching with that blob, I'm so unaccustomed to that blob, that at least right now I would prefer to have the one projector with the hard shadow. Because I'm accustomed to dealing with a shadow. No, it's not worth it to me."
>
> "Because if someone is giving even just like a normal presentation, I would be too distracted looking at the distortion then focus on what they are saying."
>
> "I actually think that for professional use that third one *(AVRP)* looks kinda rough. I wouldn't really use it because you can see a big blotch on the screen, and it kinda looks, when you are trying to do a presentation whatever, it's distracting and looks unprofessional.
>
> *2nd participant:* I agree, I think it needs some polish before it's a viable solution."

Four participants felt that PVRP was good enough:

> "I think the option number two *(PVRP)* was superior anyways and I don't know if I would ever want to switch from that.
>
> *2nd participant:* Yeah, I don't see why you wouldn't ever use option number two *(PVRP)*."
>
> "I like the third one *(PVRP)* so it wouldn't matter to me if I had the first *(AVRP)* one, so I would say no."
>
> "I liked the second one *(PVRP)* more so I wouldn't, I'd get the... *(furniture*)"

Two other participants just didn't "think it's worth the money" or didn't think that "the technology is quite there yet."

After giving their answers, the participants who didn't want to spend $500 were then asked to negotiate with the SmartBoard salesman and tell how much they felt adding AVRP to their system was worth. Three participants felt it wasn't worth anything, while one participant would offer $25 "just to have it." The remaining participants made offers ranging from $50 to $250, with the average near $150.

7.7 Study Similarities and Contrasts

The Aerospace and Hangman studies differed mostly by the task performed but also by the participant background and demographics. The different task directly affected how the groups interacted with the board (singularly or in groups) and the type of interactions with the board (GUI element manipulations vs. inking strokes). The aerospace graduate students were generally older than the primarily undergraduate participants in the Hangman study. The aerospace student groups were recruited from existing class and lab groups, and generally had experience working with each other on similar problems before the study. Participants in the Hangman study were recruited individually, and only rarely did two or more people in the groups know one another. When comparing data gathered across the two studies this task and participant differences did result in some differences in the dependent variables. For most of the metrics (especially those sampled by the individual questionnaire) we believe that the majority of differences are due to the change in task, and not to the change in participant demographic or prior friendship status.

User Rating & Image Quality

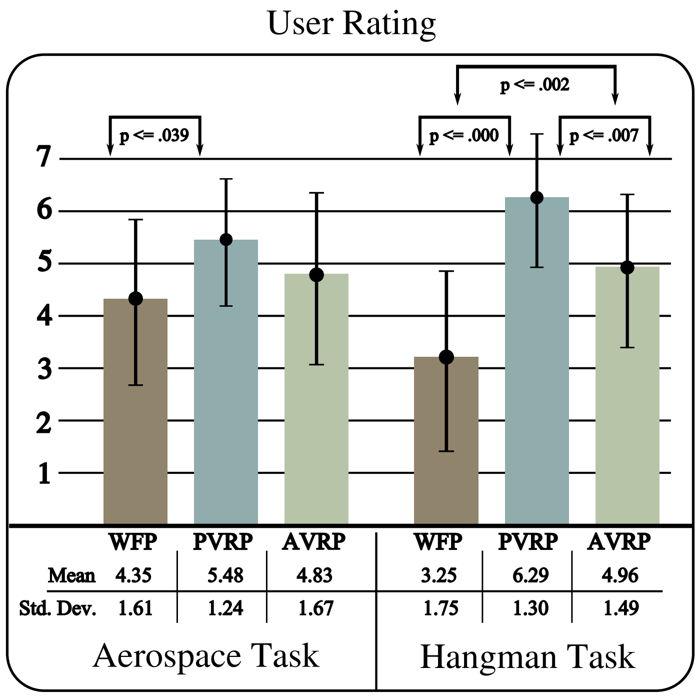

Figure 49: User rating differences between studies.

Participants in the Hangman study rated PVRP higher and WFP lower than those in the Aerospace study (Figure 49). Although the data trends are the same, all three differences in the Hangman study are statistically significant.

While no difference in image quality was detected in the Aerospace study, participants in the Hangman study rated PVRP higher and WFP/AVRP lower than those in the Aerospace task, leading to a statistically significant difference. These statistically significant differences trend in the same direction as the data in the Aerospace task (Figure 50). We attribute the larger rating differences to users having more time to passively observe the display while in the "audience" role in the Hangman task. An alternative explanation is that the graduate student population in the Aerospace study were more conservative and less likely to report as wide a difference in opinion on the Likert scale questions. Overall, the user rating and image quality questions show consistency in their trends across studies and agree with user sentiment expressed in the focus group interviews.

User Comfort

Participants in the Hangman study rated the comfort level of WFP lower than in the Aerospace

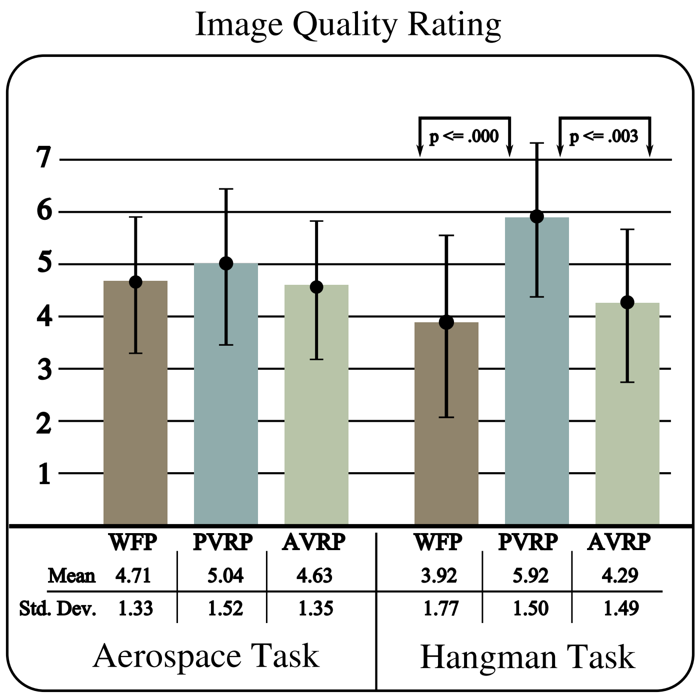

Figure 50: Image quality rating differences between studies.

study, resulting in a statistically significant difference between WFP and AVRP/PVRP in the Hangman study (Figure 51). Due to the few participants who reported being affected by blinding light from the projectors in both tasks, the results of the user comfort question are called into question. We believe that the majority of users in both studies answered this question based upon factors other than blinding light. The focus group interviews indicated that even in the Hangman study (where the driver was more likely to turn towards the projectors when interacting with the audience) the majority of users didn't notice, or were not bothered by the light from the projectors.

Perceived Value of AVRP

One of the largest contrasts between the two studies is the number of participants who were willing to spend $500 to have the option of using AVRP in a hypothetical scenario. More than twice as many people were willing to pay $500, or were undecided, in the Aerospace study than were in the Hangman study (Figure 52). A χ^2 analysis comparing these distributions indicates that the differences are significant $\chi^2 = 8.6$, *p <=0.013*. This is most likely caused by differences in the tasks. The difference in user preference between PVRP and the other conditions was stronger the Hangman study than in the Aerospace study. Many of the Hangman participants who were not

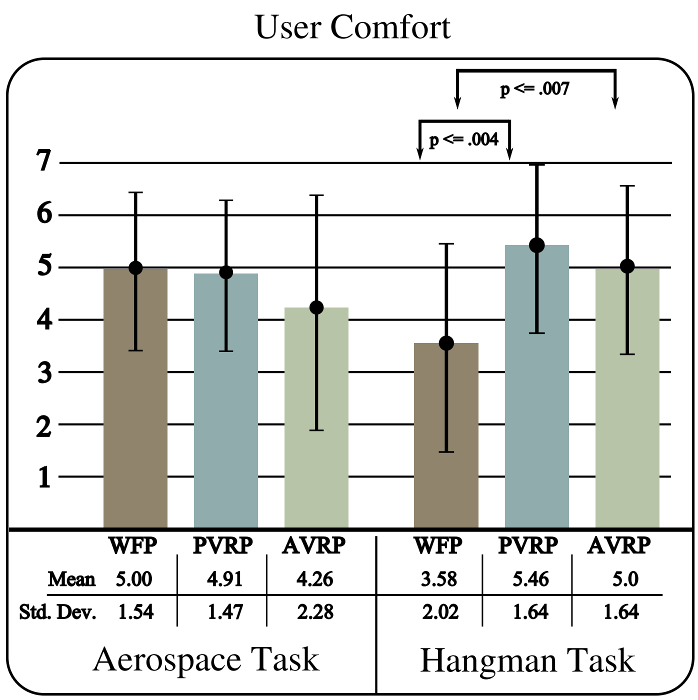

Figure 51: User comfort rating differences between studies.

willing to pay an extra $500 for the AVRP condition justified their decision by stating that PVRP was good enough. It is likely that these participants, having never experienced PVRP suffering from multiple occluders, and not having suffered ill effects from blinding light, felt that they didn't need the option of using AVRP.

Another explanation that we can not rule out is that the Aerospace Study participants, as Aerospace Engineering grad students, were predisposed to value the AVRP technology more than the Hangman students, who represent a broader (and younger) demographic.

7.8 Reflections on Research Methodology

As with any endeavor, ways to improve the studies become clear in hindsight. In the following sections we outline problems with equipment and procedures, as well as methodological changes that would have made comparison between the two studies more practical. Overall we are very pleased with how smoothly the studies ran, and although the study tasks could have been chosen so that the two studies would have worked better together to examine the role group configuration played on the dependent variables, the individual studies served the purpose for which they were originally designed.

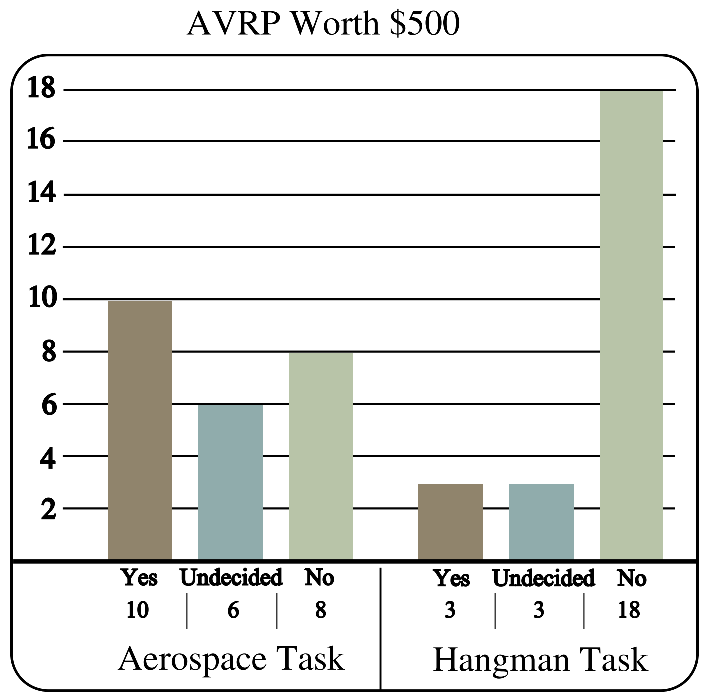

Figure 52: Is AVRP worth $500 extra?

7.8.1 Equipment & Researcher Reliability

When performing user evaluations using research prototypes a constant danger is the failure of the prototype. We are happy to report that the WinVRP application (used to implement all three projection conditions) performed flawlessly, and the computer, projectors, and task application software suffered no failures during any user study. We attribute this reliability to the previous deployment of the WinPVRP application to the School of Aerospace Engineering. Because WinVRP is built upon the WinPVRP code-base (with the addition of the AVRP algorithm) it benefited from the extensive testing, user feedback and iterative improvements that went into the WinPVRP application. During an early pilot one projector bulb (lamp) imploded, which prompted us to keep a spare lamp on-hand during the actual studies, but thankfully it was not needed.

During one focus group interview the battery in the digital voice recorder that was the primary source of audio for transcript generation was depleted, causing the audio recorder to fail to capture the last several minutes of discussion. Luckily, the audio track from the video recorder that was also used to document the focus group discussion served as a backup, and the complete focus group interview was transcribed from the two recordings. Our procedure was subsequently amended to include a battery voltage check of the wireless microphone (for capturing audio during the task sessions) and digital voice recorder (for focus group audio) batteries before each study. We recommend the use of video recording as a backup for audio recording of focus group interviews. In addition to covering for equipment failure, video was useful in several instances to review the video to see gestures made by participants–for example, to disambiguate statements when participants used expressions such as "well, like she said said *(pointing at participant 2)*" .

The only case of preventable data loss that occurred during the study was when the researcher forgot to turn on the wireless microphone that was hung above the SmartBoard and fed into the video recorder for the first condition in one Aerospace study. Fortunately, we were not planning on analyzing the audio data from the experimental sessions, and the video alone was sufficient for the analysis of changes in user touches in Section 7.5.6.

7.8.2 Reflections on Task Selection

The two studies (Aerospace & Hangman) were designed to investigate relative differences in our three projection conditions under two different usage patterns of an upright interactive display. The Aerospace study was designed to investigate collaborative use of the board by a problem solving team. The Aerospace task was chosen because it was ecologically valid. Our collaborators in the School of Aerospace Engineering felt that it was exactly the type of application their students and graduates would use in future work environments on a large interactive display.

We also wanted to investigate a driver/audience configuration where a single user was driving the display while observed by a set-back audience who would participate only vocally. We did not use the same Aerospace task because the participant pool (Aerospace engineering students who had taken the appropriate design class) was limited and could not support both studies. Instead, we looked for a task that was similar to the type of collaborative activity we wished to study, but easy enough so that a general college student population could perform it well with minimal training. We initially considered the game of Pictionary™where each driver would draw a secret word and the audience would attempt to guess the word based upon drawings made by the driver, with a one minute time limit. But after piloting the task twice we found that participants were *too* engaged in the task. In pilot tests over half of our participants suffered from task blindness to such an extent that they were not aware that we had three different projection conditions. Hangman is similar to Pictionary™ but because it did not have an explicit time limit and the audience members usually took turns choosing letters, the audience members were not totally immersed in the task.

Each of these tasks worked well for their respective studies, but the difference in participant pools meant that we could not easily compare the two studies to look for differences based solely upon the task with a large degree of confidence. Additionally, even if the participant pools had been identical, the individual tasks are somewhat different. In the Aerospace task participants choose design alternatives by selecting check-boxes, while in the Hangman task the audience choose letters and the driver wrote them on a display and kept score using digital ink. These differences in the task would have complicated matters if we wished to attribute differences found between studies to only the audience *vs.* collaborative group aspect, and not the differences in task mechanics.

To be able to make direct comparisons between the studies, we should have replaced the Aerospace task with a Group Hangman task, the mechanics of which would be as close as possible to the driver/audience Hangman game. For example, the users could have crossed out letters using digital ink strokes, and the computer would take over the role of the driver, by keeping score and positioning correct letter guesses on the letter blanks. In this way, we could have used the same participant pool for both studies, and an almost identical task. This would allow us to attribute any differences in the study to only the collaborative group *vs.* driver/audience configuration. However, this would destroy the ecological validity which is a strong point of the Aerospace task. As the original goal of the two studies was primarily to evaluate the three different projection conditions relative to each other in two common usage spaces, the differing tasks and user populations was not a critical defect.

7.9 Conclusions

Section 7.1 outlined the overall research questions that motivated these studies, while Sections 7.5 & 7.6 outline specific research hypotheses that we initially attempted to examine. Some of these hypothesis were shared between the Aerospace (AS) and Hangman (HM) studies. For example, AS-HM-H1 refers to hypothesis 1 that was shared between the Aerospace and Hangman study. We repeat them in this section labeled with the study they apply to (AS or HM) and their hypothesis number before discussing them. Additional commentary on the larger research questions is spread throughout. After discussing these research questions and hypothesis in the following sections, we make several claims that we feel these studies support.

7.9.1 User Preference

One of the primary purposes of this study was to determine how users felt about PVRP and AVRP for interactive tasks, especially when compared to a more traditional front projection (single projector) option. Additionally, we wanted to compare AVRP with PVRP and WFP because this would be the first user study of a system that actively compensated for shadows and eliminated blinding light.

In the Aerospace (collaborative group) task, questionnaire data and user focus group interviews clearly show that the users prefered PVRP to the single projector condition (WFP). As detailed in Section 7.5.2, no statistically significant differences were detected between PVRP and AVRP in the Aerospace task, although many people gave reasons they disliked AVRP, primarily due to the visual artifacts, but also including intermittent failures in shadow elimination, difficulty forming a mental model about how the system worked, and an inability to maintain an image in heavily occluded conditions. WFP was disliked primarily because users cast large shadows that interfered with their use of the board. No consistent reason to dislike PVRP emerged. Two users liked WFP because they felt it was not as blurry as the dual-projector conditions. Ten users liked PVRP due to its ability to reduce shadows, and two liked it because it had less visual artifacts than AVRP. Four participants liked AVRP because it reduced shadows.

In the Hangman (driver/audience) task, questionnaire data and user focus group interviews clearly show that the users ranked the conditions from worst to best in the following order: WFP, AVRP, PVRP. (This statistically significant ranking agreed with the trends seen in the Aerospace task data.) Again, WFP was disliked primarily due to shadowing on the screen. As detailed in Section 7.6.1 the primary reasons given for disliking AVRP was visual artifacts, dimness, “ugliness”, and intermittent failures in shadow elimination. In both the Aerospace and Hangman studies, no consistent reason to dislike PVRP emerged. In the Aerospace study no consistent reason to like WFP emerged. Three participants liked AVRP due to a lack of visual shadows, while two users simply liked the novel visual effects the active compensation provided. Users liked PVRP due to reduced shadows (six users), good image quality (four users), the fact that the board was easy to see (three users), and lighting on the board (two users).

We now discuss our two hypotheses that deal with user preferences:

AS-HM-H1*: Users would prefer AVRP to PVRP because of (a) the reduction of blinding light and (b) lack of visible half-shadows when standing between the projector(s) and the board.* Hypothesis AS-HM-H1 is disproven. No difference was found in the Aerospace study, and in the Hangman study users actually prefered PVRP to AVRP. As detailed in Sections 7.5.3 & 7.6.2.1, users as a whole were not annoyed by blinding light in either of our studies, and did not notice the reduction in blinding light provided by AVRP. As a group, users did have a problem with the half-shadows produced by PVRP.

AS-HM-H2: *Users would continue to prefer PVRP to WFP for the reason identified in our previous study (Section 4.3.1): reduction of full shadows on the display.* As detailed in Sections 7.5.2 & 7.6.1, hypothesis AS-HM-H2 is supported. PVRP was prefered to WFP in both studies, and the reasons given for disliking WFP primarily included shadows on the screen. Reasons for liking PVRP and AVRP included the reduction of shadows on the screen.

7.9.2 Benefits of Redundant Illumination & Blinding Light Suppression

The projection conditions which offered redundant illumination (PVRP & AVRP) were generally prefered to the single projector (WFP) condition. In the WFP condition, analysis of user motion showed that users were avoiding the projection beam path. In the PVRP and AVRP conditions, motion was much more noticeable inside of the projector beam paths, indicating that users moved through the space with fewer restrictions when redundant illumination was present. Many users noticed and commented on the visual artifacts produced by AVRP although only a few users said that they were extremely annoying. Overall, users did not have problems with blinding light coming from the projectors in our setup. Most users claimed to have never noticed light coming from the projectors, and the majority of those that did said that it had not bothered them. Because blinding light was not a concern for our users, we are unable to determine if the elimination of blinding light (the primary feature difference between PVRP and AVRP) is subjectively beneficial to users based upon their self reported data. The only objective measure to find a positive difference between AVRP and PVRP is the analysis of the overhead camera video that collected aggregate motion data.

The Aerospace data implied that there is a difference in user behavior (exhibited by their motion through space) between PVRP and AVRP that is as large as the difference exhibited between WFP and PVRP.

We now discuss our stated hypotheses that deal with the benefits of redundant illumination and blinding light suppression:

AS-HM-H2: *Users would continue to prefer PVRP to WFP for the reason identified in our previous study (Section 4.3.1): reduction of full shadows on the display.* As stated in the previous section, hypothesis AS-HM-H2 is supported. PVRP was prefered to WFP in both studies, and the reasons given for disliking WFP primarily included shadows on the screen. Additionally, AVRP was prefered to WFP in the Hangman study (Sections 7.5.2 & 7.6.1).

AS-H4: *When using the PVRP and AVRP conditions, users would gather closer to the screen than when using WFP (due to the shadow elimination).* Hypothesis AS-H4 is not supported. Although the location of user motion differed between the three conditions (Section 7.5.5), the absolute distance from the board was not significantly affected (less than 3 inches) by the projection technology (Figure 53). However, in the Aerospace task, user motion data that was compared to a model of idealized user layout for a collaborative task showed a positive difference between AVRP and PVRP that was just as large as the difference between PVRP and WFP. This may indicate that AVRP had increased benefits above PVRP that users were not able to articulate in the focus group interviews or questionnaires.

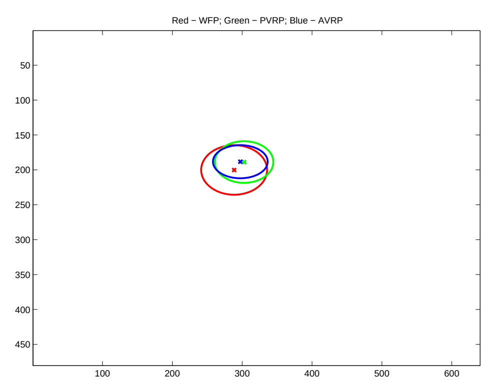

Figure 53: Location of group centroids in Aerospace study.

AS-H5: *The tendency to gather closer to the display in the dual projector conditions (AVRP/PVRP) would increase collaboration.* Hypothesis AS-H5 was not supported. We looked at the number of times the primary driver of the board changed in the Aerospace task (Section 7.5.6) but concluded that this was not a good measure of collaboration. It is likely that WFP increases the number of times the primary driver of the board changes, but this is only due to the inconvenience caused by crossing the board without redundant illumination.

AS-HM-H3*: Users would report more annoyance with projected light in the PVRP and WFP conditions when compared to the AVRP condition, because AVRP reduces blinding light.* Hypothesis AS-HM-H3 was not supported. No statistically significant difference on the Annoying Light (Comfort) questionnaire was detected in the Aerospace task. This is likely due to the majority of users in both studies not being aware of any ill effects from blinding light.

7.9.3 Claims

We present the following high level claims as a result of the studies reported in this chapter and in Chapter 4:

1. *Redundant illumination improves the user experience when compared to single projector conditions due to reduced shadows.* The studies in Chapter 4 demonstrated that users have a strong preference for warped front projection when compared to traditional front projection and that WFP and PVRP provide performance gains over traditional Front Projected displays for simple tasks due to a reduction in shadows. The studies in this chapter show that users prefer PVRP and AVRP to WFP due to the redundant illumination and shadow reduction properties.

2. *In a well constructed front projection environment using warped front projectors (singularly or in redundant pairs) with normal office illumination levels, users do not consciously suffer ill effects from projected light, and blinding light elimination may be unnecessary.* Users in the Aerospace and Hangman studies did not report annoyances caused by blinding light projected from our (off-axis) front projectors, and did not feel that AVRP provided strong advantages over PVRP due to its ability to block blinding light. However, differences in user

motion between the AVRP and PVRP conditions indicate a measurable effect on users leading to a difference in behavior that is not fully understood.

Chapter VIII

FUTURE DIRECTIONS & CONCLUSIONS

In this chapter we will discuss future research opportunities to further improve understanding of the effects that front projected displays have on users and we will then summarize the contributions of this body of work in Section 8.2

8.1 Future Directions

Going beyond our ceiling mounted projectors, other configurations should be investigated, such as an ad-hoc layout of projectors at table height. Different configurations may cause blinding light to have a more detrimental effects than in our studies with ceiling mounted projectors. Future research is needed to more fully investigate the effects that blinding light has on user behavior, preference, and performance.

Systems such as AVRP, which attempt to eliminate blinding light, currently produce visual artifacts on the screen. Although users prefered these visible artifacts to the full shadows of a single projector display, more research is needed to determine how effective such systems are at eliminating the effects of blinding light, and to determine if the visible artifacts they produce are causing other unwanted side-effects. Ideally, these artifacts can be eliminate entirely through improved photometric calibration and edge blending.

As this work was focused on constructing an output (display) system for large scale interactive surfaces, we used off-the-shelf input technologies (Liveboard, SmartBoard) that themselves suffer from cost and portability issues. Just as this work has developed technology to build easily portable and reconfigurable displays, future work needs to address the input problem, developing inexpensive and easy to deploy methods for detecting user input over large displays.

Looking further into the future, rollable wallpaper displays that incorporate touch sensing technology may allow for the easy and inexpensive deployment of large scale wall sized displays. But regardless of where these displays are deployed, users will wander elsewhere carrying only their equivalents of laptops and cell phones. These future portable computing devices are likely to have miniature laser projectors and the ability to work together to build displays that are larger, brighter, and more robust in the face of occlusions and shadows.

8.2 Conclusions

The overall motivation for this work was to enable the deployment of large scale interactive displays into everyday life. Although direct imaging displays such as plasma display panels and LCD displays have grown in size and become more affordable as manufacturing technologies improve, projection is still the most affordable way to build large displays. Projectors still hold cost, size, and portability advantages over direct imaging displays, and current trends seem to indicate that these advantages will remain constant over the next decade. While examining front projected displays, we identified two major problems (occlusion leading to shadows, and blinding light striking the user) that detracted from their usability. Our early user evaluation work showed that users disliked shadows and blinding light. We also observed performance decreases in interactive tasks due to shadows (Chapter 4).

We reduce shadowing on the screen by a combination of off-axis (warped) front projection and the use of redundant illumination achieved by calibrating multiple redundant front projectors using computer vision to produce a virtual rear projection display (Chapter 3). We use computer vision to detect when users are blocking a projector, and dynamically prevent light from striking users while correcting the resulting shadow using redundant projectors to maintain a stable image on the display. In addition to the technical development, we present a comparison of our AVRP implementation to previous systems that mitigated shadows and blinding light (Chapter 5). By implementing our algorithms on commodity hardware graphics accelerators we are able to achieve interactive frame-rates (75 Hz or faster) so that we can legitimately evaluate the technology in user studies.

We have made the source code of our implementation available to developers and other researchers as part of the PROCAMS toolkit. The PROCAMS toolkit includes abstractions that allow developers to build virtual rear projection displays without needing to understand the underlying computer vision, 3D graphics hardware acceleration, or geometric calibration problems. The PROCAMS toolkit ships with several demonstration applications that are useful for understanding how the toolkit should be used. Also included in the PROCAMS toolkit is the WinVRP software, which is the software used in our user studies. This software will allow other researchers with the appropriate hardware (a SmartBoard, high speed USB camera, infrared filter and lights, and two projectors) to replicate our user studies. In addition, one example application in the PROCAMS toolkit, WinPVRP, is also distributed separately as a stand-alone application and is suitable for end users. WinPVRP is designed to let an end user with the appropriate hardware (minimum of one projector, optional second projector and camera) construct a WFP or PVRP display "out-of-the-box" with no programming effort (Chapter 6).

In our user studies we found that redundant front projectors significantly improved the user experience over traditional front projected displays (Chapter 7). In our controlled laboratory studies we operated under normal office lighting levels and the projectors were ceiling mounted. In this configuration we found that the largest gain was due to the elimination of shadows on the display. Users did not report having problems with blinding light, but they still showed differences in behavior between the PVRP and AVRP conditions.

In summary, *by using a projector-camera system to mitigate shadows, a virtual rear projected display improves upon the user experience with respect to a traditional front projected display.* This confirms half of the thesis statement, while disproving the blinding light clause. We made the following contributions with this work:

1. Technology development to support passive and active front projection technologies for interactive surfaces (Chapters 3 & 5).

2. A software toolkit (PROCAMS) and example applications enabling others to experiment with virtual rear projection technology and replicate our work without having to re-create our implementation (Chapter 6).

3. User evaluations of passive and active front projection technologies for interactive surfaces in controlled laboratory experiments (Chapters 4 & 7).

Appendix A - Survey Instruments & Brainstorming Illustrations

Demographic Questionnaire

4th February 2003

Participant #: ____

1. __Male   __Female
2. ____Age
3. __Right-Handed   __Left-Handed
4. If you need glasses or contacts, are you wearing them now?
   __My vision needs no correction.
   __I am wearing corrective eye-ware. (__Glasses   __Contacts)
   __I am not wearing my corrective eye-ware. (Vision: ____/ 20)
5. Please rank your experience in using the following:
   1 = No Experience, 4 = Moderate Experience, 7 = Daily Use

   Chalkboards and/or Whiteboards:
   1 2 3 4 5 6 7

   Computer Video Projectors (Presentations, etc)
   1 2 3 4 5 6 7

   LiveBoards or SmartBoards:
   1 2 3 4 5 6 7

   Mimio or eBeam electronic pens:
   1 2 3 4 5 6 7
6. Do you currently have a headache, arm injuries, or any other illness or injuries that may make it difficult to finish the study?
   __NO   __YES

1

Figure 54: Demographic Questionnaire **-** Preliminary User Study (Chapter 4).

Task Questionnaire

6th February 2003

Participant #: ____

Condition 1:

How would you rate the image quality of the display technology?
Poor Quality = 1 2 3 4 5 6 7 = Excellent Quality

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

Please rate your willingness to use this display technology on the following scale:
Absolutely unacceptable = 1 2 3 4 5 6 7 = Completely acceptable

Condition 2:

How would you rate the image quality of the display technology?
Poor Quality = 1 2 3 4 5 6 7 = Excellent Quality

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

Please rate your willingness to use this display technology on the following scale:
Absolutely unacceptable = 1 2 3 4 5 6 7 = Completely acceptable

Condition 3:

How would you rate the image quality of the display technology?
Poor Quality = 1 2 3 4 5 6 7 = Excellent Quality

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

Please rate your willingness to use this display technology on the following scale:
Absolutely unacceptable = 1 2 3 4 5 6 7 = Completely acceptable

1

Figure 55: Between Condition Questionnaire - Preliminary User Study (Chapter 4).

Post-Task Questionnaire

6th February 2003

Participant #: ____

Now that you have seen all four conditions, please answer the following questions:

1. How would you rate the image quality of the display technologies?
   1 = Poor Quality
   4 = Neutral
   7 = Excellent Quality

   Condition 1: 1 2 3 4 5 6 7
   Condition 2: 1 2 3 4 5 6 7
   Condition 3: 1 2 3 4 5 6 7
   Condition 4: 1 2 3 4 5 6 7

2. Please rate the display technology on the following scale for the tasks performed:
   1 = Definite Dislike
   4 = Neutral
   7 = Liked very much

   Condition 1: 1 2 3 4 5 6 7
   Condition 2: 1 2 3 4 5 6 7
   Condition 3: 1 2 3 4 5 6 7
   Condition 4: 1 2 3 4 5 6 7

3. Please rate your willingness to use this display technology on the following scale:
   1 = Absolutely unacceptable
   4 = Neutral
   7 = Completely acceptable

   Condition 1: 1 2 3 4 5 6 7
   Condition 2: 1 2 3 4 5 6 7
   Condition 3: 1 2 3 4 5 6 7
   Condition 4: 1 2 3 4 5 6 7

1

Figure 56: Post Study Questionnaire - Preliminary User Study (Chapter 4).

Task Questionnaire

Participant ID: ____

Demographic Information:
Age: _____ Gender (circle): M F

Task Questions:

How would you rate the image quality of the projected display in...

Part 1:	Poor Quality = 1	2	3	4	5	6	7 = Excellent Quality
Part 2:	Poor Quality = 1	2	3	4	5	6	7 = Excellent Quality
Part 3:	Poor Quality = 1	2	3	4	5	6	7 = Excellent Quality

Did you find the light from the projectors to be...

Part 1:	Annoying = 1	2	3	4	5	6	7 = Unnoticeable
Part 2:	Annoying = 1	2	3	4	5	6	7 = Unnoticeable
Part 3:	Annoying = 1	2	3	4	5	6	7 = Unnoticeable

Overall, how do you rate the display technology for the task performed...

Part 1:	Definite Dislike = 1	2	3	4	5	6	7 = Liked very much
Part 2:	Definite Dislike = 1	2	3	4	5	6	7 = Liked very much
Part 3:	Definite Dislike = 1	2	3	4	5	6	7 = Liked very much

Which display technology did you like the most? ___Part 1 ___Part 2 ____Part 3

Why did you like it more than the display technologies used in the other parts?

____

____

____

Which display technology did you **DISLIKE** the most? __Part 1 __Part 2 __Part 3

What made you dislike it?

____

____

____

Figure 57: Post Study Questionnaire - User Study (Chapter 7)

Participant ID ______

In this study, you used three different projection technologies. A description of the three (not necessarily in the order you used them in the study) follows:

Two-Projector Switching

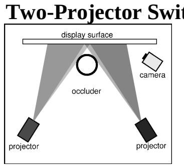

The projector on the left illuminates the screen. When a user blocks the left projector, it turns off, and the right projector fills in the shadow. Light from the projectors usually does not shine on the users.

Two-Projector Simultaneous

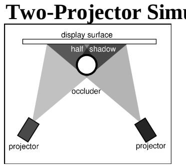

Two projectors illuminate the screen from both sides. Users create “half-shadows” where the screen is still visible within the shadow. Light from the projectors shine on the users.

One-Projector Warped

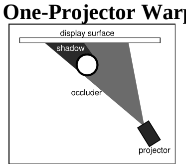

A single projector illuminates the screen from the right side. The user's shadow falls on their left. Light from the projector shines on the users.

Please tell us the order in which you used these conditions in the study. If you are unsure about the exact condition, make your best guess.

First: ____________________
Second: __________________
Third: ___________________

Now, please tell us how sure you are of the accuracy of your answers above by circling the number that best represents how sure you are:

First Condition:
Unsure my answer is accurate **1 2 3 4 5 6 7** Very sure my answer is accurate

Second Condition:
Unsure my answer is accurate **1 2 3 4 5 6 7** Very sure my answer is accurate

Third Condition:
Unsure my answer is accurate **1 2 3 4 5 6 7** Very sure my answer is accurate

Figure 58: Post Study Order Questionnaire (one example of three with rotated ordering) (Chapter 7).

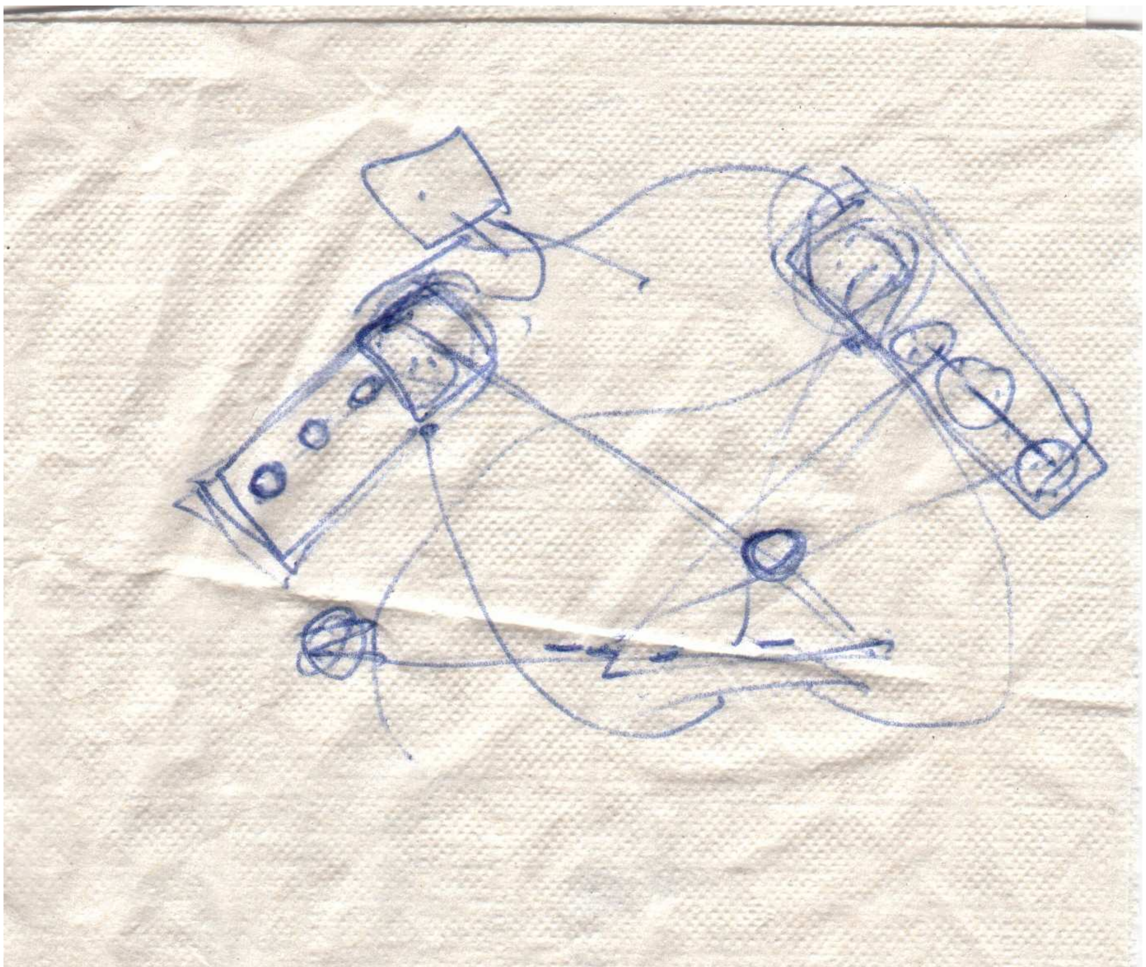

Figure 59: Initial design sketch of a virtual rear projection system.

REFERENCES

[1] Mark Ashdown, Matthew Flagg, Rahul Sukthankar, and James M. Rehg. A Flexible Projector-Camera System for Multi-Planar Displays. In *Proceedings of CVPR 2004*, pages II–165–II–172, 2004.

[2] Samuel Audet and Jeremy R. Cooperstock. Shadow removal in front projection environments using object tracking. In *Proceedings of the IEEE Workshop on Projector Camera Systems (PROCAMS 2007)*, In conjunction with CVPR 2007, Minneapolis, 2007.

[3] Ravin Balakrishnan, George W. Fitzmaurice, Gordon Kurtenbach, and William Buxton. Digital tape drawing. In *ACM Symposium on User Interface Software and Technology*, pages 161–169, 1999.

[4] D. Ballard and C. Brown. *Computer Vision*. Prentice-Hall, 1982.

[5] Deepak Bandyopadhyay, Ramesh Raskar, and Henry Fuchs. Dynamic shader lamps: Painting on real objects. In *The Second IEEE and ACM International Symposium on Augmented Reality (ISAR'01)*, 2001.

[6] Patrick Baudisch, Nathaniel Good, Victoria Bellotti, and Pamela Schraedley. Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming. In *CHI'02 Proceedings*, pages 259–266, 2002.

[7] Patrick Baudisch, Nathaniel Good, and Paul 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.

[8] Michael Black, Francois Berard, Allan Jepson, William Newman, Eric Saund, Gudrun Socher, and Michael Taylor. Digital office: Overview. In *AAAI Spring Symposium on Intelligent Environments*, 1998.

[9] S. E. Burns, K. Reynolds, W. Reeves, M. Banach, T. Brown, K. Chalmers, N. Cousins, M. Etchells, C. Hayton, K. Jacobs, A. Menon, S. Siddique, P. Too, C. Ramsdale, J. Watts, P. Cain, T. von Werne, J. Mills, C. Curling, H. Sirringhaus, K. Amundson, and M. D. McCreary. A scalable manufacturing process for flexible active-matrix e-paper displays. *Journal of the Society for Information Display*, 13(7):583–586, July 2005.

[10] Ozan Cakmakci and Jannick Rolland. Head-worn displays: A review. *Journal of Display Technology*, 2(3):199–216, September 2006.

[11] Han Chen, Rahul Sukthankar, Grant Wallace, and Kai Li. Scalable alignment of large-format multi-projector displays using camera homography trees. In *Proceedings of Visualization 2002*, pages 339–346, 2002.

[12] Scott Elrod, Richard Bruce, Rich Gold, David Goldberg, Frank Halasz, William Janssen, David Lee, Kim McCall, Elin Pedersen, Ken Pier, John Tang, and Brent Welch. Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 599–607. ACM Press, 1992.

[13] Wolfgang F. Engel. *Shaderx2: Shader Programming Tips & Tricks With Directx 9*, chapter Advanced Image Processing with DirectX 9 Pixel Shaders. Wordware Publishing, 2003.

[14] K. Fujii, M.D. Grossberg, and S.K. Nayar. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 814–821, Jun 2005.

[15] Thomas A. furness and Joel S. Kollin. Virtual retinal display, U.S. Patent 5,467,104.

[16] Kunal S. Girotra, Yong-Mo Choi, Byoung-June Kim, Young-Rok Song, Beomrak Choi, Sung-Hoon Yang, Shiyul Kim, and Soonkwon Lim. Pecvd-based nanocrystalline-silicon tft backplanes for large-sized amoled displays. *Journal of the Society for Information Display*, 15(2):113–118, February 2007.

[17] Michael D. Grossberg, Harish S. Peri, Shree K. Nayar, and Peter Belhumeur. Making one object look like another: Controlling appearance using a projector-camera. In *Proceedings of CVPR 2004*, 2004.

[18] F. Guimbretière, M. Stone, and T. Winograd. Fluid interaction with high-resolution wall-size displays. In *Proceedings of UIST*, pages 21–30, 2001.

[19] François Guimbretière. *Fluid Interaction for High Resolution Wall-Sized Displays*. PhD thesis, Stanford University, January 2002.

[20] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, New York, NY, 2000.

[21] R.A. Hayes and B. J. Freenstra. Video-speed electronic paper based on electrowetting. *Nature*, 425:383–385, 2003.

[22] Greg Humphreys, Matthew Eldridge, Ian Buck, Gordan Stoll, Matthew Everett, and Pat Hanrahan. Wiregl: a scalable graphics system for clusters. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 129–140. ACM Press, 2001.

[23] Greg Humphreys and Pat Hanrahan. A distributed graphics system for large tiled displays. In *Proceedings of IEEE Visualization*, pages 215–227, 1999.

[24] Greg Humphreys, Mike Houston, Ren Ng, Randall Frank, Sean Ahern, Peter D. Kirchner, and James T. Klosowski. Chromium: a stream-processing framework for interactive rendering on clusters. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 693–702. ACM Press, 2002.

[25] 3M IdeaBoard. http://www.3m.com/.

[26] E Ink Inc. http://www.eink.com.

[27] Hiroshi Ishii, Minoru Kobayashi, and Jonathan Grudin. Integration of inter-personal space and shared workspace: Clearboard design and experiments. In *Proceedings of Conference on Computer-Supported Cooperative Work*, pages 33–42, November 1992.

[28] Hiroshi Ishii, Minoru Kobayashi, and Jonathan Grudin. Integration of interpersonal space and shared workspace: Clearboard design and experiments. *ACM Trans. Inf. Syst.*, 11(4):349–375, 1993.

[29] Christopher Jaynes, Stephen Webb, and R. Matt Steele. Camera-based detection and removal of shadows from interactive multiprojector displays. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 10(3):290–301, May/June 2004.

[30] Christopher Jaynes, Stephen Webb, R. Matt Steele, Michael Brown, and W. Brent Seales. Dynamic shadow removal from front projection displays. In *Proceedings of the conference on Visualization 2001*, pages 175–182. IEEE Press, 2001.

[31] H. C. Jin, I. B. Kang, E. S. Jang, H. M. Moon, C. H. Oh, S. H. Lee, and S. D. Yeo. Development of 100-in. tft-lcds for hdtv and public-information-display applications. *Journal of the Society for Information Display*, 15(5):277–280, May 2007.

[32] Scott R. Klemmer, Mark W. Newman, Ryan Farrell, Mark Bilezikjian, and James A. Landay. The designers' outpost: A tangible interface for collaborative web site design. In *Proc. ACM UIST'01 Symposium on User Interface Software and Technology*, pages 1–10, 2001.

[33] Myron Krueger. *Artificial Reality 2*. Addison-Wesley Professional, 1991. ISBN:0-201-52260-8.

[34] Kai Li, Han Chen, Yuqun Chen, Douglas W. Clark, Perry Cook, Stefanos Damianakis, Georg Essl, Adam Finkelstein, Thomas Funkhouser, Timothy Housel, Allison Klein, Zhiyan Liu, Emil Praun, Rudrajit Samanta, Ben Shedd, Jaswinder Pal Singh, George Tzanetakis, and Jiannan Zheng. Building and using a scalable display wall system. *IEEE Computer Graphics and Applications*, 20(4):29–37, July 2000.

[35] Kai Li and Yuqun Chen. Optical blending for multi-projector display wall system. In *Proceedings of the 12th Lasers and Electro-Optics Society 1999 Annual Meeting*, Piscataway, N.J., 1999. IEEE Press.

[36] Simply LiveWorks. http://www.wearesimply.com/.

[37] Blair MacIntyre, Elizabeth D. Mynatt, Stephen Voida, Klaus M. Hansen, Joe Tullio, and Gregory M. Corso. Support for multitasking and background awareness using interactive peripheral displays. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 41–50, New York, NY, USA, 2001. ACM Press.

[38] Aditi Majuder and Rick Stevens. Effects of projector parameters on color variation in projection based displays. Technical Report 02-022, Department of Computer Science, University of North Carolina at Chapel Hill, 2002.

[39] A. Majumder and R. Stevens. LAM: Luminance attenuation map for photometric uniformity in projection based displays. In *ACM Virtual Reality and Software Technology*, 2002.

[40] Aditi Majumder. Properties of color variation across a multi-projector display. In *SID Eurodisplay*, 2002.

[41] Aditi Majumder and Rick Stevens. Color nonuniformity in projection-based displays: Analysis and solutions. *IEEE Transactions on Visualization and Computer Graphics*, 10(2):177–188, 2004.

[42] Aditi Majumder and Rick Stevens. Perceptual photometric seamlessness in projection-based tiled displays. *ACM Transactions on Graphics (TOG)*, 24(1):118–139, January 2005.

[43] G. Thomas McCollough, Charles M. Rankin, and Megan L. Weiner. Roll-to-roll manufacturing considerations for flexible, cholesteric liquid-crystal display (ch-lcd) media. *Journal of the Society for Information Display*, 14(1):25–30, January 2006.

[44] Ikuhisa Mitsugami, Norimichi Ukita, and Masatsugu Kidode. Multi-planar projection by fixed-center pan-tilt projectors. In *Proceedings of the IEEE Workshop on Projector Camera Systems (PROCAMS 2005)*, In conjunction with CVPR 2005, 2005.

[45] Thomas P. Murtha, Stefanie Ann Lenway, and Jeffrey Hart. Industry creation and the new geography of innovation: The case of flat panel displays. In Martin Kenney and Richard Florida, editors, *Locating Global Advantage: Industry Dynamics in the International Economy*. Stanford University Press, 2004.

[46] Elizabeth D. Mynatt, Takeo Igarashi, W. Keith Edwards, and Anthony LaMarca. Flatland: New dimensions in office whiteboards. In *CHI'99 Proceedings*, pages 346–353, 1999.

[47] S. Nayar, H. Peri, M. Grossberg, and P. Belhumeur. A projection system with radiometric compensation for screen imperfections. In *First IEEE International Workshop on Projector-Camera Systems (PR OCAMS-2003)*, 2003.

[48] Satoshi Okutani and Michiya Kobayashi. Quantitative evaluation of display characteristics of amoled displays. *Journal of the Society for Information Display*, 14(12):1119–1125, December 2006.

[49] Emilee Patrick, Dennis Cosgrove, Aleksandra Slavkovic, Jennifer Ann Rode, Thom Verratti, and Greg Chiselko. Using a large projection screen as an alternative to head-mounted displays for virtual environments. In *CHI'00 Proceedings*, pages 478–485, 2000.

[50] Elin Ronby Pederson, Kim McCall, Thomas P. Moran, and Frank G. Halasz. Tivoli: An electronic whiteboard for informal workgroup meetings. In *ACM INTERCHI '93*, pages 391–389, 1993.

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

[52] Claudio Pinhanez and Mark Podlaseck. To frame or not to frame: The role and design of frameless displays in ubiquitous applications. In *Proceedings of Ubiquitous Computing (Ubicomp)*, pages 340–357, 2005.

[53] Ramesh Raskar, Michael S. Brown, Ruigang Yang, Wei-Chao Chen, Greg Welch, Herman Towles, W. Brent Seales, and Henry Fuchs. Multi-projector displays using camera-based registration. In *IEEE Visualization*, pages 161–168, 1999.

[54] Ramesh Raskar, Greg Welch, Matt Cutts, Adam Lake, Lev Stesin, and Henry Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays. *Computer Graphics*, 32(Annual Conference Series):179–188, 1998.

[55] George Robertson, Mary Czerwinski, Patrick Baudisch, Brian Meyers, Daniel Robbins, Greg Smith, and Desney Tan. The large-display user experience. *IEEE Computer Graphics and Applications*, 25(4):44–51, 2005.

[56] George Robertson, Eric Horvitz, Mary Czerwinski, Patrick Baudisch, Dugald Ralph Hutchings, Brian Meyers, Daniel Robbins, and Greg Smith. Scalable fabric: flexible task management. In *AVI '04: Proceedings of Advanced Visual Interfaces*, pages 85–89, 2004.

[57] E. Fred Schubert. *Light Emitting Diodes*. Cambridge University Press, 2nd edition, 2006. ISBN:0521865387.

[58] Smart Technologies SmartBoard. http://www.smarttech.com/products/index.asp.

[59] R. Sukthankar, R. Stockton, and M. Mullin. Smarter presentations: Exploiting homography in camera-projector systems. In *Proceedings of International Conference on Computer Vision*, page 247, 2001.

[60] Rahul Sukthankar, Tat Jen Cham, Gita Sukthankar. Dynamic shadow elimination for multi-projector displays. In *Proceedings of CVPR*, pages 151–157, 2001.

[61] Tat Jen Cham, James M. Rehg, Gita Sukthankar, Rahul Sukthankar. Shadow elimination and occluder light suppression for multi-projector displays. In *CVPR Demo Summary*, 2001.

[62] Tat Jen Cham, James M. Rehg, Rahul Sukthankar, Gita Sukthankar. Shadow elimination and occluder light suppression for multi-projector displays. In *Proceedings of Computer Vision and Pattern Recognition*, pages 513–520, 2003.

[63] Jay Summet, Gregory D. Abowd, Gregory M. Corso, and James M. Rehg. Virtual rear projection: Do shadows matter? In *CHI'05 Extended Abstracts*, pages 1997–2000, 2005.

[64] Jay W. Summet, Matthew Flagg, Tat-Jen Cham, James M. Rehg, and Rahul Sukthankar. Shadow elimination and blinding light suppression for interactive projected displays. *IEEE Transactions on Visualization & Computer Graphics (TVCG)*, 13(3):508–517, May/June 2007.

[65] Jay W. Summet, Matthew Flagg, James M. Rehg, Gregory D. Abowd, and Neil Weston. Gvu-procams: Enabling novel projected interfaces. In *ACM Multimedia*, pages 141–144, October 2006.

[66] Desney S. Tan, Mary Czerwinski, and George Robertson. Women go with the (optical) flow. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 209–215, New York, NY, USA, 2003. ACM Press.

[67] Desney S Tan, Mary Czerwinski, and George Robertson. Large displays enhance optical flow cues and narrow the gender gap in 3d virtual navigation. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 48(2):318–333, 2006.

[68] Desney S. Tan, Darren Gergle, Peter Scupelli, and Randy Pausch. Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction*, 13(1):71–99, 2006.

[69] Desney S. Tan and Randy Pausch. Pre-emptive shadows: Eliminating the blinding light from projectors. In *CHI'02 extended abstracts*, pages 682–683, 2002.

[70] John C. Tang and Scott L Minneman. Videowhiteboard: video shadows to support remote collaboration. In *Proceedings of CHI: Human Factors in Computing Systems*, pages 315–322. ACM Press., April 27th - May 2nd 1991.

[71] Khai N. Truong, Shwetak N. Patel, Jay W. Summet, and Gregory D. Abowd. Preventing camera recording by designing a capture-resistant environment. In *Proceedings of Ubiquitous Computing (Ubicomp)*, 2005.

[72] Phillip R. Waier, editor. *Building Construction Cost Data 2002*. Robert S. Means Co., 2002.

[73] Grant Wallace, Otto J. Anshus, Peng Bi, Han Chen, Yuqun Chen, Douglas Clark, Perry Cook, Adam Finkelstein, Thomas Funkhouser, Anoop Gupta, Matthew Hibbs, Kai Li, Zhiyan Liu, Rudrajit Samanta, Rahul Sukthankar, and Olga Troyanskaya. Tools and applications for large-scale display walls. *IEEE Computer Graphics and Applications*, 25(4):24–33, July 2005.

[74] 3M Digital WallDisplay. http://www.3m.com/.

[75] Mark Weiser. The computer for the 21st centry. *Scientific American*, pages 94–104, September 1991.

[76] NEC Projectors with Geometric Correction. http://www.nec-pj.com/products/others/distorted.html. Accessed June 2007.

[77] NEC WT-600. http://www.nec-pj.com/products/wt/.

[78] K. Yamamoto, T. Oguchi, K. Sasaki, I. Nomura, S. Uzawa, and K. Hatanaka. Fabrication and characterization of surface-conduction electron emitters for sed application. *Journal of the Society for Information Display*, 14(1):73–79, January 2006.

[79] Ruigang Yang, David Gotz, Justin Hensley, Herman Towles, and Michael S. Brown. Pixelflex: a reconfigurable multi-projector display system. In *Proceedings of the conference on Visualization '01*, San Diego, California, October 2001. unc.

[80] Sang Kyeong Yun. Spatial optical modulator (som): Samsung's light modulator for next-generation laser displays. *Journal of the Society for Information Display*, 15(5):321–333, May 2007.