DESIGN OF A MULTI-PROJECTOR DISPLAY SYSTEM

Daniel Sud

Department of Electrical and Computer Engineering
McGill University, Montréal

May 2005

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfilment of the requirements of the degree of
Master of Engineering

Your file Votre référence
ISBN: 0-494-12652-3
Our file Notre référence
ISBN: 0-494-12652-3

Abstract

This thesis presents a new multi-projector front-projection display system that can be used for an immersive environment. The implementation of such a system faces new challenges that were avoided by design in CAVE-like systems. These include removing the perspective warping from misaligned projectors and adjusting the intensity to create a uniformly lit display in regions where projectors overlap. Also, when the display is in use, occluding objects between the projector and the display surface cause shadows on the display that must be removed. These three issues have been addressed in this work based in part on algorithms from the field of computational geometry and on techniques from existing projector systems.

Résumé

Cette thèse présente un nouveau système de projection frontale à multiples projecteurs qui peut être utilisé pour un environment immersif. La mise en oeuvre d'un tel système présente de nouveau defis qui avaient été contournés par les systèmes de type CAVE. Ceux-ci incluent corriger de la perspective pour des projecteurs non-alignés et ajuster l'intensité des regions illuminées par de multiples projecteurs afin de créer une image uniformément éclairée. En cours d'utilisation, les objets présents entre les projecteurs et l'écran créent des ombres qui doivent êtres enlevées. Ces trois problèmes ont été résolus dans ce travail basé en parti sur des algorithmes de géometrie computationelle ainsi que des techniques provenant de systèmes de projection existants.

Acknowledgements

I would like to thank my research supervisor Dr Jeremy R. Cooperstock for providing funding and giving me the opportunity to participate in this research. His group offers a very stimulating environment where many masters and doctoral students are working on cutting edge technology and paradigms that have the potential to become part of the communication standards of the following decades. I would also like to thank him for heading McGill's Robocup team that allowed me and fellow students such as Francois Cayouette to build a team of soccer-playing Sony Aibos that competed internationally in Seattle in 2001 and Fukuoka, Japan in 2002. With respect to the work presented in this thesis, I would like to thank Maria Nadia Hilario for participating in the early geometric calibration efforts and for the occlusion detection that allowed a closed-loop system to be implemented. I would also like to thank Stephen P. Spackman for the beneficial technical and programming advice, as well as for his interesting random pieces of knowledge.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

Introduction

1.1. Problem Statement

The goal of the Shared Reality Environment is to give remote participants the impression of being in the same virtual room. For this experience to be convincing, remote participants must appear life-size rather than a few inches tall on a computer screen as in current videoconferencing systems. The large display needed for this immersive environment can be created using multiple projectors.

The goal of this thesis is to develop a multi-projector display system, including methods for 1. removing warping of images due to perspective distortion from misaligned projectors, 2. ensuring a uniformly lit display even when multiple projections are used and 3. removal of shadows when one of the projectors is occluded. The system is intended to be used as a spatially immersive environment. Earlier immersive displays were built using a CAVE-like system[**6**] where these three problems were avoided by design. These systems consisted of three or four screens that were each illuminated by a single mechanically aligned projector. They used rear-projection so no shadows were created by users in the environment. One downside to this approach is that it required a room much larger than the environment itself.

With front-projection, it is possible to use the walls of the room as display surfaces so the environment can be made as large as the room itself. Since it may not be possible to place projectors perpendicular to the screens, the image they project might not be rectangular. While modern projectors offer some form of perspective correction, it can generally be performed only on a single axis through manual adjustments that are fairly time consuming. A software solution can be made much faster and even completely automated with the use of a camera. It may be necessary to use multiple projectors to cover a single surface if it is larger than what a single projector can illuminate. In this case, it will be necessary to overlap the images slightly so that no gap appears while also avoiding the need for perfect alignment. Overlapping regions will appear brighter since they are being illuminated by more than one projector. By dimming these regions in the corresponding projectors, the display can be made uniformly lit. Another use for multiple projectors on a single surface is removing shadows from users blocking the projectors. In this case, as much overlap as possible is desired so that additional projectors will be able to fill in the regions where shadows appear. By performing occlusion detection using camera images, this can be made into a closed-loop system that dynamically removes shadows while the display is in use.

1.2. A Brief Review of Projective Geometry

Projective geometry is at the core of any projector system. Unless a projector's line of sight or optical axis is perfectly orthogonal to the display surface, the image it projects will appear distorted. Similarly, a perspective distortion will occur in the images from an unaligned camera observing the display. Both these phenomena are visible in Figure 1.1. In this figure the projected rectangle appears as an irregular quadrilateral on a whiteboard. The camera's alignment also causes the left of the

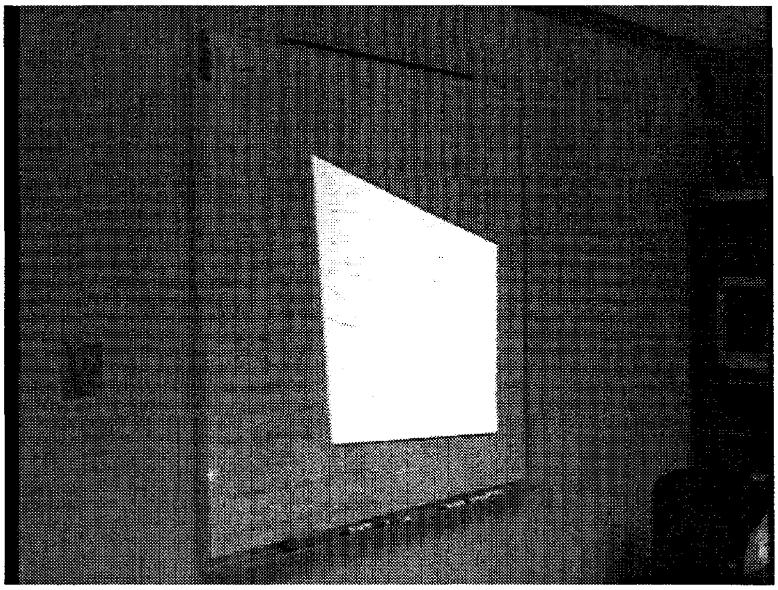

FIGURE 1.1. Camera image observing a misaligned projector illuminating the whiteboard

whiteboard to appear larger than the right side. If this misalignment is only along the horizontal or vertical axis, the distortion is known as keystoning. Recent projectors allow this type of distortion to be manually corrected. For a more general misalignment, along both axes, keystone correction is insufficient to unwarp the image. However, it is possible to compute the perspective transformation causing the distortion and remove it by applying the inverse of this transformation to the image. These two transformations cancel each other, and thus, the resulting projection appears rectangular on the display surface. To obtain these transformations, it is not necessary to know the internal or intrinsic parameters of a projector nor its position and orientation, otherwise known as its extrinsic parameters. Since a camera can be modeled as a projector, its position can also be compensated for by warping camera images. It will now be explained how to compute the required transformations by expressing images in the projective space.

A point in the image plane or the planar display surface can be represented by a pair of coordinates (x, y) in the two dimensional plane $\mathbb{R}^2$. If $\mathbb{R}^2$ is considered as a vector space, the point is in fact a vector. The non-linear mapping from the image plane to the display surface can be represented using linear equations by expressing points as elements of the projective space $\mathbb{P}^2$. In, the projective space $\mathbb{P}^2$, points are written in homogeneous coordinates as a triple (x', y', t), where $t \neq 0$, $x = \frac{x'}{t}$ and $y = \frac{y'}{t}$. Any vector (kx', ky', kt) is equivalent to (x', y', t) for $k \neq 0$. In consequence, (x, y) in $\mathbb{R}^2$ can be expressed in many ways in $\mathbb{P}^2$ but a convenient representation is simply $(x, y, 1)$ where $t = 1$. Once the points have been represented in homogeneous coordinates, it is possible to represent the mapping from one plane to the other as a matrix. Following the convention, boldface symbols such as $\mathbf{x}$ will be used to represent a column vector and express points using column vectors such as $(x', y', t)^T$, which corresponds to the transpose of the row vector (x', y', t). Theorem 2.10 in Multiple View Geometry [**9**] states that

> A mapping $h : \mathbb{P}^2 \mapsto \mathbb{P}^2$ is a projectivity if and only if there exists a non-singular 3x3 matrix $\mathbf{H}$ such that for any point in $\mathbb{P}^2$ represented by a vector $\mathbf{x}$ it is true that $h(\mathbf{x}) = \mathbf{H}\mathbf{x}$.

This means that an invertible 3x3 matrix can be found that will map points from the image plane to the display surface. It also means that the inverse of this matrix will map points from the display surface back to the image plane. A projectivity can also be known as a collineation, a projective transformation or a homography. The calculation of such a transformation is presented below.

Given that a 3x3 matrix maps points in $\mathbb{P}^2$ from one plane to another, the transformation is as follows:

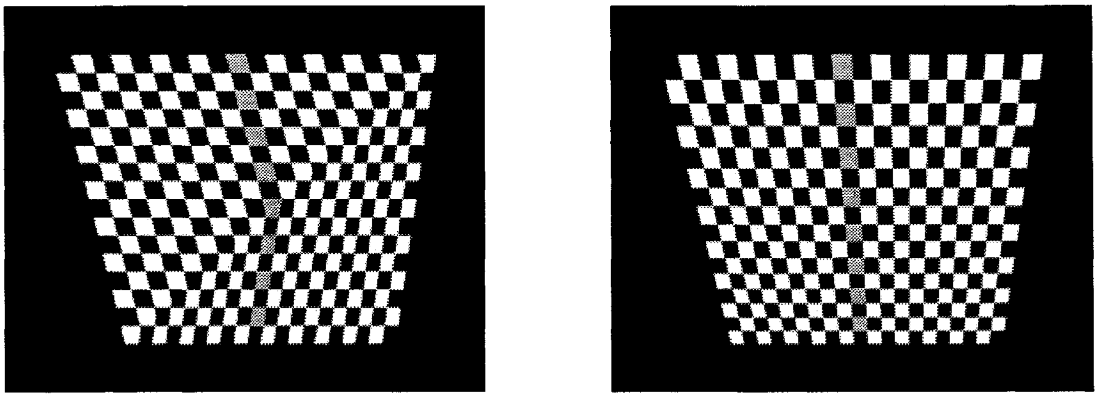

(a) Point correspondences used to move corners only

(b) Projective transformation used on entire surface

FIGURE 1.2. Warping the projector framebuffer to appear rectangular on the display surface

$$\begin{pmatrix} x1' \\ x2' \\ x3' \end{pmatrix} = \begin{pmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{pmatrix} \begin{pmatrix} x1 \\ x2 \\ x3 \end{pmatrix}$$

In short, this can also be written as $\mathbf{x}' = \mathbf{Hx}$. All the values of $\mathbf{H}$ can be determined from point correspondences. While $\mathbf{H}$ has 9 entries, it is only defined up to scale since multiplying it by a non-zero scale factor does not affect the projective transformation. When returning the point $(x1, x2, x3)^T$ to the real plane, $x = \frac{x1}{x3}$ and $y = \frac{x2}{x3}$ so there will be no difference if $x1$, $x2$, and $x3$ are all multiplied by a non-zero scaling factor k. Therefore $\mathbf{H}$ only has 8 degrees of freedom or independent variables. Similarly, points in homogeneous coordinates have only two degrees of freedom. This means that $\mathbf{H}$ can be determined from four pairs of corresponding points since each pair provides two degrees of freedom. Figure 1.2 shows that point correspondences alone are insufficient to correctly warp the projector content. In Figure 1.2(a), the corners have been moved but linear texture mapping, the operation responsible for drawing the interior of the quadrilateral, results in an incorrect image. The outline is correct but the surface is rendered differently in the two triangular parts of the quadrilateral. Lines that should be straight appear broken at the common edge of the triangles. The problem is fixed in Figure 1.2(b) by computing the projective transformation from the point correspondences and applying it to the entire surface.

As mentioned above, using four point pairs will determine **H** exactly. However, it will be imprecise when dealing with pixels instead of points since pixels represent a small area instead of a single point. The solution in this case is to use many more pixel correspondences and to find a matrix **H** that best fits the input using estimation techniques. The preferred approach is the Gold Standard algorithm, described in chapter 4 of Multiple View Geometry [**9**].

There are some limitations to the preceding approach. First, it assumes that the display surface is planar. If this is not the case, a different approach will be needed. One such approach is described in section 1.3. If the display surface is a composition of planar surfaces, then a homography can be computed for each one and the projective distortion can be removed independently for each plane. It is also assumed that the camera or projector can be accurately modeled using the basic pinhole model where the camera or projector is modeled as a single point and the lens introduces no distortion. If this is not the case and radial distortion is present, it will first need to be removed to produce a linear image before point correspondences can be obtained. The second limitation is a projector or camera's depth of field. If the angle between the display surface normal and the projector's optical axis is too large, parts of the image will appear blurred. There is very little that can be done about this other than trying to reposition the projector in question. Similarly, in a camera at an angle or where the focal length is not matched to the display size, many projected pixels will appear in the same camera pixel thus limiting its precision. Again, repositioning the camera is the only way to improve this situation.

1.3. State of the Art in Multi-Projector Systems

In recent years, many efforts have been made to create projector-based displays, most notably at the University of North Carolina at Chapel Hill [**16**] , at the University of Kentucky [**10**] and at the Cambridge Research Laboratory [**18**]. The initial step in all systems is the geometric calibration of projectors and cameras. Most rely on point correspondences to recover the homography between projectors, cameras and the display surface. A calibration pattern or a series of patterns is projected and the resulting image on the display surface is automatically observed by a camera. One method to obtain sub-pixel accuracy in the projector to camera homography is to project circles and locate their centroid in camera images [**19**]. Another method [**16**], also used in the work presented here, is to project a single checkered pattern where the corners of each cell can also be found with sub-pixel accuracy. Once all the homographies have been found, the warping homography that will produce the corrected output for each projector is generated. Steele goes a step further in the Metaverse project[**17**]. With the additional requirement of having a camera fixed to a projector, he is able to recover the three-dimensional position and orientation of the projector with respect to the planar display. The warping homography can then be decomposed into the appropriate rotation, translation and scaling transformations. This approach also continuously monitors and corrects the geometry of the display while it is in use. On the downside, the monitoring approach in its current form is only applicable for a single projector.

To locate features on the display surface, some approaches make use of landmarks such as contrasting dots placed at the corners of the intended display or take advantage of a contrasting border that surrounds the display [**19**]. Alternatively, it is interactively specified [**18**] by having users click on the corners. In a multi-planar system, the edge between planes is located by observing discontinuities in projected lines [**3**]. More recently, a projector is manually positioned and other projectors are aligned with it [**10**].

All the previously mentioned approaches make the assumption that the display surface is planar or piecewise planar. In the latter case, the image is broken down into regions and a homography is computed for each one. An iterative refinement can then be applied to improve consistency between planes [**3**]. It is also possible to create an undistorted display on a non-planar surface [**20**]. This was done by projecting a series of horizontal and vertical bands to recover as many point correspondences as possible over the entire surface. A mesh is then computed and a homography is computed for each cell, which is treated as piecewise planar. With this approach it is possible to create a display on surfaces that are curved or irregular. One limiting factor with this approach is that the display will be correct only from the camera's point of view so its placement is crucial.

1.3.1. Projector Mosaic. When multiple projectors are used to form a larger screen, they are positioned so that the images overlap slightly. This avoids the need to mechanically align the projectors while ensuring that every part of the display is lit. Since the overlapping regions will appear brighter, it is necessary to adjust the intensity in those regions. Current approaches use a pixel-by-pixel intensity adjustment [**16, 14**]. Whether for projector mosaic or shadow removal, it is generally assumed that projectors contribute equally to the illumination of the display [**16, 18**]. Subsequently, the intensity is adjusted to add to unity for every location on the display. Majumder produces a complete radiometric calibration [**14, 13**], taking into account interprojector as well as intraprojector variation, or the variation of intensity within a single projector. It can be concluded from this calibration method that the intensity in a single projector is brighter at the center than at the edges; therefore a truly uniform display must be set to an intensity that is attainable for every pixel.

A downside to this approach is that using the outermost pixels to calibrate a display dramatically reduces its dynamic range.

1.3.2. Systems with Shadow Removal. A major problem of calibrated front-projected displays is that users or objects between the display and the projector will create occlusions that mask part of the display. Distracting shadows will be visible and some information will be lost. To compensate for this problem, multiple projectors can be positioned to illuminate the same area. This way, once occlusion has been detected, another projector can compensate by filling in the missing area [**18, 11**].

In a two projector system, it suffices to increase the intensity in both projectors in occluded regions. Therefore, both projectors could share the exact same intensity map in a fully overlapped display [**18**]. This system presents two problems. First, it will be necessary to detect overly bright areas once the occluding object is removed. Second, the light on users will also be increased. To avoid blinding them, shadow removal has been extended to include occluder light suppression [**5**]. With occluder light suppression, one projector must turn on while the other turns off. If both projectors were displaying the same image, it is not initially known by the system which projector is occluded. The answer must be found in a few frames by trial and error. With both projectors set at half-intensity, the information remains available on the display even when one of the projectors is occluded. An alternative is to have a projector illuminate the display at full intensity and another be completely off [**11**]. While the information will temporarily disappear when the light is blocked, the shadow and the light on the occluder can be removed without any guesswork. This approach also simplifies the detection process by making shadows much more obvious.

To detect occlusion, one approach is to have a camera image of every frame that will be projected [**18**]. Then, occlusion detection consists of a simple image differencing. While this is feasible in the context of a presentation, it is not suited for immersive environments where the projected content is generated dynamically. In this case, it is necessary to use color transfer functions that will allow prediction of the camera image based on the projectors' framebuffers [**11**], which are the rectangular images in video memory that are sent to the projectors. The actual camera image is then compared to the predicted one by the system to determine if occlusion is present.

Occlusion regions are usually treated on a pixel-by-pixel basis, much like for intensity adjustment [**16**]. More recently, the Metaverse team experimented with a region-based method [**10**] where occluded regions were represented using a bounding box. An advantage of the region-based method is in its compact representation that requires only four coordinates instead of an entire image map. When the projector system is distributed on multiple computers connected in a network, this method will use considerably less bandwidth, allowing the system to be more scalable.

1.4. Thesis Outline

In this thesis, a new projector system is proposed that builds on existing systems and proposes some interesting innovations including improved geometric calibration, intensity blending and shadow removal methods. Chapter 2 covers the geometric calibration of individual projectors through the use of homographies. It also discusses an efficient algorithm to compute the largest rectangle, which is useful to optimize the size of the display. Chapter 3 deals with combining multiple projectors and intensity blending in overlapping regions. It also presents a new passive shadow removal scheme. In Chapter 4, a closed-loop system is discussed involving occlusion detection to remove shadows while the display is in use. Finally, Chapter 5 gives the conclusions of this thesis and provides a summary of the contributions made by this work.

CHAPTER 2

Geometric Calibration of a Single Projector

2.1. Introduction

The first step to creating a front-projection display is the geometric calibration of individual projectors. Unless a projector is orthogonal to the screen, the image it projects will appear distorted. The goal of the geometric calibration phase is to unwarp the projected image so it appears rectangular on the wall. As described in section 1.2, the way to calibrate a projector is to compute the projective transformation or homography that will undo the distortion caused by the projector's misalignment. For this purpose, two different methods were developed in this work. The first method is an interactive one where the user clicks on the desired corners of the screen. With this approach, a fully overlapped display can be created very quickly to experiment with shadow removal. The topic of shadow removal itself is dealt with in chapters 3 and 4. The second geometric calibration method developed for this work uses a camera to determine the outline of the projected image on the display surface. By locating multiple projected images on a common display surface, a large display can be created by combining them. Both these methods are presented in the following sections.

For both geometric calibration methods, a rectangular image can be projected once the warping homography is obtained by using an existing method [**16**]. First, an orthographic projection is set up, meaning that the projection matrix will be the identity matrix. This matrix is then multiplied by the warping homography. The framebuffer, which initially contained a rectangular image will appear distorted and the initially distorted image will now appear rectangular on the display surface.

2.2. Calibration using Manual Corner Selection

The first calibration method is an interactive one where the user directly selects the locations of the new corners of the image within the area of the projected image. Initially, the system projects a checkered pattern that takes up the entire projector framebuffer. Using a mouse to move the cursor directly on the display surface, the user then clicks on the four corners in a clockwise direction starting from the top left. These can be chosen by placing markers on the display surface before the calibration process takes place. As each corner is clicked, the surface is immediately redrawn without computing a new homography to reflect the updated corner positions. This allows the user to see the evolving outline of the display. Once all four corners have been clicked, the warping homography can be computed. In this case, the warping homography $\mathbf{H}_{warp}$ is simply a mapping from the old corners p to the new ones p'. This homography, $\mathbf{H}_{pp'}$, is then multiplied with the projection matrix, which is used to render the display correctly in subsequent frames. This method has already been employed in an existing system [**18**] and can even be found in certain projector models manufactured by NEC under the name 3D Reform[**15**]. A fully overlapped display can be created easily using this method by clicking on the same four corners

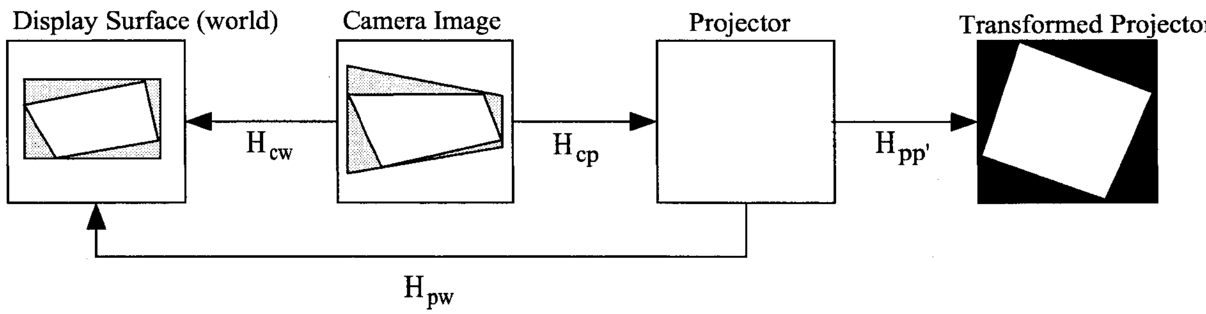

FIGURE 2.1. Different views in the geometric calibration of a single projector

for another projector. A disadvantage of this method is that the resulting display is not guaranteed to be rectangular unless the markers used to determine the corner locations have been very carefully positioned or a rigid frame is used. Also, when multiple projectors are used to form a large display, it will be necessary to accurately position a large number of markers to obtain a calibrated display. In this case, it is preferable to automate the process by including a camera and using the display surface itself as a common frame of reference for all projectors.

2.3. Automatic Calibration using a Camera Image

A more elaborate method was also developed for cases where more than two projectors are used and when they are not fully overlapping. Instead of directly selecting the corners of the display, this approach recovers the shape of the projected image on the display surface using a camera. The largest rectangle within the illuminated area is then selected and its corners become those of the calibrated display. As in the manual approach, the warping homography is then computed as the mapping from the old corners to the new ones. Figure 2.1 illustrates all the views that are needed for this calibration method as well as the homographies relating them. The process shown in the figure is as follows.

To locate projectors on the display surface, it is necessary to use a camera positioned in a way that the entire illuminated area is visible in camera images as well as at least four features from the display surface. A few intermediate steps are needed before the mapping between the projector and display surface is obtained. The first step of this process is to recover the mapping between the camera and the display surface (world) $\mathbf{H}_{cw}$. This is done by locating features of the display surface in the camera image. This can be done using image processing techniques if the features consist of clearly identifiable markers or a contrasting border. Alternatively, a user can manually select the corners of the display surface by clicking on them in a camera image displayed on a computer screen. Once at least four features have been located, the homography can be recovered from point correspondences.

The next step of this geometric calibration process consists of recovering the mapping between the camera and projector $\mathbf{H}_{cp}$. This is done by projecting a calibration pattern or a series of patterns and detecting feature points in the camera image. There are several techniques that can be used to achieve sub pixel accuracy such as projecting a series of circles [**19**] or a singled checkered pattern [**16**]. In this work, a checkered pattern is projected and Bouguet's Camera Calibration Toolbox for Matlab [**4**] is used to extract the grid corners and compute an accurate homography. Having many more than four pixel correspondences will significantly increase the accuracy of the homography by reducing the effect of outliers, that is, pixel correspondences inconsistent with the rest of the data set. For example, if a 16x20 grid is used, there are 15x19 or 285 point correspondences available since only those in the interior of the grid are considered. All the point correspondences can then be used to accurately compute the corresponding homography.

The final step of this process is to recover the homography between the projector and world $\mathbf{H}_{pw}$. While this cannot be measured directly, it can be obtained by combining the previous two homographies as follows:

$$\mathbf{H}_{pw} = \mathbf{H}_{cp}^{-1}\mathbf{H}_{cw}$$

At this point, the homography describing how the projected image appears on the display surface has been obtained but not the one needed to restore a rectangular image. One possible approach is to use $\mathbf{H}_{wp}$, the inverse of $\mathbf{H}_{pw}$ and multiply it by a scaling factor to maximize the display area. This method requires the use of a heuristic method [**19**] to determine the scaling factor; however this may not give the optimal solution and further, the necessary details are not provided by the authors.

This work proposes a new deterministic approach that is guaranteed to provide the largest display area. In this approach, the largest rectangle within the projected image is determined in world coordinates. Using $\mathbf{H}_{wp}$, the corners of the largest rectangle are then transferred back to projector coordinates p'. The warping homography $\mathbf{H}_{warp}$ then becomes a mapping from the original corners to the newly obtained ones $\mathbf{H}_{pp'}$.

2.4. Calculating the Largest Rectangle

As explained in the previous section, the display area can be maximized by selecting the largest rectangle on the display surface. In previous work [**16**], computing the largest rectangle has been done by discretizing the area and scanning it for a close to optimal solution. This approach, which was also used in this work as an early prototype, has the advantage that it can be implemented very quickly. On the other hand, it is quite inefficient and can take several seconds to generate a solution. While acceptable for offline calibration, it would definitely need to be improved if

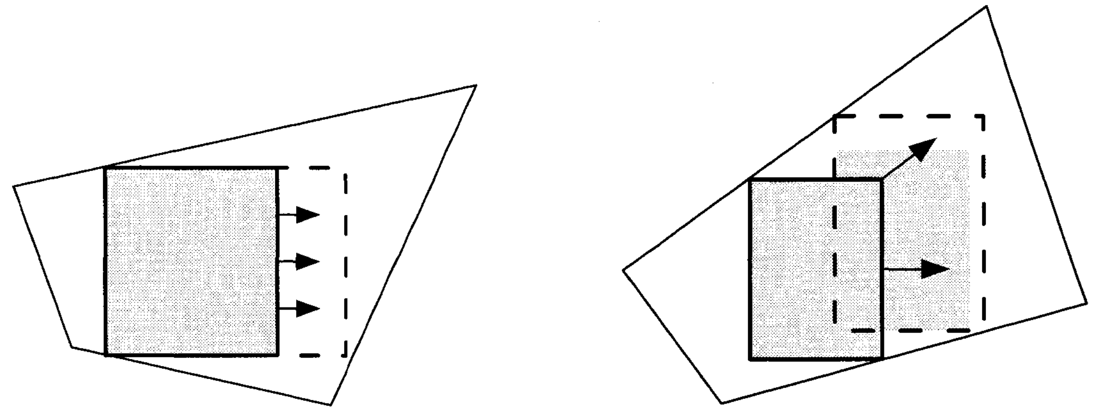

(a) The two corners are not diagonally opposed

(b) Neither corner is on a vertex of the polygon

FIGURE 2.2. A rectangle with 2 corners on the boundary of P can expand in certain cases

the geometry of the display was continuously monitored (cf. [**17**]). This has led to the examination of an efficient algorithm [**2**] that can find the largest rectangle in a convex polygon in $O(log\ n)$, where n is the number of vertices in the polygon.

2.4.1. Characterization of the Solution. This algorithm computes the largest rectangle LR inscribed in a convex polygon by taking advantage the geometric characteristics of the solution. They observe that LR will either have two or three corners on the boundary of the convex polygon P. The case where all four corners are in contact with P can be treated exactly like the three-corner case. When exactly two corners of LR are on the boundary of P, these will be diagonally opposed. If this was not the case, the rectangle could be made larger simply by moving the edge not in contact with P as in Figure 2.2(a). It is also observed that in the two-corner case, at least one of the corners of LR must be located at a vertex of P. Otherwise, the rectangle could slide along the edges of P and either expand or eventually make contact at a third corner, which would then be treated as such. This case is illustrated in Figure 2.2(b). In the case where three corners of LR are on the boundary of P, it is possible for none of the corners to be located at a vertex of the polygon. Figure

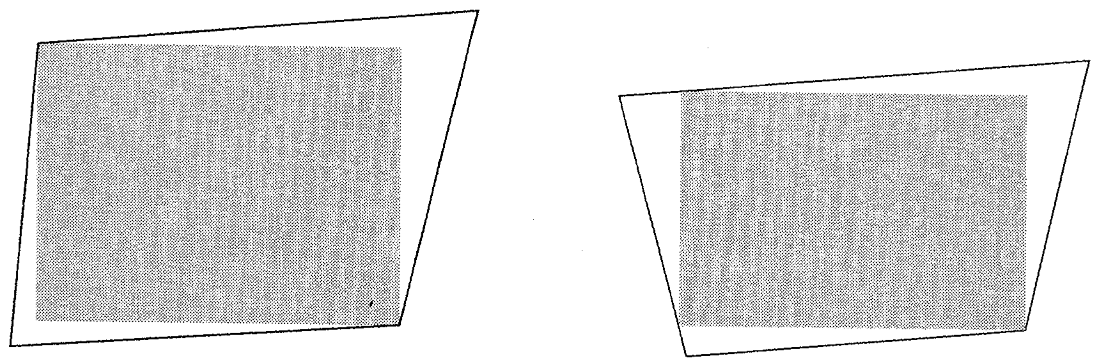

(a) Largest rectangle with 2 corners on the boundary of P

(b) Largest rectangle with 3 corners on the boundary of P

FIGURE 2.3. The largest rectangle inscribed in a convex polygon can have 2 or 3 corners on its boundary

2.3 shows examples of the largest rectangle with either two corners on the boundary in Figure 2.3(a) or three in the case of Figure 2.3(b).

2.4.2. Prune-and-Search for the Solution. Using results from previous work [**12**], the authors show that the solution can be found by expressing the problem as the search for the fixed-point of a composition of two functions in the two-corner case and as the fixed-point of a composition of three functions in the three-corner case. With the appropriate functions, a prune-and-search, or tentative prune-and-search in the three corner case, is then applied to find a solution in logarithmic time. Details have been left out of this work to avoid overwhelming readers but those interested can refer to the original paper [**2**] for more information. There are two cases for the largest rectangle with two corners on the polygon, one for each diagonal of the rectangle, and four cases for the largest rectangle, one for each combination of three corners, that must all be examined. While it is assumed that the polygon is in general position, meaning that no edges of the polygon are parallel, it is possible to introduce perturbations [**8**] allowing the algorithm to be used even when this is not the case.

2.4.3. Simpler Algorithm Based on the Characterization. In the interest of simplicity, it is possible to develop an algorithm that takes into account the characterization of the solution outlined above. By omitting the prune-and-search and tentative-prune-and-search, the implementation is greatly simplified. This new algorithm searches for a solution at every vertex of the polygon. This finds any solution with two corners on the boundary of the polygon and some of the solutions with three corners on the boundary. An additional search at all combinations of three edges of the polygon reveals any remaining solution where the rectangle is not in contact with any vertex. In the case where the display is created using a single projector, the polygon consists of a quadrilateral, so this involves looking for a solution at four vertices and at four groups of three edges. While not optimal, this method is significantly faster than the discretization method where hundreds of tests are made when scanning for the solution.

2.5. Conclusion

In this chapter, two methods were presented for the geometric calibration of individual projectors. The first method relies on selecting the corners directly and is ideally suited for fully-overlapping projectors. For other cases, a more general method was presented that uses a camera to locate projectors on the display surface. This method also maximizes the size of the display by computing the largest rectangle. For this purpose, an efficient algorithm is introduced, as well as a simpler version, based on the characterization of the solution. However, this approach assumes that only one projector is employed to generate the display. In chapter 3, additional projectors are introduced and the resulting issues are examined.

CHAPTER 3

Combining Multiple Projectors

3.1. Introduction

It may be required to use more than one projector to create a display. The addition of more projectors permits two different goals to be achieved and these will determine how the projectors should be positioned. The first is the creation of a display that is larger than what a single projector can produce. For this case, the images will need to be overlapped slightly to remove the need for a perfect mechanical alignment, which can be a very time-consuming process. Such a display can be made larger with each additional projector by keeping most areas of the display covered by only one of them and a seamless illumination can be achieved by adjusting the intensity in the overlapping region. The second objective is to be able to remove shadows from users blocking the projectors. In this case the entire display will need to be covered by multiple projectors. Each of these two cases is examined in detail in the following sections.

3.2. Geometric Calibration in a Large Display

When the display is larger than the area illuminated by a single projector, then it will be necessary to use a slightly different technique from the one presented in the preceding chapter. Instead of computing the largest rectangle within a single projector, it will be computed inside the union of all available projectors. This is done by first locating all projected images on the display surface as in Figure 3.1(a). Since the union of convex polygons is not itself convex, the algorithm described in the previous chapter cannot be used directly. While there exists an algorithm to compute the largest rectangle in any polygon [**7**], its implementation can be quite complex. Instead, the non-convex union can first be clipped to become convex by removing all reflex vertices. These vertices, defined by an internal angle larger than 180 degrees, appear only at the intersection points since they are not present in convex polygons. They can be removed by knowing the relative placement of the projected images. For example, if two images are positioned side by side to form a wide display, the union of the two images can be clipped along horizontal lines at all reflex vertices. For images one above the other, reflex vertices can be removed by clipping along a vertical line. The resulting polygon will be convex, as seen in Figure 3.1(b) and the largest rectangle can be computed normally in this area to create an optimally sized display. The result is shown in Figure 3.1(c). For individual projectors, instead of calculating the largest rectangle within the projected area in display coordinates, it will be necessary to compute the bounding box of the projector on the display surface as shown in Figure 3.1(d). Mapping these points back to projector coordinates will produce the new corners needed to compute the warping homography. With the appropriate homography, the entire area illuminated by the projector can then be used as part of a larger display.

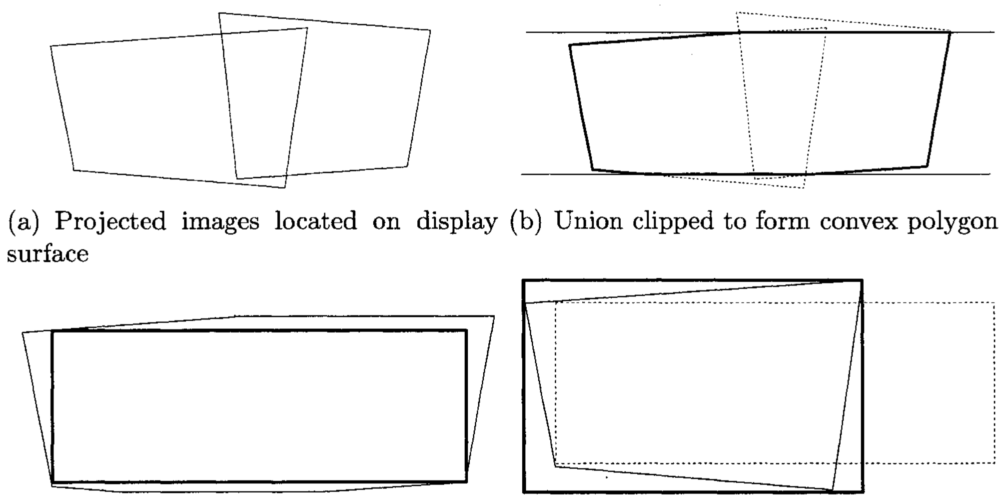

(a) Projected images located on display surface

(b) Union clipped to form convex polygon

(c) Largest rectangle selected in convex polygon

(d) Bounding box used to calibrate individual projectors

FIGURE 3.1. Geometric calibration of projectors in a large display

3.3. Region-Based Intensity Blending

When combining multiple projectors to form a larger display, some overlap is required in the images. This is done to remove the need for mechanical alignment and to completely eliminate the gaps between units. This is an advantage of a projector-based display over traditional monitors where the frame is visible between units. One problem that will arise from this setup is that the region covered by two projectors will appear much brighter than the remainder of the display as can be seen in Figure 3.2(a). It is therefore necessary to locate this region and adjust the intensity of the projectors.

Since the outline of the images projected by each projector is a convex polygon, the region of higher intensity can be located by computing the intersection of convex polygons. An efficient algorithm for this operation is described in the following section. Once the region has been found, different approaches can be taken. Either the region can be assigned to one of the projectors and the others turned off or each projector can contribute a part of the total illumination. It has been found that this

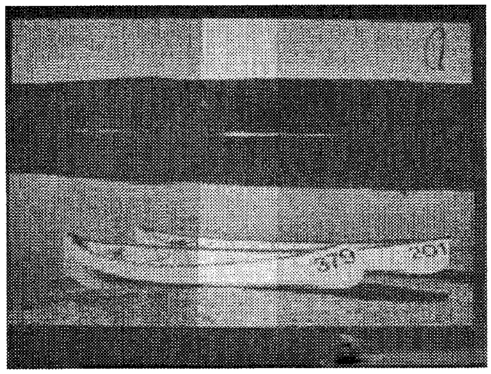

(a) Both projectors at full intensity

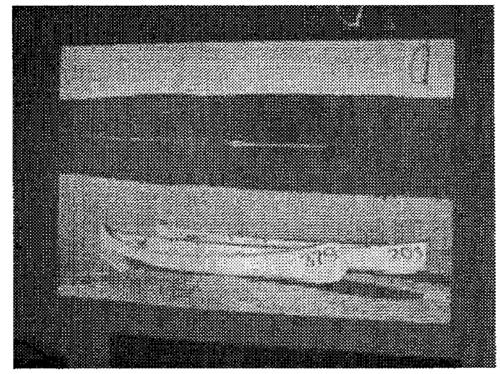

(b) Both projectors with intensity blending

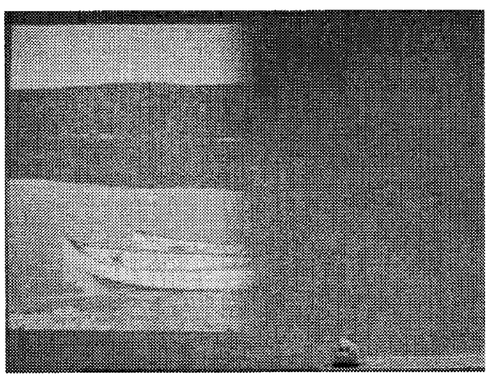

(c) Contribution of the left projector

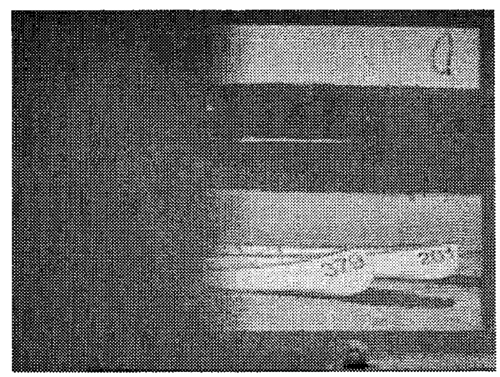

(d) Contribution of the right projector

FIGURE 3.2. Region-based intensity blending in a dual projector setup

method can lead to artifacts at the edges of the region [**16**]. Unless pixels from different projectors line up exactly, there could be an overly bright edge or a gap that is less than a pixel wide where regions meet. The solution proposed previously and also used in this work is to fade from one intensity to the other using a linear ramp. This will reduce the visible artifacts between regions. Existing systems use a pixel-by-pixel approach for intensity blending. At each pixel the intensity of the first of two projectors is $\frac{d1}{d1+d2}$ where d1 and d2 are the distances to the closest boundary of the first and second projector images in world coordinates. Since the second projector will have an illumination of $\frac{d2}{d1+d2}$ at that point, the total illumination will add up to unity over the entire display.

In this work, a region-based approach is proposed to do the same work, including the linear ramp at the edges. In the case of two side-by-side projectors that have already been geometrically calibrated, the overlapping region will be located in the middle of the display. It can be made rectangular by clipping the sides and assigning them to a single projector. In the remaining rectangular portion, the intensity can be reduced linearly from full intensity on one side of the rectangle to completely off on the other side. As in the pixel-by-pixel method, the intensity at every point in the overlapping region adds to unity thus creating the impression of a uniformly lit display. With this method however, it is not necessary to adjust the intensity value at every pixel. Instead, a black rectangle is placed in the image over the overlapping region and its alpha value or opacity is only set at the vertices. On the side where the projector is completely on, the rectangle vertices are set to fully transparent, with an alpha value of zero. The two other vertices are set to fully opaque on the opposing side. When the display is being rendered, the correct intensity at every pixel is determined by the video hardware that automatically performs a linear interpolation between the vertices of the rectangle and lets the appropriate amount of light from the display behind it pass through. The experimental result of this simple technique can be seen in Figure 3.2. In Figure 3.2(a), both projectors are initially illuminating the display at full intensity, creating an undesirable brighter region in the center. This problem has been corrected in Figure 3.2(b) where intensity blending has been applied. The contribution of individual projectors creating this result can be seen in Figures 3.2(c) and 3.2(d).

It is also possible to create a region-based method to correct the intensity for non-rectangular intersection regions. Once the convex polygon representing the intersection region has been found, it can be assigned to either one of the projectors or shared by both at half-intensity. To create a seamless transition between regions, bands must be created where the intensity is faded from one region to the next. This is done by translating the edges of the intersection region inward by a fixed amount. This is similar to methods used for expanding a polygon when doing silhouette edge rendering for toon shading [**1**] except that edges are translated towards the interior instead of the exterior. This results in a smaller polygon inside the intersection region. The intensity of this smaller polygon can then be specified for each projector by setting the alpha values at its vertices. The alpha values also need to be set at the vertices of the larger intersection polygon. Finally, the rendering pipeline on the video card automatically performs the fading effect by linear interpolation between the vertices of the polygons.

3.4. Computing the Intersection of Convex Polygons

When creating a doubly covered display using two projectors for shadow removal, the display will be contained entirely within each projector so the geometric calibration technique presented in the previous chapter can be used. The only difference is that the largest rectangle will not be the one from either projector but rather the one that fits both. In other words, the largest rectangle in the intersection of the two projectors is required. This leads to the investigation of how to compute the largest rectangle in a convex polygon.

Computability of the intersection of convex polygons is useful to determine the different regions in a projector-based display because misaligned projectors will not project a rectangular image but a convex polygon. Also, the intersection of convex polygons is itself a convex polygon. This means that once projectors have been individually located on the display surface, the necessary information will be available to identify all the overlapping regions automatically.

There exist a few methods with linear complexity that serve this purpose, one of the most straightforward, by Toussaint [**22**], is built on other basic computational geometry algorithms. This method works by first finding the convex hull of the two polygons. Then bridges are located on the convex hull. Bridges are edges of the convex hull that are not edges in either of the original polygons. Since there is a one to one correspondence between bridges and intersection points, all the intersection points can be obtained once all the bridges have been found. The polygons are then broken into chains between these intersection points. Finally, merging the inner polygonal chains between intersection points results in the desired intersection. The union of convex polygons can be obtained using the same method but by merging the outer polygonal chains.

To obtain the convex hull of the union of two convex polygons, the rotating calipers method can be used to find a solution in linear time [**21**]. Rotating calipers are parallel lines of support, one on each polygon. By rotating them around the polygon, bridges are found when the lines overlap. The area located between the bridge and the two polygons has a special structure and is known as a sail polygon. By taking advantage of its structure, it is easy to triangulate a sail polygon and find the corresponding point of intersection.

3.5. Passive Shadow Removal

A system that performs passive shadow removal can be created simply by overlapping multiple projectors. In such a system, no detection takes place so shadows cannot be completely removed but their effect can be minimized. In the simplest case where two projectors fully overlap to create a display, the intensity of each projector can be reduced simply to 50%. This way, when only one of the projectors is occluded in an area of the display, the information in that area will be dimmer but will remain visible, as the other projector is unoccluded. With passive shadow removal alone, nothing is done at runtime to deal with occlusion as it appears. Since shadows will still be visible and distracting, it is preferable to supplement passive shadow removal with an active removal process to detect and remove them fully by increasing the intensity of the unoccluded projector. Unfortunately, passive shadow removal introduces difficulties for the automatic detection of shadows.

First, if both projectors are contributing exactly the same information, the system will not know which one is occluded by examining the display when attempting active removal. The solution in this case is to use trial and error to determine which projector is occluded, by alternating an increase of intensity between the two. The occluded projector is then determined by identifying in which of the previous two frames the original image was best restored. It is possible to improve this process by maintaining the probability that each projector is occluded. Once it has been observed that a particular projector is occluded more often, one could start by assuming that it is being occluded and only try the other option if the shadow is not removed. Another alternative would be to assume that new shadows at the edges of an existing one are due to the same projector being occluded. The advantages of using either of these two methods are that, statistically, they require less guesswork and shadows can be removed faster.

The other problem with passive shadow removal, when used in conjunction with an active removal stage, is that it makes the detection much more difficult. This can be seen in Figure 3.3 in which the occluded region of the left image appears completely black, whereas, with two overlapping projectors, as in the right image, it is merely dimmer than the original. If passive shadow removal is the only method used to eliminate occlusion, this will not be a problem. It becomes one when passive shadow

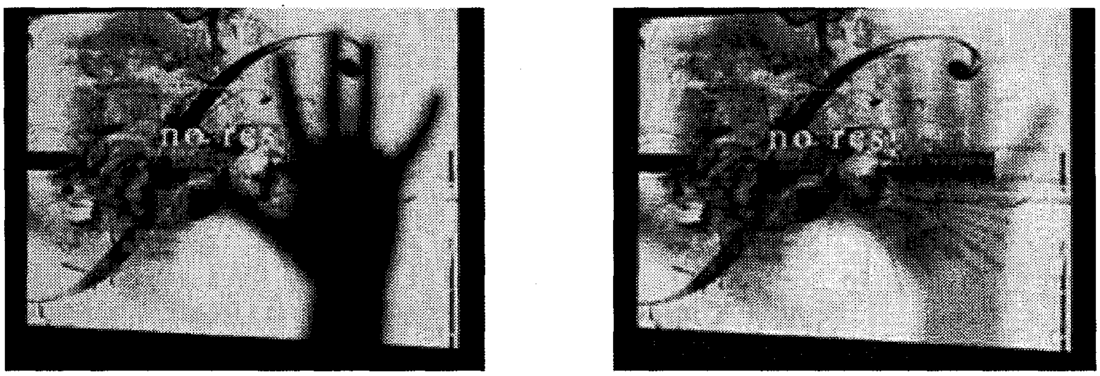

(a) Single projector illuminating the display

(b) Two fully overlapping illuminating the display

FIGURE 3.3. How occlusion appears in different projector setups

removal is used as a prelude to active removal methods. In this case, occlusion detection, which is needed to fully remove shadow, will be more challenging as the region of occlusion still contains a reduced-intensity version of the projected display. For this reason, existing approaches that attempt to remove shadows actively do not generally illuminate the same area of the display with multiple projectors simultaneously.

In this work, a new scheme is presented that attempts to resolve both problems created by the use of passive shadow removal. Instead of illuminating the display with two projectors contributing equally, the color channels are separated among projectors. For example, one projector can project the red channel at full intensity while a second projects the blue and green channels, also at full intensity. Using this system, a shadow is easily determined by the complete absence of certain colors in a given area once any leakage and background illumination have been taken into account. This can be seen in Figure 3.4(a), where, observing that only blue and green are visible in the occluded area, it can be determined that the red projector is occluded. The opposite case is shown in Figure 3.4(b) where only red is visible in the occluded area.

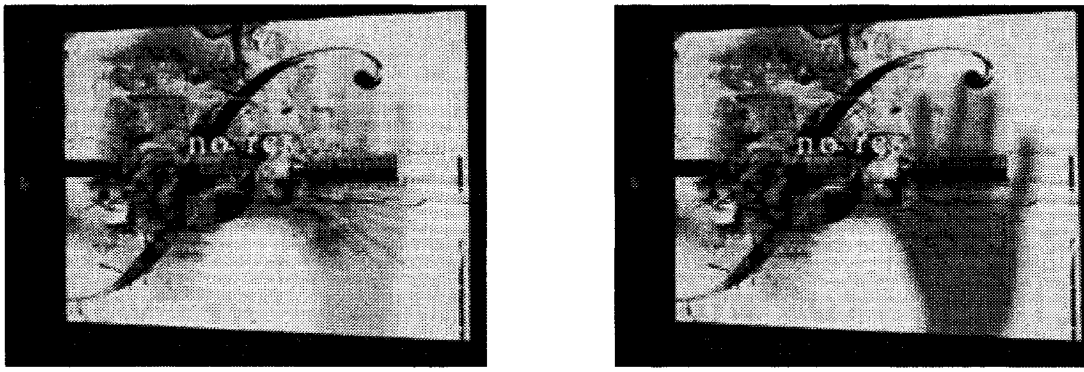

(a) Occlusion of the red projector (b) Occlusion on the green and blue projector

FIGURE 3.4. Improved passive shadow removal

In addition to making detection easier, the method preserves the advantage of passive shadow removal, which is to maintain some visible information in the occluded region. For this method, as for all sections of this work, it is assumed that each projector in a dual projector setup provides an equal amount of illumination to the display. This may not necessarily be the case. As can be seen in Figure 3.5, one projector was manually set to form a larger image, and thus, its light was more diffuse in the part used for the display. Even if the projectors were identical and at similar distances from the display surface, since the red projector is effectively brighter, the image appears too pink. This experiment illustrates the need to examine better models of the projector contributions in future work.

3.6. Conclusion

This chapter introduced the use of multiple projectors to form a single display. Multiple projectors can be used either to form a larger display or one where shadows can be removed by overlapping them. In the case of larger displays, intensity blending is needed in certain regions to create a uniformly lit display and the intersection of convex polygons can be used to locate those regions. For an overlapped display, the

FIGURE 3.5. Projectors not contributing equal illumination to the display

intersection of convex polygons can also be used to obtain the largest possible screen. Without active detection, shadows can be removed passively by illuminating the same region with multiple projectors. However, doing so makes the task of actively removing shadows more difficult. For this reason, a new passive removal scheme was introduced that can be combined with the active removal of shadows without requiring guesswork and without making the detection task overly difficult. The process of occlusion detection followed by active shadow removal is covered in chapter 4.

CHAPTER 4

Shadow Removal with Occlusion Detection

4.1. Introduction

In the previous chapter, passive shadow removal was introduced where multiple projectors continually illuminate the entire display. If one of them becomes blocked, the occluded area appears darker but remains visible. In order to remove occlusion completely, it is necessary to detect it using a camera observing the display. The camera's role is to make sure that the observed display remains consistent with the predicted model. Once occlusion is detected, these regions are communicated back to the process driving the projectors in the form of image maps that are then used to adjust the intensity accordingly. In the next section, occlusion detection is discussed.

4.2. Occlusion Detection

There are two categories of methods to detect occlusion. The first, and also the simplest, is to take a snapshot of the display when it is known that there is no occlusion. This process is repeated for every image that will be displayed before the system is actually used. When in use, the current image is simply compared with the existing snapshots. Image differencing, where each pixel in the current image is compared with the unoccluded snapshot, is all that is needed to detect the presence of occlusion. While fairly simple, this technique is not usable for all applications. It is ideally suited for a slideshow or presentation, where all the content that will be displayed while the system is in use is available beforehand. However, if the content is generated dynamically during use, it will not be possible to use this approach. The immersive environment applications for which it is required to use the projector system fall into the latter case so different detection methods must be sought.

The second category of detection methods involves generating a predicted camera image from the projector framebuffers. With this technique, it is possible to remove shadows even if the projector content is not known a priori. Different algorithms have been proposed to generate the predicted camera images. Each consists of projecting colored patterns and deriving a color transformation that maps the projector framebuffer into a camera image. The initial method in this category proposed matching the samples to an exponential function [**11**]. More recently, we have been developing an alternative color transfer function [1]. The early work involved using a color lookup table where each color in the projector framebuffer is individually mapped to a camera color. More recently, a color transfer matrix was proposed, which consists of a linear approximation to the camera response. One advantage of this method over the previous ones is that generating the predicted camera image can be done more rapidly as it can be handled by the video hardware. The purpose of the generated camera image is to predict what the unoccluded display should look like. Occlusion can then be detected by comparing it with the actual camera image.

[1] This work is being conducted by Maria Nadia Hilario.

To generate predicted camera images, each projector framebuffer must be transformed to the camera's point of view, followed by a color transformation approximating the camera's response to the projected colors. The final predicted image is obtained by accumulating all transformed projector framebuffers into a single image. This method requires that the projector framebuffers be available, as was the case for this work, where two projectors and a camera were all connected to a single computer.

A downside to this method is that it does not scale well when more projectors are involved. Since each computer can only control a limited number projectors, additional computers will be needed if more projectors are added. In this case, it will be necessary to transmit projector framebuffers between computers. Since these images in uncompressed form can be of a few megabytes in size, transmitting them at every frame will considerably slow down the system. It would be tempting to use a simpler method to generate predicted camera images where the display content is directly transformed to the camera view instead of first being rendered for each projector. By generating a camera image directly from the source content, there would be no need to transmit projector framebuffers, resulting in considerable bandwidth savings. Unfortunately, this method is unlikely to produce results of good quality for two reasons. First, some image degradation can occur when each projector is rendered. For example, if high resolution content is rendered on a low resolution projector, the camera will actually observe a low resolution display. If the generated camera image used the high resolution image directly, the detection process would produce many false positives resulting in a noisy detection. Second, the camera response will be different for each projector so the color transformation stage must be handled independently for each. Again, failure to take this into account will increase the error in the generated predicted camera image and increase the noise level in the detection.

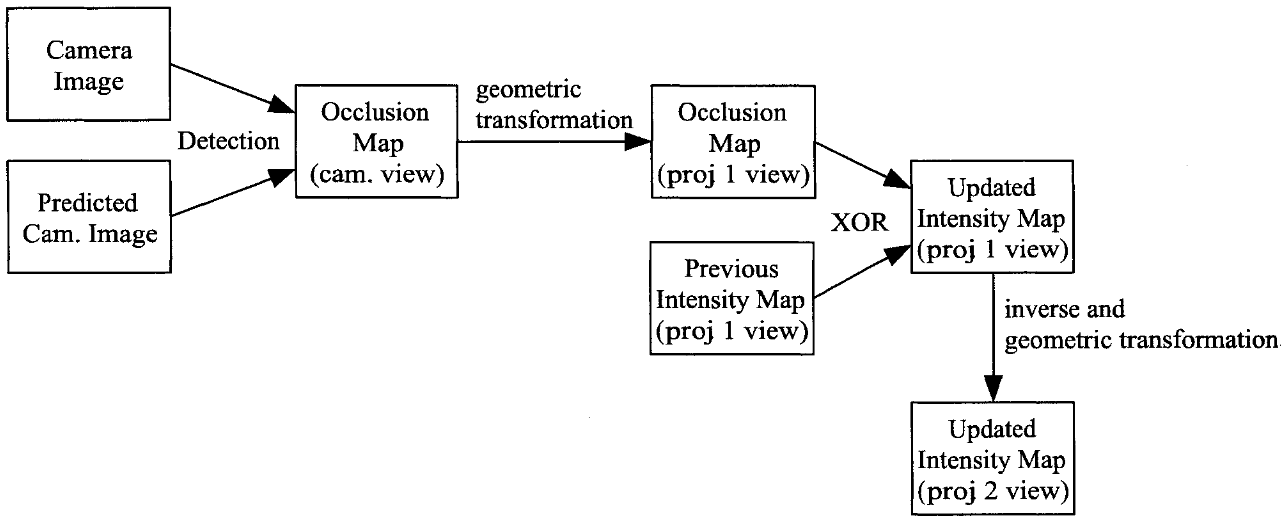

FIGURE 4.1. Process of updating projector intensity maps from detected occlusion

One possibility to allow a scalable system is to add cameras to each computer that is controlling projectors. Each camera would then be responsible for detecting occlusion on the part of the display illuminated by the projectors connected to the same machine. This system would permit both the use of projector framebuffers to generate predicted camera images and the use of multiple computers to control a large number of projectors.

4.3. Image-Based Method for Shadow Removal

Once occlusion has been detected, an occlusion map must be communicated to the process controlling the projectors. The simplest way is for the camera to send a binary image indicating in which pixels occlusion was found. The image is then transformed for each projector that must then use it to turn individual pixels on or off. A simple algorithm has been developed that takes care of the decision process in a two projector setup.

Initially, a first projector illuminates every pixel in the display while the second projector is completely off. Individual pixels can be turned on or off for each projector with the use of a binary intensity map that acts as a stencil, letting light through only for those pixels for which the projector is responsible. For every frame in the closed loop system, a detection process first creates a binary occlusion map identifying which camera pixels differ from the predicted camera image. This occlusion map is then geometrically transformed to the point of view of the first projector. It is then combined with the projector's binary intensity map using the exclusive-or (XOR) operation to produce an updated intensity map. The updated intensity map is then inverted for the second projector. By this operation, every area of the display turned off in the first projector will be turned on in the second one and vice versa. Finally, the inverted intensity map is geometrically transformed for the second projector. A model of this system is presented in Figure 4.1.

To render each projector framebuffer, the intensity map is placed over the geometrically transformed content. The newly obtained projector framebuffers are then displayed by the projectors and are also used to create the predicted camera image for the next frame, thus creating a closed-loop system. The goals of this system are for every pixel to be lit by a single projector and for the projectors to switch when occlusion is detected at a particular pixel. There are four cases to consider depending on whether or not a camera pixel is occluded and whether the first projector is on or off for that pixel. The use of the XOR operation to update the projector's intensity map allows all cases to be dealt with in a single operation. These cases are presented in table 4.1.

Experimental results can be seen for this system in Figure 4.2 where four consecutive camera frames demonstrate occlusion removal using the XOR operation. The system used for this experiment consisted of two projectors at a resolution of 1280x1024 and a camera with a 640x480 resolution all connected to a single machine. The system ran at a framerate of 1 frame per second on average when performing

Occlusion	Initial Intensity	Updated Intensity	Description
0	0	0	No occlusion so projector stays off
0	1	1	No occlusion so projector stays on
1	0	1	Occlusion detected so projector turn on
1	1	0	Occlusion detected so projector turns off

TABLE 4.1. All cases for updating the projector intensity map

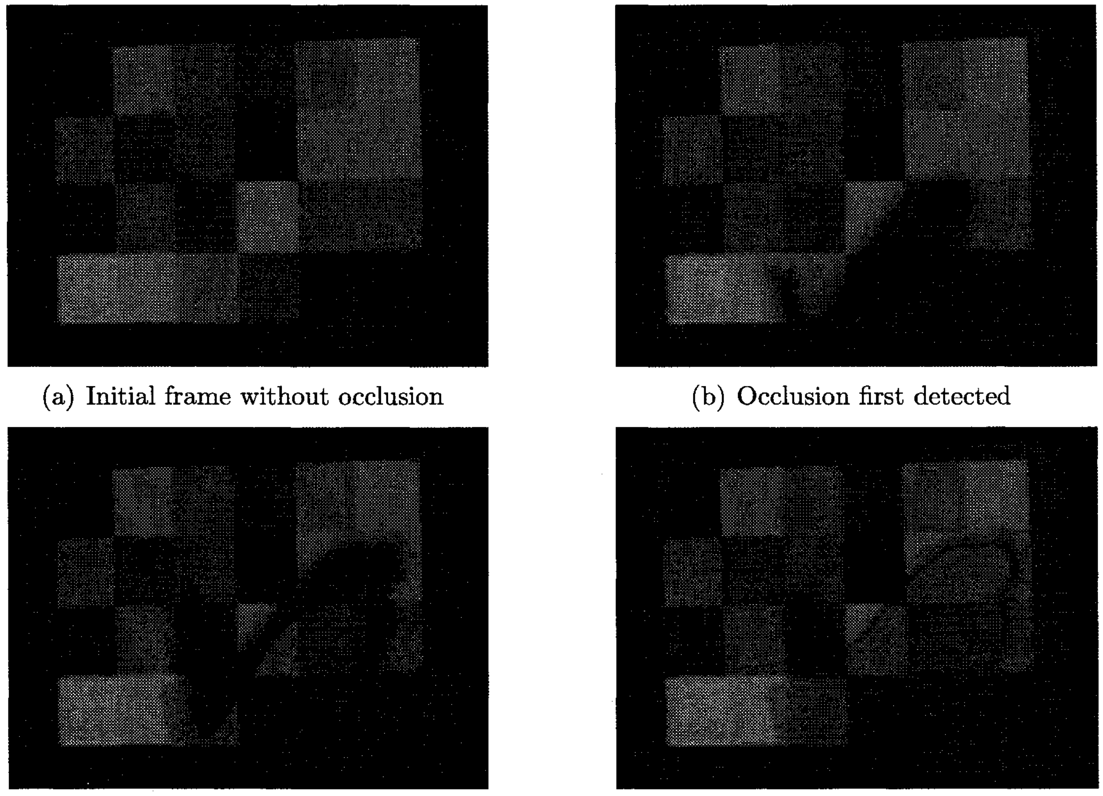

(a) Initial frame without occlusion

(b) Occlusion first detected

(c) Previous occlusion removed but new one visible on its border

(d) Again, only the occlusion from the previous frame is removed

FIGURE 4.2. Consecutive camera image frames in an occlusion detection and removal closed-loop system

detection and removal of shadows and 3 frames per second when performing detection alone. These results are similar to the 2 frames per second obtained by another group [**10**]. They were able to increase the performance up to nine frames per second by reducing the resolution of camera images and using approximate bounding boxes to represent occluded areas. Other methods are currently being investigated to increase the performance of the system even further. These include breaking up the display into tiles and using statistical sub-sampling to determine whether each tile is occluded or not. Also, the camera observing the display could be placed on a separate machine to allow generated predicted camera images and projected framebuffers to be rendered in parallel instead of in series as is currently the case. This would require the use of a region-based method to represent shadows, which is briefly introduced in the following section.

4.4. Region-Based Method for Shadow Removal

The disadvantage of the image-based technique is that sending images from the camera to the projectors requires significant amounts of bandwidth when they are connected to different computers. In the simple system explained above, the two projectors and the camera were all connected to a single computer. The system would become quite inefficient if multiple cameras were connected together in a network of computers. For a system that is more scalable, it is best to turn to a region-based method to communicate occlusion.

In such a system, the binary occlusion image produced by cameras would first be processed to interpret occluded regions as polygons. A very simple form of this method is simply to communicate the bounding box of occlusion regions. This compact representation can then be propagated efficiently to multiple computers that can then reconstruct the occlusion image. This method would also allow occlusion regions to be expanded to allow intensity blending and a smoother transition between projectors.

4.5. Future Directions for Occlusion Removal

This work only dealt with the case where two projectors are used. Some systems have already started tackling the problem of integrating multiple projectors [**10**]. The solution proposed is to use a suitability factor to determine which projector is the best one to illuminate a given area. When it is occluded, the next best one is then selected.

Once occlusion has been removed, there is no way to tell if the occluding object is still present and whether using the best projector can be resumed. The suggestion for this case is to shrink the regions illuminated by secondary projectors thus returning to the primary ones over time. With a region-based method, it would be possible to maintain a history of detected occlusion. The contribution of less desirable projectors could then be removed by reprocessing the occlusion regions after a certain amount of time.

Another problem faced by current occlusion detection and removal methods is that they only correct occlusion after it has become visible. If the occluding object is moving, a new occluded region will appear at the edges of a previously corrected one. Since tracking algorithms are being developed for gesture-based interactions, these could be used to preempt shadows by predicting where occlusion will appear in future frames. Pushing this idea even further, members of the research group are investigating 3D reconstruction methods to build a fully three-dimensional model of the users in the environment using a large number of cameras. This could eventually lead to a system that predicts which projectors will be blocked and thereby remove occlusion before it ever occurs.

4.6. Conclusion

This chapter presented the active removal of shadows using occlusion detection. This work focused on removing occlusion from arbitrary images by predicting the camera images instead of simply trying to restore the content in a closed set of slides. A simple implementation of shadow removal in a two-projector setup was demonstrated. In this system, the projector contribution map was built from camera occlusion maps using the XOR operator. The last two sections discussed what needs to be done to create a scalable system and improve on current techniques. Namely, high-level representation of shadows would allow them to be communicated efficiently between machines. Also, by anticipating where shadows will occur, they could be preempted to create a permanently unoccluded display.

CHAPTER 5

Conclusions

This work presents a new projector-based display system. It addresses the three problems of geometric calibration, intensity blending and shadow removal faced by front-projection systems. For geometric calibration, two different systems are proposed. The first is a manual approach where a user has direct control over the position of the display. This method, which does not even require a camera, can be used to quickly set up a fully overlapped display that can perform shadow removal. The second method developed relies on locating projected images on a common planar surface using camera images. By combining multiple projectors, this method can be used to create a large display. Its size is maximized by taking advantage of an efficient algorithm to compute the largest area rectangle. When creating a large display, intensity blending is required in regions where projected images overlap to produce a uniformly lit image. For this purpose, a simple rectangular region-based method is presented instead of the pixel-by-pixel method used in existing systems.

Without occlusion detection and active shadow removal, it is still possible to remove shadows passively by illuminating the display with multiple projectors. When used in combination with active shadow removal, the existing method introduced problems by making occlusion detection much more difficult and also making it impossible to determine which projector was being occluded without any guesswork. These problems have been addressed by the introduction of an improved passive shadow removal scheme. By separating the color channels among projectors, the new method simplifies the detection task while allowing the system to determine which projector is occluded in a single step. A closed-loop system for active shadow removal with occlusion detection is also presented. The occlusion detection relies on predicted camera images that are generated at runtime so the system can be used even when the content is created dynamically during use.

The system currently implemented uses two projectors and a camera, all connected to a single computer. To scale this system to multiple computers, efficient communication and synchronization between machines is needed. Towards this goal, polygonal region-based techniques have been introduced to replace current pixel-by-pixel approaches for both intensity blending and shadow removal. It now remains to develop a fully functional implementation of these methods. Also, current shadow removal methods are based entirely on removing occlusion seen in previous frames. Using tracking techniques and eventually 3d reconstruction, future methods should investigate the removal of occlusion before it appears.

REFERENCES

[1] Tomas Akenine-Moller, Tomas Moller, and Eric Haines, *Real-time .rendering*, A. K. Peters, Ltd., 2002.

[2] Helmut Alt, David Hsu, and Jack Snoeyink, *Computing the largest inscribed isothetic rectangle*, Proc. 7th Canad. Conf. Comput. Geom., 1995, pp. 67–72.

[3] M. Ashdown and R. Sukthankar, *Robust calibration of camera-projector system for multi-planar displays*, Technical Report HP Labs HPL-2003-24, 2003.

[4] Jean-Yves Bouguet, *Camera calibration toolbox for matlab*, http://www.vision.caltech.edu/bouguetj/calib_doc/index.html, 2004.

[5] Tat-Jen Cham, Jim Rehg, Rahul Sukthankar, and Gita Sukthankar, *Shadow elimination and occluder light suppression for multi-projector displays*, Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, June 2003.

[6] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti, *Surround-screen projection-based virtual reality: the design and implementation of the cave*, Proceedings of the 20th annual conference on Computer graphics and interactive techniques, ACM Press, 1993, pp. 135–142.

[7] K. Daniels, V. Milenkovic, and D. Roth, *Finding the maximum area axis-parallel rectangle in a polygon*, Proc. 5th Canad. Conf. Comput. Geom., 1993, pp. 322–327.

[8] Herbert Edelsbrunner and Ernst P. Mucke, *Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms*, Symposium on Computational Geometry, 1988, pp. 118–133.

[9] R. I. Hartley and A. Zisserman, *Multiple view geometry in computer vision*, second ed., Cambridge University Press, ISBN: 0521540518, 2004.

[10] 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 **10** (2004), no. 3, 290–301.

[11] Christopher Jaynes, Stephen Webb, R. Matt Steele, Michael Brown, and W. Brent Seales, *Dynamic shadow removal from front projection displays*, Proceedings of the conference on Visualization '01, IEEE Computer Society, 2001, pp. 175–182.

[12] David G. Kirkpatrick and Jack Snoeyink, *Tentative prune-and-search for computing fixed-points with applications to geometric computation.*, Fundam. Inform. **22** (1995), no. 4, 353–370.

[13] A. Majumder and R. Stevens, *Color nonuniformity in projection-based displays: Analysis and solutions*, IEEE Transactions on Visualization and Computer Graphics **10** (2004), no. 2, 177–188.

[14] Aditi Majumder and Rick Stevens, *Lam: luminance attenuation map for photometric uniformity in projection based displays*, Proceedings of the ACM symposium on Virtual reality software and technology, ACM Press, 2002, pp. 147–154.

[15] NEC, *3d reform*, http://www.nec-pj.com/technology/setup.html, 2004.

[16] Ramesh Raskar, Jeroen van Baar, and Jin Xiang Chai, *A low-cost projector mosaic with fast registration*, Proceedings of The 5th Asian Conference on Computer Vision (Melbourne, Australia), 2002.

[17] Matt Steele, Stephen Webb, and Christopher Jaynes, *Monitoring and correction of geometric distortion in projected displays*, Master's thesis, Dept. of Computer Science, University of Kentucky, 2002.

[18] Rahul Sukthankar, Tat-Jen Cham, and Gita Sukthankar, *Dynamic shadow elimination for multi-projector displays*, Proceedings of IEEE Computer Conference on Computer Vision and Pattern Recognition, 2001.

[19] Rahul Sukthankar, Robert Stockton, and Matthew Mullin, *Smarter presentations: Exploiting homography in camera-projector systems*, Proceedings of the International Conference on Computer Vision, 2001.

[20] Jean-Philippe Tardif, Sebastien Roy, and Martin Trudeau, *Multi-projectors for arbitrary surfaces without explicit calibration nor reconstruction*, Proc. IEEE 4th International Conference on 3-D Digital Imaging and Modeling (3DIM'2003) (Banff, Canada), October 2003, pp. 217–224.

[21] G. Toussaint, *Solving geometric problems with the rotating calipers*, Proc. MELECON '83 (Athens), 1983.

[22] Godfried T. Toussaint, *A simple linear algorithm for intersecting convex polygons*, The Visual Computer **1** (1985), no. 2, 118–123.