PixelFlex: A Reconfigurable Multi-Projector Display System

Ruigang Yang*, David Gotz*, Justin Hensley*, Herman Towles*
Department of Computer Science
University of North Carolina at Chapel Hill

Michael S. Brown †
Department of Computer Science
University of Kentucky

Abstract

This paper presents *PixelFlex* – a spatially reconfigurable multi-projector display system. The *PixelFlex* system is composed of ceiling-mounted projectors, each with computer-controlled pan, tilt, zoom and focus; and a camera for closed-loop calibration. Working collectively, these controllable projectors function as a single logical display capable of being easily modified into a variety of spatial formats of differing pixel density, size and shape. New layouts are automatically calibrated within minutes to generate the accurate warping and blending functions needed to produce seamless imagery across planar display surfaces, thus giving the user the flexibility to quickly create, save and restore multiple screen configurations.

Overall, *PixelFlex* provides a new level of automatic reconfigurability and usage, departing from the static, one-size-fits-all design of traditional large format displays. As a front-projection system, *PixelFlex* can be installed in most environments with space constraints and requires little or no post-installation mechanical maintenance because of the closed-loop calibration.

CR Categories: I.3.3 [**Computer Graphics**]: Picture/Image Generation - *Digitizing and scanning, Display algorithms, Viewing algorithms*; I.4.1 [**Image Processing and Computer Vision**]: Digitization and Image Capture - *Imaging geometry, Camera calibration*; B.4.2 [**Input/Output and Data Communications**] Input/Output Devices - *Image display*.

Additional Keywords: large-format projection display, camera-based registration and calibration

1 INTRODUCTION

In recent years, with increased computer performance and the advancement of projector display technology, a number of large-format display systems have been built by research and commercial institutes. The compelling visuals and higher resolutions of these displays make them ideal for a variety of applications in scientific visualization, entertainment, business, and education.

While these systems are very effective at providing large scale imagery to users, installation and operation is often a tedious undertaking. Because of design constraints, most are rear-projection systems that require substantial floor space to accommodate. Moreover, continuous maintenance can

---

*{ryang,gotz,hensley,towles}@cs.unc.edu
†brown@cs.ust.hk (now at the Hong Kong University of Science and Technology)

Figure 1: The top image shows *PixelFlex* with a wide area configuration, while the bottom image shows a stacked configuration.

be required to maintain geometric alignment. Because of the tremendous setup effort needed, once these systems are installed, the spatial layout is finalized.

This papers presents *PixelFlex* - a spatially reconfigurable projector-based display system (shown in Figure 1) which is optimized for planar or nearly planar screens. *PixelFlex* is composed of ceiling-mounted projectors, each with computer-controlled pan, tilt, zoom and focus. A single camera is used for closed-loop calibration. *PixelFlex* allows for installation in small room environments, while providing the flexibility and versatility to change display layout for different users or applications. For example, during group collaboration, a user might desire a wide-area, full wall display. Later, a smaller, brighter, higher pixel density display may be desired. In other cases, a wide-area display using an extra projector to create a high-resolution inset may be needed. In the future, it may be possible to allow two overlapping layers to create a passive stereo system similar to the one described in [2].

This automatic reconfigurability allows users to easily create, save and restore a multitude of display layouts in minutes, literally at the touch of a button.

The remainder of this paper is organzied as follows. Sec-

tion 2 provides an overview of related work. Sections 3, 4 and 5 detail the three major aspects of developing the *PixelFlex* reconfigurable display system including:

- **Components for Reconfigurability**: Physically arranging and controlling the projectors, pan-tilt units and camera for a reconfigurable display involves many details.

- **Automatic Calibration**: Seamless display using multiple, overlapping projected images requires precise geometric registration and photometric calibration, which *PixelFlex* achieves quickly and automatically using a single camera and computer vision techniques.

- **Rendering Applications**: We have developed two applications for the *PixelFlex* system - an X Windows desktop and an OpenGL 3D viewer. These applications also represent two different rendering techniques - a more general two-pass image warping technique for handling non-linear lens distortion over the full optical zoom range of the projectors and a one-pass algorithm applicable when the projector optics (zoom position) is set at the linear "sweetspot".

The paper concludes in Sections 6 and 7 with a discussion of results and conclusions.

2 BACKGROUND AND RELATED WORK

Although newer technologies may lead to larger thin-panel displays and eventually the promise of displays that could be applied like wallpaper, the use of light projectors is currently the most effective way to build large-scale, high-resolution displays. There are several commercially available projector-based large scale display systems, including the well-known CAVE[5, 12] environment, as well as a variety of video walls and dome products [26, 27, 17].

Owning and operating a large format display is often an expensive endeavor, requiring rigid display surface construction with precise projector alignment and constant maintenance. This, compounded with expensive rendering hardware, has limited the use of such systems to only a handful of well-funded research institutes. Addressing these issues, there are on-going research efforts to make large format displays more accessible. Much of this work can be divided into two categories: *distributed rendering* and *geometric registration*.

2.1 Distributed Rendering

Operating systems, such as MacOS, Unix, and MS Windows 2000, have support for extending their desktop to multiple displays. Using this feature with carefully aligned projectors, users can create tiled displays[1]. However, the scalability of such arrangements is often limited by the OS window manager and/or the number of video output channels available on a single machine.

This has led to efforts to explore the use of multiple rendering nodes, often PCs with high-end graphics cards, to create a single logical display. The key challenge with these systems is to get the *distributed* nodes to coordinate as a seamless rendering engine. There are two notable efforts in this area.

First is Li's *et al.* [15, 22, 14] Scalable Display Wall at Princeton. Their research addresses several challenges, including resource allocation, parallel visualization algorithms, and user-interface metaphors for the display ([14] gives a comprehensive overview to this work). Their implementation provides several application support layers, including a Virtual Display Driver to allow Window's applications and an Window's OpenGL implementation.

Second is Humphrey's *et al.* Infomural[11] and WireGL[9] research at Stanford into scalable distributed display architecture. Their effort focuses on efficient algorithms to minimize network load and thus provide efficient scalability. The WireGL software [10] provides an easy to use distributed-OpenGL implementation with available source code which is cross-compatible with several OS platforms.

2.2 Geometric Registration

While distributed rendering research is allowing large scale displays walls to be created from a set of commodity PCs, the construction of these displays is still quite tedious, requiring precise projector overlap, often needing orthogonal projection to the screen. This arguably is the most prominent drawback of large format display design. Research into techniques for automating this registration process is helping alleviate this time consuming setup.

A general solution to the seamless display problem was presented by Raskar *et al.* [19]. In this approach, a series of calibrated stereo cameras are used to determine the display surface and individual projector's intrinsic and extrinsic parameters in a common coordinate frame. The result was an exhaustive description of the entire display environment. Although this approach allowed for a general solution, the computational effort and resources needed to implement this approach introduce their own level of complexity.

Chen *et al.* [3] provides a mechanism to help reduce mechanical alignment by calculating a corrective *projective function* (a 3×3 co-lineation from projector space to the display space) for each projector. These equations are solved by observing corresponding projector pixels and lines via an un-calibrated camera with controllable zoom and focus, mounted on a pan-tilt unit. Simulated annealing is used to find a global solution that minimizes the overall pixel position and line slope error between adjoining projector segments. This approach requires substantial image data and computation. As reported in their paper, data collection and a final solution can take over 30 minutes to compute. Furthermore, this approach corrects the imagery for *slightly misaligned* projectors, it is not clear if it can handle large misalignment, as in our system.

Surati [25] presented a solution that also used a camera to establish the relative geometry of multiple projectors. Using a camera that had been "calibrated" by looking at a regularly spaced grid placed in front of the display surface, subsequent projector imagery can be registered to the grid. Surati's solution was designed for a planar surface; concurrent research [20] showed that this technique could be used to create geometrically correct imagery on arbitrary display surfaces. The only requirement of this technique is that a single camera be able to observe the entire display.

While most registration methods treat the calibration task as a pre-process that is done only once when the configuration is fixed, Yang and Welch [30] presented an on-line and continuous registration method that auto-calibrates the display while the system is being used. Using a camera to observer the entire display, they iteratively refine the estimate

of the display surface shape based on image-based correlation between the known projector image and the observed camera image. While this method takes a relatively long period of time to converge, it can be used to continually refine the display surface geometry even if changes occurs during usage.

2.3 Other Work

The IEEE Computer Graphics and Applications Vol 20, Number 4, 2000 is a special issue on large-format displays, in which the above-mentioned groups and researchers from the University of Chicago, Argonne National Laboratory[8], Lawrence Livermore National Laboratory [23], Sandia National Laboratories [6], and AT&T Shannon Laboratory [28] present their recent research and experiences in building such display systems.

2.4 Spatially Reconfigurable Display

While the current research is promising to make large-scale display more affordable with less rigid design constraints, the display designs themselves are limited to a static, one-size-fits-all design philosophy.

The idea behind *PixelFlex* is to allow for an automated, reconfigurable, large scale display that can change its display layout to accommodate a variety of desired viewing arrangements. In that sense, it is similar to the steerable projection system [4] developed at IBM Research. Their system employs a single projector, while our system uses an array of projectors to create a single logical display device.

3 COMPONENTS FOR RECONFIGURA-BILITY

The *PixelFlex* system currently includes a configuration-control PC, a camera, and eight Proxima 6850 LCD projectors driven by multiple pipes of an SGI Infinite Reality system. The projectors have a resolution of 1024×768 and output 1500 ANSI lumens. In front of each projector, a front-surface mirror is mounted on a pan-tilt unit (PTU). These eight combined assemblies are mounted on the ceiling in two rows of four for front-projection onto a white diffuse wall approximately seven feet away. Front-projection allows the system to fit into much smaller areas by eliminating the need for room behind the display surface.

Both the projectors and the PTUs are connected through serial links to the control PC for computer control of each projector's optic functions (zoom and focus) and mirror orientation (pan and tilt). A picture of the projector array with a closeup view of the mirror-PTU assembly is shown in Figure 2.

An NTSC camera is mounted across the room from the display surface where it can observe the entire display area. It is connected to a video capture card in the control PC.

In our typical wide-area configuration, the display area is approximately 12 feet by 5 feet, with an average spatial resolution of 25 DPI and 15% overlap. Limited by the zoom range, high-resolution insets can increase the spatial resolution to approximately 40 DPI.

From the configuration-control PC, a user interface permits one to change individual projector's optical settings, as well as steer its light via the mirror-PTU unit. Once a desired layout has been created, the layout settings can be saved to a configuration file. These configuration files can

Figure 2: This image shows the *PixelFlex* projector array. The inset is a closeup view of the mirror and PTU in front of a projector.

be loaded via the control panel at a later time to restore the *PixelFlex* system to any saved configuration.

4 AUTOMATIC DISPLAY CALIBRATION

When the user requests a change in the display configuration, the new display layout needs to be *calibrated*. Projector calibration of the PixelFlex system involves accurate computation of the mapping function from projector image coordinates to world display coordinates. We present a simple, yet accurate geometric registration procedure using a single video camera which satisfies the mapping requirements of two different rendering algorithms.

Photometric calibration includes automatically determining the display overlap regions and intensity responses of each projector. This data is used to compute an alpha mask used in the rendering process to attenuate the light contribution of individual projectors in these overlapping regions, thereby producing a more photometrically seamless display.

We also present our method for determining the optimal zoom setting for each projector that minimizes radial distortion. The results from these measurements are used when rendering with our one-pass algorithm.

4.1 Geometric Registration

The goal of the geometric registration procedure is to create a mapping between each projector's image coordinates and the display's global coordinates. There are three main steps in the geometric registration process: (1) camera to display surface registration, (2) projector registration via structured light, and (3) post processing of the registration data for the appropriate rendering algorithm. Due to space limit, we will provide an overview of the process here, more details of the registration process are presented in the technical report [7].

4.1.1 Camera to Display Surface Registration

We use a single, standard NTSC video camera to determine the mapping between the projectors' imagery and the display surface. Our camera needs to see the entire screen and thus uses a wide-field-of-view lens which inevitably suffers from radial distortion. We remove this non-linearity by

computing the lens' distortion factors using Intel's OpenCV computer vision library [13]. Using the undistorted camera image, we register it to the display surface by observing four fiducial placed on the display surface defining the display's rectangular coordinate system. After finding the corresponding positions in the undistorted camera image of the fiducials, we can define a 3×3 homography transformation to transform observed camera points to the global display coordinate system[1]. This procedure needs to be performed only one time unless the camera is moved.

4.1.2 Projector Registration via Structured Light

Once the camera-to-display transform is known, we find a mapping between each projector's pixel coordinate system and the camera's image by taking regular samples in the projector's pixel space and linearly interpolating between them. This is done by projecting a regular array of structured light patterns from each projector and viewing them with the camera. We have found empirically that a 10×10 array of circular features, with Gaussian distribution luminance, is adequate for our current setup. The centroid of the projected features can be determined to sub-pixel accuracy in the camera's image. The sub-sampled data provides a piecewise linear approximation to all pixel space distortions, including keystone distortion, projector lens distortions, and irregular display surface geometry.

The structured-light procedure defines a mapping from each projector's image space to the camera's image space. Using the camera-to-display registration computed in the previous section, the results are mapped into the display's global coordinate system. A similar approach was presented in [29].

4.1.3 Post Processing for Rendering Algorithms

The geometric registration process generates a sub-sampled mapping between projector image coordinates and global display coordinates. We process this initial mapping into the form needed by either of two rendering algorithms.

Optimized One-Pass Algorithm - When the display surface is planar and the projector complies with the pinhole camera model, i.e., lens distortions in the projector are minimal, individual projector imagery can be aligned with a 3×3 homography transformation[18]. This transformation is computed for each projector using samples from the projector-to-display mapping determined during the registration process. This transform corrects the image keystoning caused by off-axis projection. In section 4.3, we define an automatic procedure for finding the optimal zoom setting at which the lens distortion is minimal.

Generalized Two-Pass Algorithm - Our two-pass rendering algorithm is a more general solution that corrects for both linear and non-linear distortions such as non-planar display surfaces and radial lens distortion. The actual non-linear mapping function to pre-warp the projected image is implemented with a piecewise linear approximation. We break the projector image into a tesselated mesh upon which we texture the desired image. The mesh structure and texture mapping coordinates are derived from the mapping data obtained in the structured light registration procedures described in the previous section. In summary, the desired image is first rendered into host memory in the first pass,

---

[1]Chen *et al.* [3] showed that this procedure can be accomplished with an uncalibrated controllable camera mounted on a pan-tilt unit.
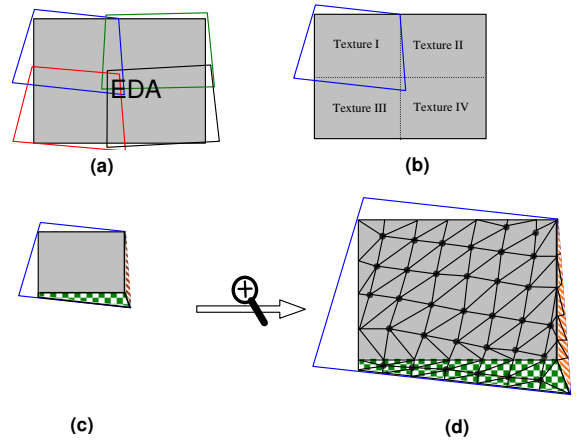

(a)   **(b)**

(c)   **(d)**

Figure 3: Post-processing for a 2×2 projector array. (a) the maximum inscribed rectangle is the Effective Display Area (EDA); (b) the EDA is divided into four texture patches; (c) A projector contains three texture patches; (d) Re-triangulation with normalized texture coordinate.

and texture mapped in a second pass onto this tesselated structure.

Since the projectors are casually aligned, the outer boundary of the unified display is not normally rectangular. So to produce a rectangular display, we determine the maximum inscribed rectangular area. The resultant area defines the Effective Display Area (EDA) on the projection screen. Pixels outside this area are blanked. An example of the EDA is shown in Figure 3(a) for a four-projector array.

The size of the EDA modulates the dimensions of the image rendered into host memory on pass one. In our widescreen format using eight projectors, we would typically render a 4096 x 1536 texture image on the first pass. Texture hardware limitations require us to break this large texture into smaller 1024 x 1024 sub-textures for the second texture rendering step. If a projector must reference multiple sub-textures, then this requires us to also sub-divide the geometric mesh into patches that reference a single sub-texture. This is necessary to insure texture coordinates of a graphic primitive reference only one texture. Figure 3(b-d) illustrates this sub-division for both the geometric tesselated mesh and the texture patch for a four-projector array.

4.2 Photometric Calibration

There are two major tasks involved with photometric calibration. The first task is the measurement of each projector's intensity response. This data is used to create a color lookup table that linearizes each projector's intensity response.

The second task is the determination of the display overlap regions. This data is used to compute an alpha (blending) mask used in the rendering process to attenuate the light contribution of individual projectors in these overlapping regions so as to produce a photometrically seamless display.

To linearize the projector's intensity response, we adopt the techniques introduced by Majumder [16]. We use a spectroradiometer to accurately measure the luminance responses of each channel of a projector. The inverse of each channel's response is loaded into the graphic hardware's
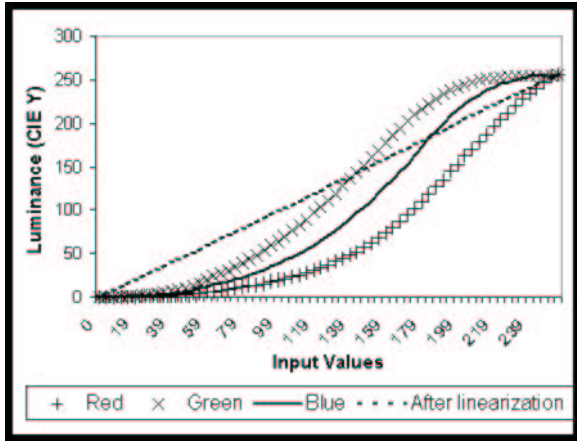
Figure 4: Luminance response of the *PixelFlex* projectors (Proxima 6850) before and after non-linearity correction.

color look-up table for real-time intensity correction. Figure 4 shows the input-output responses before and after non-linearity correction.

Our alpha weighting function is based on the distance formulation presented by Raskar *et al.*[19], however, we add a pixel-density attenuation factor to account for different pixel densities. Following their notations, the final alpha weight $A_m(u,v)$ associated with projector P_m's pixel (u,v) is evaluated as follows:

$$A_m(u,v) = \frac{a_m(m,u,v)p_m^n}{\sum_i a_i(m,u,v)p_i^n} \qquad (1)$$

where a_i is the distance related alpha function computed for projector P_i, p_i is the average pixel density for projector P_i, and n is the constant attenuation factor specified by the user. An attenuation factor of 2 is typical. A larger attenuation factor favors projectors with a higher pixel-density. The modified weight function still guarantees that the alpha values for all projectors sum to unity for each screen point, while allowing more flexible control of the blending smoothness in regions where projectors with different pixel densities overlap.

4.3 Optical Linearity Evaluation

A projector's geometric non-linearity is mainly exhibited in term of lens distortions. During tests of our system, we determined that radial distortion is dependent on the zoom setting of the projector. The lens distortion changes from barrel distortion to pincushion distortion as we change the optical zoom from the narrowest to the widest field-of-view. This implies that there is an optimal zoom setting at which lens distortion is minimal.

Because our one-pass rendering algorithm does not correct for non-linear distortions, we find the zoom setting that minimizes these distortions. Operating the projectors at the optimal zoom setting allows us to efficiently render imagery using our one-pass rendering algorithm that corrects for the geometric distortion created by non-orthogonal projection.

Using linearized images from our system camera, we have developed a line-fitting, computer vision methodology to evaluate a projector's optical non-linearities. The left graph of Figure 5 shows how the optical non-linearity varies as a function of zoom setting for the Proxima 6850 projector.
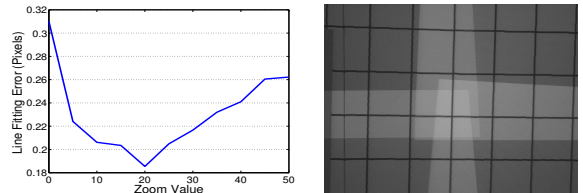

Figure 5: Left: Projector optical non-linearity as a function of zoom setting; Right: Accurate geometric registration is achieved with the one-pass algorithm in this four projector corner region.

The point of minimum non-linearity is close to the middle of the zoom range. At this zoom setting, the spatial distortion of the projected image is usually within a pixel, which allows us to use the more efficient one-pass rendering algorithm while maintaining accurate geometric registration. The right image in Figure 5 shows the level of geometric linearity and registration we can achieve using the one-pass algorithm. Alpha blending has been turned off in this image to clearly identify the four projector overlap in this critical registration region.

5 RENDERING APPLICATIONS

We present two applications for our system - an X Windows Desktop and an OpenGL 3D Viewer. The X Windows Desktop was developed using the two-pass rendering algorithm described in section 4.1.3, and as such supports the full optical reconfigurable range of *PixelFlex*. This desktop allows most existing X Windows applications to run on our system without modification. Unfortunately, it currently does not support OpenGL programs so we separately developed an OpenGL 3D Viewer application. The latter employs our single-pass rendering technique, which is applicable when projector's zoom optics are set at the linear sweetspot.

In both applications, photometric blending of the projector overlap regions is achieved by multiplying the resulting framebuffer image with the pre-computed alpha mask. This can be performed using standard texturing hardware with minimal performance cost. Note that the rendering cost of the second-pass texturing and alpha blending steps are independent of scene complexity.

5.1 X Windows Desktop

The X Windows Desktop supported on *PixelFlex* is provided through a modified version of the Virtual Network Computer(VNC) software from AT&T Laboratory Cambridge [21]. The basic VNC software allows the desktop of one machine to be shared on another machine. The VNC server software also allows an arbitrarily sized, virtual X Windows screen to be created in host memory. We modified the VNC software to create VNC clients that utilize the two-pass rendering algorithm described in section 4.1.3. While inherently intended to run networked on different machines, nothing precludes a VNC server and clients from running on the same platform.

On *PixelFlex*, a VNC X Windows server is first started on the SGI platform. Depending on the desired resolution, one or more rendering clients running on the same SGI machine connect to the server and request updates only for their respective portion of the high-resolution screen. These screen viewports are then pre-warped for display in a second ren-
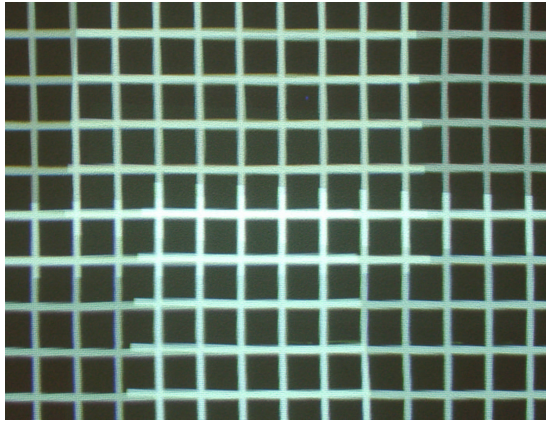
Figure 6: The worse-case registration quality using our two-pass mapping algorithm with photometric-blending off.

dering pass and blended with the alpha mask. This allows X-based applications to be displayed on the *PixelFlex* display without modification, as they would be on any X server.

5.2 OpenGL 3D Viewer

We have implemented an OpenGL 3D viewer to demonstrate the high performance, single-pass rendering algorithm for the case of linear display setup, i.e., planar display surfaces and no radial distortions in the projectors. The mapping from the unified display to the underlying distributed rendering engine are encapsulated in a single C++ class. While not totally transparent, this allows most existing OpenGL applications to be modified to run on *PixelFlex* by changing only a few lines of code.

In this implementation, advantage is taken of SGI's OpenGL Multipipe SDK [24] to hide the complexities of using multiple display pipes on the SGI InfiniteReality 2 platform. The API provides transparent access to accrued hardware resources and allows 3D applications to run with greater flexibility in a multipipe configuration without re-compilation.

6 RESULTS AND FUTURE WORK

6.1 Results

We have assembled a working *PixelFlex* prototype in a conference room at UNC-Chapel Hill, and have successfully used it in a number of configurations.

The default configuration for *PixelFlex* is a wide area configuration with an effective pixel resolution of approximately 3500×1300. This configuration is useful for group collaboration and demonstrations. In Figures 11 and 12, a user is working with the X Windows desktop on *PixelFlex*. The large format display allows him to view a high-resolution image from an astrophysics simulation while checking relevant information on the web using Netscape. Figure 8 shows the underlying triangulated mesh in the left half of the desktop that is used in the 2-pass rendering technique.

Our OpenGL 3D viewer is shown in Figure 9. The rendering is done on an SGI InfiniteReality 2 system using two graphics pipes with four output channels each. It can be easily integrated into existing OpenGL applications without degrading the rendering performance.

Figure 7 shows the OpenGL 3D Viewer running while *PixelFlex* is in a stacked configuration. In the future, with the proper light polarizing filters, a similar stacked configuration could be used for stereoscopic viewing.

Figure 10 shows a view of the X Windows desktop with a high-resolution inset. The smaller picture in the lower-left part of the figure shows a zoomed-in view of the physical high-resolution inset. In this example, the inset allows viewers to see the micro-print on a twenty dollar bill [2] more clearly. This example demonstrates how user control over pixel placement can benefit visualization tasks.

While a complete quantitative measurement of *PixelFlex* is an on-going project in our group, we presented a qualitative view of the accuracy of the geometric registration algorithm. Figure 6 shows the worse-case registration quality using our two-pass mapping algorithm, while the right image of Figure 5 shows the registration in the one-pass rendering case. Both images are taken of a four corners region of four overlapping projectors. Photometric blending was turned off to more clearly identify the overlap region. If projector-to-projector registration were ideal, one would see no C^0 or C^1 discontinuities in the grid pattern of these images. In most cases our positional error over the entire display was less than one pixel; but as seen in the worse-case image, registration errors of up to two pixels sometimes occur. With photometric blending turned on, these differences are less obvious.

6.2 Future Work

While *PixelFlex* represents a large step towards a flexible, spatially dynamic display system, there are a number of issues that need to be addressed in future work.

While the geometric registration we have been able to achieve is very good, we must continue our efforts to fully quantify the registration quality. In this area, we plan to explore the benefits a higher resolution calibration camera may afford. Related to this is a deeper understanding of overall screen resolution especially in the overlap regions where non-coincident pixel centers from multiple projectors complicate the sampling structure.

The most visually concerning issue is that of photometric uniformity. Our system incorporates only the basic notion of blending across projector intensity values in overlapping regions. Overcoming large color differences between projectors and even within a single projector remain a major challenge as does the more fundamental issue of matching black level across the entire display. Reconfigurability further complicates the photometric issues because of changes in brightness, pixel density, and overlap regions. Photometric uniformity is an ongoing research topic in our group. While much of this involves precise measurements with an expensive spectroradiometer, one goal of this research is to determine how low-cost color cameras can be used in the photometric correction process.

We also recognize the need to migrate our system away from a single SGI host and towards a low cost PC-cluster architecture. The most significant challenge in this direction is the development of a software architecture that will support a unified display API. We would like to extend WireGL for our rendering architecture so as to provide truly transparent OpenGL application support; and develop a distributed X Windows desktop architecture that is more efficient than our current VNC implementation. These combined efforts will

---

[2]The "Currency Demo" concept originated with the Scalable Display Wall team at Princeton University.

lead to a more affordable and easier to use display system that transparently supports both desktop and full-screen 3D applications.

7 CONCLUSIONS

In this paper we have presented techniques to build *PixelFlex* – a reconfigurable multi-projector display system. The main idea is to build a large scale display system that allows *the end-users*, not the system designer, to change the display layout on an application driven basis. To realize this vision, we have

- developed a projector array whose display layout can be easily reconfigured

- automated the calibration/registration process using computer vision techniques via closed-loop camera operation

- developed two applications that transparently map underlying *PixelFlex* components to a seamless, geometrically correct, unified display the end-user sees

Combining these techniques, we believe we have built and demonstrated the first large scale display system that has the ability to *automatically reconfigure* for different users or applications. It is a versatile front-projection display system that 1) can be installed in a variety of rooms, 2) is easy for end users and application developers to operate, and 3) requires little maintenance.

8 ACKNOWLEDGMENTS

References

[1] G. Bishop and G. Welch. Working in the Office of "Real Soon Now". *IEEE Computer Graphics and Applications*, 20(4):76–78, 2000.

[2] W. Chen, H. Towles, L. Nyland, G. Welch, and H. Fuchs. Toward a Compelling Sensation of Telepresence: Demonstrating a portal to a distant (static) office. In *IEEE Visualization 2000*, Salt Lake City, UT, USA, October 2000.

[3] Y. Chen, H. Chen, D. W. Clark, Z. Liu, G. Wallace, and K. Li. Automatic Alignment of High-Resolution Multi-Projector Displays Using An Un-Calibrated Camera. In *IEEE Visualization 2000*, Salt Lake City, UT, October 8-13 2000.

[4] Pinhanez Claudio. Using a Steerable Projector and a Camera to Transform Surfaces into Interactive Displays. In *ACM Conference on Human Factors in Computing Systems (CHI 2001)*, pages 369–370, March 2001.

[5] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. *Computer Graphics*, 27(Annual Conference Series):135–142, 1993.

[6] J. A. Friesen and T. D. Tarman. Remote High-Performance Visualization and Collabortation. *IEEE Computer Graphics and Applications*, 20(4):45–49, 2000.

[7] D. Gotz. The Design and Implementation of PixelFlex: A Reconfigurable Multi-Projector Display System. Technical Report TR01-025, University of North Carolina at Chapel Hill, 2001.

[8] M. Hereld, I. R. Judson, and R. L. Stevens. Introduction to Building Projection-based Tiled Display Systems. *IEEE Computer Graphics and Applications*, 20(4):22–28, 2000.

[9] G. Humphreys, I. Buck, M. Eldridge, and P. Hanrahan. Distributed Rendering for Scalable Displays. In *IEEE Supercomputing 2000*, Dallas, TX, Nov. 4-10 2000.

[10] G. Humphreys, M Eldridge, Ian B., G Stoll, M Everett, and P Hanrahan. WireGL: A Scalable Graphics System for Clusters. In *Proceedings of SIGGRAPH 2001*, August 2001.

[11] G. Humphreys and P. Hanrahan. A Distributed Graphics System for Large Tiled Displays. In *IEEE Visualization 1999*, San Francisco, October 1999.

[12] Fakespace Systems Inc. http://www.fakespacesystems.com/index.html.

[13] Intel. Open Source Computer Vision Library (OpenCV), http://www.intel.com/research/mrl/research/opencv/.

[14] K. Li, H. Chen, Y. Chen, D. W. Clark, P. Cook, S. Damianakis, G. Essl, A. Finkelstein, T. Funkhouser, T. Housel, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, P. J. Singh, G. Tzanetakis, and J. Zheng. Early Experiences and Challenges in Building and Using A Scalable Display Wall System. *IEEE Computer Graphics and Applications*, 20(29–37):671–680, 2000.

[15] K. Li and Y. Chen. Optical Blending for Multi-Projector Display Wall System. In *Proceedings of the 12 th Lasers and Electro-Optics Society 1999 Annual Meeting*, November 1999.

[16] A. Majumder, Z. He, H. Towles, and G. Welch. Color Calibration of Projectors for Large Tiled Displays. In *IEEE Visualization 2000*, Salt Lake City, UT, USA, October 2000.

[17] The University of Minnesota. Power Wall. http://www.lcse.umn.edu/research/powerwall/ power-wall.html.

[18] R. Raskar. Immersive Planar Display using Roughly Aligned Projectors. In *IEEE VR 2000*, New Brunswich, NJ, USA, March 2000.

[19] R. Raskar, M. S. Brown, R. Yang, W. Chen, G. Welch, H. Towles, B. Seales, and H. Fuchs. Multi-Projector Displays Using Camera-Based Registration. In *IEEE Visualization*, pages 161–168, San Francisco, October 1999.

[20] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. *Computer Graphics*, 32(Annual Conference Series):179–188, 1998.

[21] T. Richardson, Q. Stafford-Fraser, K. R. Wood, and A. Hopper. Virtual Network Computing. *IEEE Internet Computing*, 2(1):33–38, 1998.

[22] R. Samanta, J. Zheng, T. Funkhouser, K. Li, and J. P. Singh. Load Balancing for Multi-Projector Rendering Systems. In *SIGGRAPH/Eurographics Workshop on Graphics Hardware*, August 1999.

[23] D. R. Schikore, R. A. Fischer, R. Frank, R. Gaunt, J. Hobson, and B. Whitlock. High-Resolution Multi-Projector Display Walls. *IEEE Computer Graphics and Applications*, 20(4):38–44, 2000.

[24] SGI. OpenGL Multipipe SDK 1.0, http://www.sgi.com/software/multipipe/sdk.

[25] R. Surati. *Scalable Self-Calibration Display Technology for Seamless Large-Scale Displays*. PhD thesis, Massachusetts Institute of Technology, 1999.

[26] Trimension Systems. http://www.trimension-inc.com.

[27] Panoram Technologies. http://www.panoramtech.com.

[28] B. Wei, C. Silva, E. Koutsofios, S. Krishnan, and S. North. Visualization Research with Large Displays. *IEEE Computer Graphics and Applications*, 20(4):50–54, 2000.

[29] R. Yang, M. S. Brown, W.B. Seales, and H. Fuchs. Geometrically Correct Imagery for Teleconferencing. In *ACM Multimedia 99*, Orlando, FL, Nov 1999.

[30] R. Yang and G. Welch. Automatic Projector Display Surface Estimation Using Every-Day Imagery. In *9th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, Plzen, Czech Republic, February 2001.

PixelFlex: A Reconfigurable Multi-Projector Display System
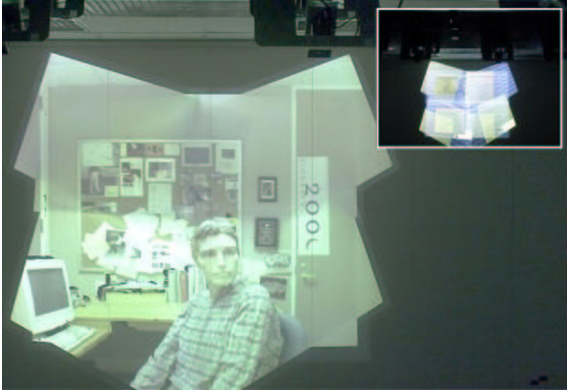

Figure 7: 3D model viewer running in a stacked configuration of *PixelFlex*. Insert image shows projector overlap.
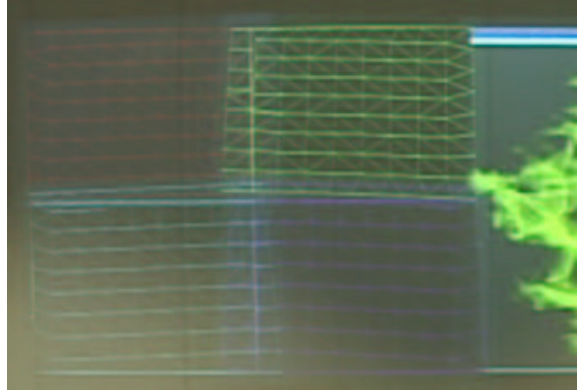

Figure 8: This image reveals the underlying projector meshes used in the X Window desktop application. Triangles in the mesh are textured with the appropriate X Window framebuffer to create a seamless, geometrically correct image.
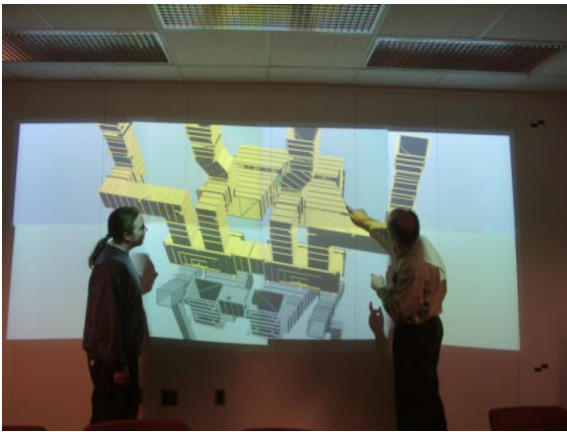

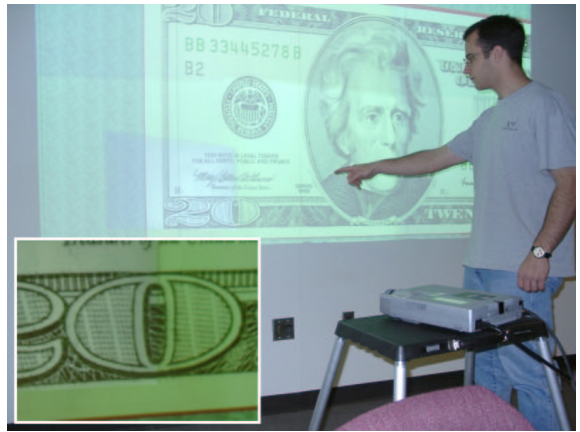Figure 9: 3D viewer application showing portions of a power plant model.

Figure 10: User points to a high-resolution inset area in *PixelFlex* display. The inserted image shows a closed-up view of the high-resolution detail.
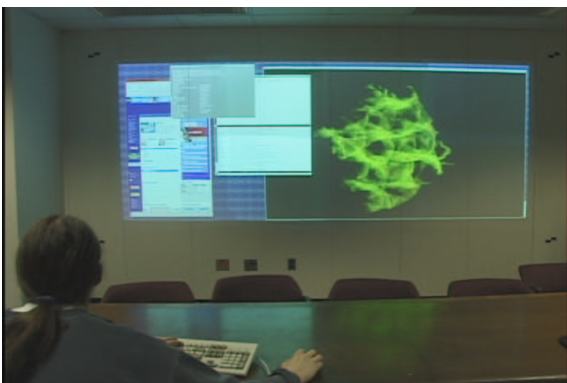

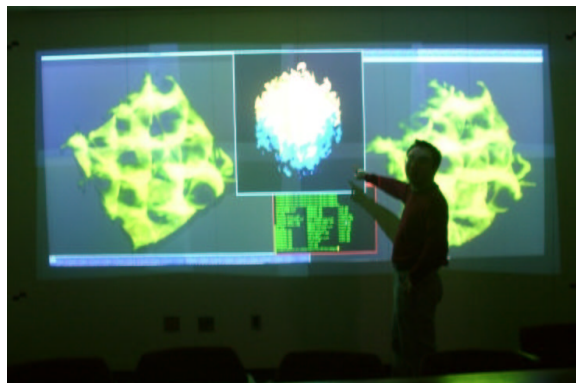Figure 11: A user running the system with a wide-area configuration.

Figure 12: A user using *PixelFlex* to view DOE ASCI/Flash data.